

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

Christopher Hansman
Imperial College London

Jonas Hjort
Columbia University
& BREAD & NBER

Gianmarco León
Universitat Pompeu Fabra
& Barcelona GSE & IPEG

Matthieu Teachout
Columbia University

October 11, 2017

Abstract

We study the relationship between exporters' organizational structure and output quality. If only input *quantity* is observable, theory predicts that vertical integration may be necessary to incentivize suppliers to increase input *quality*. Using data on suppliers' behavior, supplier ownership, supply transactions, and manufacturers' output by quality grade and exports from the Peruvian fishmeal industry, we show the following. After integrating with the plant being supplied and losing access to alternative pay-per-kilo buyers, suppliers take more quality-increasing and less quantity-increasing actions. Integration consequently causally increases output quality, and manufacturers integrate suppliers when facing high relative demand for high quality grades.

JEL codes: D2, O1

*c.hansman@imperial.ac.uk, hjort@columbia.edu, gianmarco.leon@upf.edu, mrt2138@columbia.edu. We thank Charles Angelucci, David Atkin, Wouter Dessein, Dave Donaldson, Bob Gibbons, Marina Halac, Amit Khandelwal, Rocco Macchiavello, Ameet Morjaria, Andrea Prat, John Van Reenen, Eric Verhoogen, Xiao Yu Wang, Daniel Wolfenzon, and Chris Woodruff for very helpful comments and suggestions; seminar participants at Columbia, LSE, MIT, UPF, the NBER Working Group on Organizational Economics, the 2015 Summer School in Development Economics (SSDEV), and the 2016 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 and Cesar Casahuamán 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 (SEV-2015-0563) and grant ECO2014-55555-P.

1 Introduction

Why do so many of our economic transactions occur within firm boundaries (Antràs, 2003; Lafontaine & Slade, 2007)? Research on this question goes back to Ronald Coase’s seminal contributions. The contracting issues emphasized in the body of theoretical work following Coase (1937) are especially prominent in environments where firms attempt to source *high quality* inputs (see Woodruff, 2002; Gibbons, 2005). Because producing high quality *output* typically requires high quality inputs, this suggests that quality upgrading—an essential element of export-driven economic development¹—may be an overlooked motivation for vertical integration.

Designing the incentives of a supplier who faces a tradeoff between producing inputs of high quality or in high volumes is a quintessential challenge for firms that helped motivate Holmstrom & Milgrom (1991)’s seminal work on multitasking (see also Holmstrom & Tirole, 1991; Holmstrom, 1999). They elegantly demonstrate how attaching high powered incentives to input *quantity*—which is typically easier to measure—may *disincentivize* quality-increasing supplier actions. As a result, weakening incentives over quantity may be necessary to ensure that suppliers do not neglect quality. Theories of the firm have since noted that, in a variety of situations, the best or only way to do so may be to bring the suppliers inside the firm (Holmstrom & Milgrom, 1994; Baker *et al.*, 2002; Gibbons, 2005). However, direct evidence on the empirical relevance of Holmstrom & Milgrom’s classical ideas—and their implications for quality-oriented firms’ organizational structure—has remained elusive.

In this paper, we test the hypothesis that (a) integrating suppliers enables firms to incentivize actions that increase input quality, and (b) that vertical integration is therefore a strategy firms use to produce output of high quality. The context we study, Peruvian fishmeal manufacturing², provides an unusual opportunity to do so. This is both because of characteristics of the industry—independent and integrated suppliers deliver inputs of hard-to-observe quality that are converted into a vertically differentiated product—and because unique data allows us to directly observe the entire chain of production, from supplier actions to output quality grades. Testing (a) requires identifying how being integrated affects a supplier’s behavior. Testing (b) requires identifying *both* how vertical integration (X) affects output quality (Y) ($X \rightarrow Y$) and how (demand for) output quality affects vertical integration ($X \leftarrow Y$).³

We begin by presenting a simple theoretical framework that follows prominent theories of the firm in explaining how integration may change a supplier’s incentives to take quantity- and quality-increasing actions. Guided by the framework’s predictions, we estimate how integration affects supplier behavior, focusing particularly on “switchers”—suppliers who supply the same plant before and after being integrated (or sold). We then investigate if integration causally affects output quality, quantifying integration as the number of suppliers owned or, alternatively, the fraction of inputs that are sourced from integrated suppliers. We implement an IV strategy that exploits plausibly exogenous variation (across space and time) in the local presence of independent suppliers: a scarcity of independent options mechanically shifts plants

¹An influential existing literature has shown that access to high quality inputs can help downstream manufacturers in poor countries upgrade the quality of their products and export to more profitable, richer countries (see, among others, Hallak, 2006; Verhoogen, 2008; Goldberg *et al.*, 2010; Brambilla *et al.*, 2012; Kugler & Verhoogen, 2012; Manova & Zhang, 2012; Bastos *et al.*, 2016).

²Fishmeal 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).

³The $X \leftarrow Y$ direction of causality may confound evidence (otherwise) supporting $X \rightarrow Y$ (and vice versa) if the two are not separately identified. In general, the data and variation needed to identify *both* the effectiveness of a firm strategy and the determinants of its use are rarely available. Most existing studies therefore restrict attention to either the former (see e.g. McKenzie *et al.*, 2008; Bloom *et al.*, 2013; Hardy & McCasland, 2015) or the latter (see e.g. Bastos *et al.*, 2016; Park *et al.*, 2009; Brambilla *et al.*, 2012).

toward sourcing a higher share of their inputs from integrated suppliers. Finally, we ask if quality motives are in fact an important determinant of organizational structure. We instrument for an individual firm’s output quality using foreign, quality specific, demand shocks to quantify the extent to which firms vertically integrate when demand shifts towards higher quality grades. Figure 1 gives a visual overview of the paper.

Peru’s fishmeal sector is an ideal setting to study the relationship between organizational structure, supplier behavior, and output quality both because unique data is available on every stage of the production chain, and because of two characteristics of the industry. First, fishmeal manufacturers face textbook contracting challenges: the quality of the product’s primary input—fish—is difficult to observe, and the presence of outside options—other fishmeal manufacturers who may value input quality less—complicates controlling the incentives of an independent supplier.

Second, changes in firms’ structure (whether the suppliers used are integrated or independent), the quality of the output produced, and foreign demand for quality all occur with high frequency during the period we study. This allows us to exploit independent-to-integrated and integrated-to-independent (within-relationship) switchers for identification, and to provide what to our knowledge is the first evidence on the extent to which firms change their organizational structure in order to produce goods of a targeted type.

Three unusual types of data make our analysis possible. We directly observe the behavior of suppliers. Fishing boats are required to transmit GPS signals to the regulatory authorities while at sea. We use GPS-based proxies for input quality- and quantity-increasing actions to explore how a supplier’s behavior differs when integrated.

Additionally, we directly observe output quality. Firms (each of which consists of multiple plants) are required to report each plant’s monthly production of fishmeal in the “prime” (high) quality and the “fair average” (low) quality range. Furthermore, because we observe the price prevailing across the full, granular range of quality in a given week×export port in auxiliary data, we can also infer a continuous measure of each export shipment’s precise quality grade. With these measures we avoid relying on the quality proxies used in the existing literature, which risk conflating quality with mark-ups and horizontal differentiation (see [Atkin et al. \(2017\)](#) for an exception and e.g. [Goldberg & Pavcnik \(2007\)](#); [Khandelwal \(2010\)](#); [Hallak & Schott \(2011\)](#) for discussion of these risks).

Finally, we observe all transactions between plants and suppliers in data from regulatory authorities. Domestic transactions data are rare; observing the transactions of an entire industry, including within-firm transactions, is especially unusual. In combination with plant production and customs data, the supply transactions data allow us to track the flow of goods across three different stages of a global value chain.

Our results are as follows. A *given supplier supplying a given plant* delivers lower total quantities, but higher quality inputs (fresher fish), when integrated with the plant. The change in behavior when a supplier is acquired or sold by the owner of the plant being supplied thus provides empirical support for the intuition captured in our model, namely that vertical integration weakens suppliers’ incentives to produce quantity in ways that benefit quality.⁴

We then show that vertical integration increases the quality of a firm’s output. The relationship holds

⁴We also test whether it is integration per se that influences the behavior of a supplier as opposed to the *repeated interactions* that result from integration. To do so we analyze the behavior of independent suppliers that have long-lasting (quasi exclusive) relationships with a specific firm. In contrast to the change in behavior when a supplier is acquired or sold, independent suppliers do not deliver inputs of higher quality if they become engaged in repeated interactions with the plant supplied (see Appendix B for more details). Note also that our results indicate that what independent suppliers in the Peruvian fishmeal manufacturing industry lack is not the *knowledge* necessary to produce high quality inputs, but rather the *incentives* to do so. We show, for example, that a given supplier produces high quality inputs only when supplying its owner firm, and not when owned by one firm but supplying another.

whether we measure integration as the number of suppliers a firm owns or as the share of inputs that are sourced from integrated suppliers (“Share VI”) at the time of production, and whether we define quality based on our directly observed measures or on unit prices. Furthermore, the relationship between Share VI and output quality also holds across individual plants, both in the industry as a whole and within a large firm that shared internal data with us.⁵

To show that the estimated effect of integration on quality can be interpreted causally, we develop an instrument for a plant’s Share VI based on the local presence of independent suppliers—exploiting the fact that independent boats move up and down the coast of Peru 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 any intentions to produce high quality output—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. As a result, the plant will increase its Share VI. To ensure that the instrument is not contaminated by the strategy of the plant (or firm) in question⁶, we follow a leave-firm-out approach, using independent suppliers serving *other* local firms as a proxy for independent suppliers’ local presence. This approach yields results that are very similar to OLS estimates.

In the final part of the paper, we investigate whether firms in fact change their organizational structure to meet output quality objectives. To do so, we utilize the fact that individual export destinations tend to purchase very specific quality grades, and construct instruments for firm specific demand for high quality output based on shocks to demand from these export destinations. Specifically, we instrument for the quality composition of the demand a firm faces at a given point in time by interacting its initial export destinations with importer countries’ total, leave-firm-out imports from Peru at that later point in time.⁷ The results show that downstream manufacturers acquire more suppliers and thereby increase their Share VI, when faced with greater *relative* demand for high quality output. In other words, a firm’s need to produce quality is a meaningful determinant of its degree of vertical integration.

Different settings entail different contracting environments (see e.g. [Kosová et al. , 2013](#)), and the relationship between integration and quality we document will not apply in every context. However, because input quality is so frequently difficult to observe (and hence incentivize), the challenges we describe here are typical of industries producing vertically differentiated output—particularly in settings where contracts are difficult to enforce.⁸ 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. Our findings thus appear to reflect an association between vertical integration and manufacturing output quality that is not unique to Peruvian fishmeal manufacturing.

⁵Since we find the same relationship between Share VI and output quality across plants within one large firm—and also because the instrument in our IV strategy (described in the following paragraph) differentially impacts different plants within the same firm—we can rule out confounds due to firm×season level shocks (e.g. to productivity or demand) whose correlation with organizational structure (X) and output quality (Y) does not reflect a direct ($X \rightarrow Y$) effect. Note also that controlling for supplier characteristics leaves the estimated effect of Share VI essentially unchanged and that suppliers that get integrated (sold) supply the acquiring (selling) firm/plant almost as frequently pre- (post-) integration. We thus identify the effect of suppliers’ integration status itself.

⁶Of course, Peruvian fishmeal manufacturers do have some control over which particular plants their integrated suppliers deliver to on a given day. Studying the endogenous component of the match between suppliers and plants at a given point in time is beyond the scope of this paper: here we instead exploit the exogenous component for identification.

⁷We follow many fruitful applications of such an approach in the trade literature, including [Park et al. \(2009\)](#); [Brambilla et al. \(2012\)](#); [Bastos et al. \(2016\)](#).

⁸In general, 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](#)).

Our study bridges and advances the literatures on organizational structure and quality upgrading. The body of empirical work on the causes and consequences of firms' choice of structure in developing countries, which began in earnest with Woodruff (2002), is small (see also Natividad, 2014; Macchiavello & Miquel-Florensa, 2016).⁹ 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). What distinguishes our paper from existing studies in this literature is that we are able to trace out the entire chain of production: we identify a force that leads firms to change their organizational structure, but also show how induced changes in structure influence supplier behavior and firm performance. Additionally, we focus on a firm objective that is especially important in countries whose exports to rich countries remain limited: producing high quality output.

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

While especially pressing in developing countries, contracting challenges complicate the operations of firms everywhere. This connects our paper with the broader literature on the boundaries of the firm (see Gibbons (2005); Lafontaine & Slade (2007); Bresnahan & Levin (2012) for excellent overviews¹⁰), to which we make three contributions. First, 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 and quantity-oriented actions of a given supplier interacting with a particular firm.¹¹ Second, we show how vertically integrating causally affects output quality, which to our knowledge has not been done before. 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 an exception¹²), and none that we know of from settings where downstream sectors are vertically differentiated but otherwise homogeneous.¹³ Finally, we show that the output quality-enabling function of vertical integration leads firms to acquire suppliers when relative demand for high quality grades rises, connecting this paper with

⁹Macchiavello & Miquel-Florensa (2016)'s innovative study shows how supply assurance motives influence organizational structure (and use of relational contracts) in the Costa Rican coffee industry by relating measures of ex post renegeing 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 during which 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. There is also a broader literature on business groups in developing countries (see e.g. Bertrand *et al.* (2008) and Khanna & Yafeh (2007)).

¹⁰There is also an influential 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).

¹¹In Appendix C we also provide a new form of evidence on how supplier “adaptation” depends on organizational structure by documenting how suppliers adjust their behavior to exogenous variation in production conditions driven by plankton density (see Williamson, 1975, 1985; Tadelis, 2002; Dessein & Santos, 2006; Forbes & Lederman, 2009, 2010; Baker *et al.*, 2011; Barron *et al.*, 2017). In general, evidence on how a given supplier's behavior depends on integration status is rare: Atalay *et al.* (2014) also exploit changes in integration within supplier-firm pairs, but focus on the transfer of intangible knowledge.

¹²Like us, Forbes & Lederman (2010) exploit exogenous drivers of use of integrated suppliers, showing that routes airlines self-manage have fewer delays/cancellations. Other important evidence includes Novak & Stern (2008); Gil (2009); Kosová *et al.* (2013).

¹³In settings where product differentiation is multidimensional, an analysis like ours would be difficult.

a larger body of empirical evidence on the *causes* of organizational form beginning with [Hart et al. \(1997\)](#); [Baker & Hubbard \(2003, 2004\)](#). Existing work convincingly demonstrates how firms, given their 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](#)).¹⁴ We instead study how firms change their organizational *structure*—acquiring and selling existing suppliers (and making corresponding changes in the share of inputs sourced within firm boundaries)—when their incentives change, focusing particularly on changes in incentives that arise because firms’ (quality) *objectives* change.

The rest of the paper is organized as follows. In Section 2 we provide context background, and in Section 3 we present the theoretical framework. We lay out the data we use in Section 4, and analyze how supplier behavior changes with integration in Section 5. In Section 6 we explore the association between output quality and vertical integration, and in Section 7 how firms’ integration choices change with demand for quality. Section 8 concludes.

2 Background on Peru’s Fishmeal Manufacturing Sector

2.1 Sector profile

Fishmeal is a brown powder made by burning or steaming fish (in Peru typically the anchoveta), and is primarily used as feed for agriculture and aquaculture. Peru is the world’s largest exporter of fishmeal, making up around 30 percent of the world’s 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 owned by 37 firms, but the seven largest firms account for approximately 75 percent of total production. There is heterogeneity in processing capacity, technology, and the share of production that is of high quality grade across plants, and firms differ considerably in their average number of export transactions per season, and the size and value of their shipments.¹⁵

Plants receive inputs of raw fish from fishing boats, which may be independent or owned by the firm that owns the plant. There are on average 812 active boats per fishing 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, underscoring the effort necessary to find and catch fish. Changes in installed technology are rare both for boats and plants during our data period. Table 1 shows summary statistics, providing further detail on the sector.

¹⁴[Forbes & Lederman \(2009\)](#) instead take advantage of inherently exogenous, cross-sectional variation—in weather—to instrument for airlines’ decision to integrate a particular route. Our demand results resonate with the findings of [Alfaro et al. \(2016\)](#), who compare industries across countries and time. They show that higher prices for homogeneous final products allow firms to overcome the costs of integrating (see also [Legros & Newman, 2013](#)).

¹⁵As seen in Appendix Figure A1, firm size correlates positively with average quality grade produced, consistent with [Melitz \(2003\)](#) style models in which unobserved firm heterogeneity governs firms’ targeted output quality, other production choices, and size, and changes in demand- or supply-side factors can lead to changes in the targeted output quality (see, among others, [Verhoogen, 2008](#); [Khandelwal, 2010](#); [Baldwin & Harrigan, 2011](#); [Johnson, 2012](#); [Kugler & Verhoogen, 2012](#)). Such a perspective also appears consistent with a bird’s-eye view of the evolution of the Peruvian fishmeal sector during our data period.

2.2 Product differentiation and quality

An important feature of fishmeal is that output quality effectively depends on a single—measurable—dimension: protein content. In Peru, protein content typically ranges between 63 and 68 percent. In addition, batches with protein content above a specified threshold are labeled “prime” quality. Plants report, for example, 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.

Input quality is much more difficult to quantify or measure directly. However, the ultimate protein content of fishmeal depends crucially on characteristics of the inputs: namely the freshness and integrity of the raw fish that boats deliver to plants. Freshness at the time of delivery in turn depends on several 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. Because of the relationship between freshness and output quality, fish is generally processed as soon as possible after being delivered.

In addition to the freshness of inputs, protein content depends on the technology used by plants. After boats deliver fish at a plant’s docking station, the fish is weighed, cleaned, and converted to fishmeal using one of two technologies: steam drying (hereinafter “High technology”), and exposing the fish directly to heat (hereinafter “Low technology”). As seen in Appendix Figure A2, firms and plants that use the High technology produce higher quality fishmeal on average.

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.¹⁶ In theory fishmeal can be stored for up to six or even 12 months, but we find that almost all is sold before the next production season begins, as shown in Appendix Figure A3 and discussed below.

2.3 Organizational structure

On average, 28 percent of the boats that are active in a given season are integrated with a fishmeal firm. Downstream ownership of boats slowly increased during the last decade,¹⁷ and this slow growth in ownership of suppliers occurred in parallel with Peruvian fishmeal firms (also slowly) increasing the share of their output that is of high quality grade. However, there is significant variation across firms around these trends.

A boat’s total catch in a season is governed by a quota, and each boat typically exhausts its quota. As a result, a firm can generally change the *level* of inputs that come from integrated suppliers only by buying or selling boats. Of course, a firm may change the *share* of inputs coming from integrated suppliers (“Share VI”) also by increasing or decreasing its use of independent suppliers. Following the trend in ownership, Share VI increased during our data period—by 2.9 percent from season to season. As shown in Appendix Table A2, 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.

¹⁶Our period of analysis is 2009 to 2016, or the first 14 fishing and fishmeal production seasons after the introduction of “individual, transferable quotas” (ITQs) in the Peruvian anchovy fishery.

¹⁷See Appendix Figure A4. While the long-term trend in downstream ownership of boats has been positive, we also observe some sales of boats from fishmeal firms to independent co-ops or captains in the data. There was a bigger jump in downstream firms acquiring boats with the introduction of ITQs. As discussed above, we focus on the post-ITQ period. Regulations allow only steel boats to be owned by fishmeal firms. Steel boats that are not owned by fishmeal firms are generally owned by co-ops (“armadores”), while wooden boats are generally owned by individuals or families.

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

In Figure 2 we show the fraction of trips boats deliver to various firms and plants. The bottom part of the figure focuses only on "switchers". Suppliers that get integrated or sold deliver to the acquiring/selling firm around 63 percent of the time *when independent* (i.e. before getting acquired or after getting sold), and around 81 percent of the time when integrated. Similarly, switchers deliver to the plant (within the acquiring firm) they deliver to most frequently around 41 percent of the time when independent and around 45 percent of the time when integrated.¹⁸

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.¹⁹

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 at a given point in time independent suppliers are paid a price per metric ton of fish delivered that is essentially fixed: Port \times Date fixed effects explain 99 percent of the price variation across transactions. We are not aware of formal contracts between independent suppliers and firms that specify explicit delivery requirements.

To sum up, three elements of the Peruvian fishmeal setting are particularly salient: (i) input quantity is observable at the time of delivery, but (ii) input quality is not, and (iii) formal contracts appear to be difficult to write. In the next section, we develop a simple model to help rationalize the organizational structure and contracts we see in practice and to understand the relationship between integration and output quality.

¹⁸The 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.

¹⁹The 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 can be captured in 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 world 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—which depends on the care taken when handling the catch, the attention paid to icing and keeping the fish cold, the duration of time since the fish was caught, and many other factors—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 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, there are other firms that typically produce low quality grades, and hence provide a (presumably less quality sensitive) alternative for suppliers to deliver their catch.²⁰

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. However, 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 \(2005\)](#), 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

²⁰A natural question that our model abstracts from is why different firms might want to be producing different quality levels simultaneously. However, with heterogeneity in production costs and in access to markets with different demands, one might imagine a situation in which it is efficient for different firms to specialize.

independent. If the suppliers are integrated, the firms that own them have the right to the good after the actions are taken. If the suppliers are independent, they retain the right to the good. They bargain with the high quality firms over whether to deliver the good or consign it to its alternative use.

We assume that suppliers have two potential actions $\{a_1, a_2\}$, with costs $c(a_1, a_2) = \frac{1}{2}a_1^2 + \frac{1}{2}a_2^2$. 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 $\gamma, \delta \geq 0$.²¹ 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.²² 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 contractible²³, 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:

²¹More 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.

²²This 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.

²³Alternatively, 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).

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

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

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

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

It is 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

This 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 at any point, 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, 2005) 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 impact 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 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,²⁴ 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 for a low quality firm that is aiming to produce quantity. The market already provides strong incentives for quantity, and integration

²⁴We postulate that P —the value of the inputs in their alternative use—is subject to random shocks, orthogonal to any action the supplier takes. As the size of this shock grows, the first best becomes impossible to sustain with an independent supplier.

would only interfere with and lessen the strength of these incentives. Accordingly, quantity focused firms prefer independent suppliers. A similarly formulated model, with the roles of high and low quality firms switched (e.g. $P \gg Q$), provides precisely this result.

Of course, the theoretical role of vertical integration is a contentious topic, and this framing, which 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), is not the only potential model that would produce a relationship between integration and output quality.²⁵ In actuality, integration is a complex organizational change whose causes and consequences likely operate through the mechanisms emphasized in a number of models. 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 for supplier behavior in the next section, we see these alternative theories as complementary to our primary story, rather than contradictory.

The framework presented in this section motivates three empirical predictions that we test in the remainder of the paper. First, and following most directly from the model, the actions of a supplier should change when the supplier is integrated. In particular, if integration is used as a tool to lower suppliers' incentives for quantity, we should observe suppliers that get integrated *reducing* their effort to produce quantity, especially in ways that also benefit quality. Second, if input quality benefits output quality, then the degree to which a firm or plant utilizes integrated suppliers should causally impact output quality. Third, firms' organizational structure should respond to variation in the profitability of producing high quality output. In other words, an increase in relative demand for high quality grade output should lead to more integration.

4 Data

The primary datasets we use are the following:

Plant production. We use administrative data on all plants' production 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 fishmeal as reported by plants is based on a cut-off on the protein content ladder. From these records, we construct each individual plant's "high quality share of production", our main measure of quality at the plant-month level.

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.²⁶ We also use this registry to link the production data to export data. We can do so for almost

²⁵For 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 5. 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 C.

²⁶The 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).

all firms, since the smallest firms use intermediaries to export.

Export transactions. We use detailed data on the universe of fishmeal exports at the transaction level from Peru’s customs authority. We observe the date of the transaction, the export port of transaction, the destination country, the weight of the fishmeal, the value of the transaction, and the exporting firm.

While we do not directly observe the exact protein content of each export shipment, we can approximate this precise quality grade measure unusually well. This is because we observe quality grade-specific fishmeal prices in granular (week×export port×protein content level) data recorded by a fishmeal consulting company. We infer the protein content of a firm’s exports at each point in time by comparing the unit values of export shipments to this price data.

We have no reason to believe that this inferred protein content measure should be systematically biased. First, fishmeal is a vertically differentiated but otherwise homogenous product, so prices at each quality grade should be the same across firms. Second, one of the largest firm in Peru shared its sales records with us (see below) and these records report the exact quality grade of each shipment. We do not find any significant difference between the reported and the inferred quality grade for that firm. Finally, inferred protein content is highly correlated with the “high quality share of production” directly observed for the firm’s plants in production data. In most of our analysis tables below, we also report results using export unit prices as the dependent variable for robustness.

Fishmeal can be stored for a few months but is usually sold before the next fishing season starts (See Appendix Figure A3); hence, a shipment can only be traced back to a specific fishing season and not a specific month of production. 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 and not the plant level.

Internal data from a large firm. One of the largest fishmeal firms in Peru shared its internal records on sales with us. The firm has been operating for more than a decade, and owns many plants along the coast. The sales records are detailed and include information on the shipment’s type of 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 sales records include information on the specific plant that produced a given shipment of fishmeal, so we can attribute a quality grade to a particular plant during a fishing season.

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.²⁷

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 us²⁸.

²⁷Information 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.

²⁸Only about half of the observations in the Supply transactions dataset can be matched to a GPS recording, and the missing GPS observations are not “missing at random”. 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.

5 Organizational Structure and Supplier Behavior

We begin our empirical analysis by examining how a supplier’s behavior changes when the supplier integrates with or separates from the firm being supplied. In particular, we test whether a supplier engages in behavior that increases quantity but is harmful to quality to a lesser extent when integrated, as the model in Section 3 predicts.

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 the supply transactions data described in Section 4, while the second two are constructed from our boat GPS data. The total quantity supplied is a direct measure of actions taken by the supplier to increase quantity. However, this variable also implicitly 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 vertically, 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 crucially on the time between catch and delivery—is paramount for the protein content of fishmeal. Captains must balance traveling further and longer to catch more fish against ensuring freshness. Because all three of these behaviors 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 \times plant fixed effects, and hence compare the behavior of a *specific* supplier within a *specific* relationship—before and after integration (or separation). For an important part of “switches”, the supplier goes from being independently owned to being owned by the fishmeal firm supplied, but there are also cases where the supplier is sold from a downstream firm to an independent owner or from one fishmeal firm to another. In Section 7 we discuss firms’ motivation for integrating suppliers and show that an important one is the composition of demand: fishmeal manufacturers’ integrate more of their suppliers when they face high relative demand for high quality grade fishmeal and vice versa.

As shown in Appendix Table A3, 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. However, 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, our focus on *within* supplier changes in behavior is not likely to be influenced by these attributes. Recall also that we saw in Figure 2 that suppliers that get integrated or sold deliver to the acquiring/selling firm 63 percent of the time *before getting acquired* (or after getting sold): integration typically implies a simple change in the status of an already frequently used supplier.

We estimate the following regression:

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

where B_{ijt} is a measure of the behavior of supplier i , delivering to plant j , on date t . $[VI \times \text{supplies owner firm}]_{ijt}$ 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 behaviors and Supplier \times 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 2 corroborate our predictions. Column 1 of Panel A shows that, when integrated and supplying a parent plant, a boat delivers on average about ten percent lower quantity 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.

Although our strategy examines changes in supplier behavior *within* a supplier \times plant pair, one concern is that plants may simply choose to integrate suppliers who have already begun changing their behaviors—in effect a violation of a parallel trends assumption. In Panel B of Table 2, we address this concern by examining suppliers who are always integrated but are sold from one fishmeal firm to another during our sample period. While not frequent, we do observe integrated suppliers delivering to plants owned by a different firm.²⁹ Furthermore, some of these suppliers are later purchased by these different firms. As a result we are able to examine how an *integrated* boat's behavior changes with respect to a given plant, when it is acquired by (the parent firm of) that plant.³⁰ We find quite similar, even slightly larger, effects as compared to Panel A. These results suggest that a supplier *increases* its focus on delivering quality when delivering to the plant that now owns it, and also *decreases* its focus on delivering quality when delivering to a plant that used to own it.³¹

Additionally, the results appear to be the result of integration itself, as opposed to any long term relationship that happens to coincide 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.

A last potential concern 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. Such a story would raise a conceptual question: if there is no input quality benefit, and integration lowers input quantity, then why integrate at all? In Section 6 we address this concern. We document that plants' *output* quality responds to integrating existing suppliers in exactly the manner we expect if the behaviors we see do indeed improve input quality.

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. Since the *timing* of acquisitions and sales of suppliers is not randomly assigned, we cannot entirely rule out the possibility that firms e.g. acquire existing suppliers exactly at times when the suppliers' behavior would have become more quality-

²⁹ As seen in Figure 2, Panel (a), integrated suppliers deliver to other firms just over 10 percent of the time.

³⁰ To implement, we run the same specification as in Equation 1, but define $[VI \times \text{supplies owner firm}]$ to be equal to one if the supplier is (i) always owned by a fishmeal firm, and (ii) currently delivering to its parent firm.

³¹ These results also indicate that a story in which integration enables knowledge transfer from manufacturers to their suppliers (see e.g. Atalay *et al.*, 2014) is unlikely to be the primary explanation behind the difference in supplier behavior when integrated.

oriented even if they had remained independent. However, the fact that we find evidence consistent with the predictions of our model motivates the empirical investigation of the model’s other predictions in the next two sections. There we make use of explicitly quasi-random variation in the use of integrated suppliers and demand for high quality fishmeal.

6 Vertical Integration and Output Quality

6.1 Estimating how vertical integration affects output quality

In light of the model in Section 3 and the change in behavior when a supplier gets integrated that we documented in Section 5, we expect vertical integration to be an effective strategy for producing high quality output. In this section we show a series of regressions demonstrating a robust relationship between changes over time in firms’ organizational structure—how vertically integrated Peruvian fishmeal manufacturers are—and changes in their output quality. We then show IV evidence indicating that this relationship arises because integration causally increases output quality.

To quantify vertical integration, we consider both the number of suppliers that a firm owns, and the degree to which the firm actually uses those suppliers—the “Share VI” in the firm’s inputs. While the former represents a more traditional notion of organizational structure, the latter is more directly relevant when asking whether integration causally increases output quality. If vertical integration impacts quality because it incentivizes suppliers to deliver higher quality inputs, then it should matter not just if a firm owns its suppliers, but the degree to which it sources inputs from integrated versus independent suppliers at the time of production. In what follows below, we run specifications both at the firm level, which enables us to consider both definitions, and at the plant level, where we focus on “Share VI.”

We estimate regressions of the form:

$$\text{Quality}_{it} = \alpha + \beta_1 \text{VI}_{it} + \beta_2 \text{HighTech}_{it} + \gamma_i + \delta_t + \varepsilon_{it} \quad (2)$$

where Quality_{it} measures the quality of the output produced by firm or plant i during season or month t . As discussed in Section 4, we observe a series of different quality measures corresponding to our different samples. At the firm level, and for a fishmeal firm that shared internal plant level data for us, we measure output quality using protein content and log export unit prices. These two measures are averaged across export transactions at the season level, weighting by quantity. We view them as complementary: unit prices—the measure of output quality used in the existing literature—are important outcomes in and of themselves, but could partially reflect mark-ups and/or within-season fluctuations in the world price of fishmeal. Protein content instead provides a granular measure of output quality itself.³² At the plant level, we observe the monthly share of fishmeal produced that is of “prime” quality, a direct measure of output

³²While not relevant for our plant level measure of quality, our firm level quality measures require matching export transactions to input transactions. One potentially concern is that fishmeal can be stored for several months, and hence that firms might attempt to strategically time their export transactions. However, in practice inventories are small—between +10 and -10 percent of total season production—as seen in Appendix Figure A3. This is likely because 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 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. To deal with this possibility, we exclude from our sample 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.

quality whose interpretation requires no assumptions.

The type of technology the firm or plant uses to convert fish into fishmeal is an important determinant of output quality (see Appendix Figure A2), and one that could plausibly correlate with VI (Acemoglu *et al.*, 2007, 2010). We thus control for installed HighTech_{it}, i.e., steam drying (High) technology.³³ Finally, γ_i is a plant or firm fixed effect, and δ_t is a time period fixed effect (season or month, depending on the sample). We thus estimate within-time period changes in output quality for those firms or plants that see a *change* in VI within a given season or month, relative to other firms/plants that do not see a change in VI at the same point in time. We cluster the standard errors at the firm level.³⁴

6.2 Vertical integration and output quality at the firm level

We begin by directly exploring the relationship between a firm’s organizational structure—as measured by the number of suppliers it owns—and the quality grade of its output. The results, shown in Panel A of Table 3, point towards a strong baseline relationship between owning suppliers and output quality: if the estimates reflect a causal relationship, they imply that a firm that acquires 10 more of its suppliers would produce fishmeal with a 0.21 percentage point higher protein content (about five percent of the approx. 63-68 percent range observed in Peru), commanding a 2 percent higher unit price.

In Panel B of Table 3 we show that, beyond simply ownership of 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 fishmeal a firm were to produce with inputs coming entirely from integrated suppliers would have 1.3 percentage point higher protein content (again on the approx. 63-68 percent range observed in Peru) than fishmeal produced by the same firm with inputs from independent suppliers. Similarly, the results imply that fishmeal produced with inputs from integrated suppliers would command nine percent higher prices. Notice that the results in Panels A and B underscore the value of our measure of protein content, relative to the conventional approach of using (only) unit prices to measure quality. While the sign and magnitude of the coefficient estimated in columns (1)-(2) and (3)-(4) are consistent, there is significantly more precision in regressions with protein content on the left-hand side. This likely reflects the tighter link between protein content and quality, and highlights the drawbacks of relying on unit prices to measure output quality.

The results in Panel B begin to rule out perhaps the most plausible alternative explanation of the relationship documented in Panel A, namely that firm level shocks—to demand or productivity—that affect output quality independently also enable or incentivize a firm to acquire or sell suppliers. For example, a growing firm might both produce higher quality output and acquire more suppliers. Of course, these results do leave open the possibility that the relationship between output quality and concurrent *use* of integrated suppliers at the firm level is due to shocks or changes in firm strategy that *independently* affect both. In the plant level analysis below we consider these concerns in detail.

In Panel C of Table 3, we rule out the possibility that observable supplier or firm characteristics might explain the results found in Panels A and B. We repeat the regressions in Panel B, but now include a series of additional controls. Controlling for the share of inputs coming from steel boats, high capacity boats, and boats with a cooling system (separately or jointly) leaves the magnitude and significance of the coefficient on

³³As discussed in Section 4, at the firm level, HighTech_{it} is equal to the share of installed capacity that is of the high type, while at the plant level, where we observe whether any high technology is installed, HighTech_{it} is instead a dummy variable.

³⁴Clustering the standard errors at the firm level is not possible in the regressions where we used internal data from one firm.

share of inputs coming from VI suppliers essentially unchanged.³⁵ This is expected, as we saw in Figure 2 that firms acquire/sell suppliers that also ex ante/ex post supply most of their inputs to the firm in question, and changes in boat characteristics are rare.

In columns (2) and (4) of Panel C we add a control for the firm’s share of total industry production. This has little impact on the estimated VI supply coefficient. This result provides further evidence against a story in which unobservable shocks induce firms to simultaneously acquire suppliers and increase output quality, without the former directly affecting the latter (as in e.g. Kugler & Verhoogen (2012)). These estimates also suggests that a “foreclosure” story in which buying suppliers helps downstream firms increase their mark-ups by excluding competitors from the market cannot explain our results (see Ordoover *et al.*, 1990; Hortacsu & Syverson, 2007), consistent with the fact that price variation within quality grade is negligible.

In Appendix B, we consider whether the results found here might be the result of an ongoing supplier-firm relationship, rather than ownership per se. This does not appear to be the case, as we do not observe the positive relationship between output quality and “Share relational” (share of inputs coming from suppliers that are independent but engaged in repeated interactions with the firm) that we have established for Share VI. Hence, it appears to be integration itself, not the relationship, that influences output quality. This result is consistent with the predictions of the dynamic version of our model, also shown in Appendix B.

While the results presented here are suggestive of a causal relationship between vertical integration and output quality, we cannot yet entirely rule out alternative explanations. In the next sub-sections we analyze the relationship between output quality and Share VI at the *plant* level, where we directly observe quality grade, potential confounds that operate at firm×time period level can be ruled out, and most importantly, where we can construct an instrument for our measure of vertical integration and thus directly address possible concerns about the causality of our estimates.

6.3 Vertical integration and output quality at the plant level: OLS

If integration increases output quality because integrated suppliers deliver higher quality inputs, then the relationship between Share VI and output quality that we observe at the firm level should hold at more granular levels. In this sub-section, we look explicitly at the interactions between suppliers and individual plants within firms. We first present OLS regressions that are roughly analogous to those at the firm level. We then describe and present an IV strategy that enables us to isolate the causal link between Share VI and quality.

We find the same positive relationship between the share of inputs coming from integrated suppliers and output quality at the plant level as we do at the firm level. In Panel A of Table 4, we use the directly observed, dichotomous measure of output quality discussed in Section 4: a plant’s total production of high quality (and low quality) fishmeal. These quantities are available for all 94 plants in the full sample and reported at the month level.³⁶ We include plant and month fixed effects and thus focus on changes in share VI across months within a given plant. The results in columns 1 and 2 imply that the share of a plant’s

³⁵We define “high capacity” as greater than the 75th percentile. 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.

³⁶Note that running the regressions in Sub-section 6.2 at month level would require an assumption about how firms manage their inventories (for example, first-in-first-out versus first-in-last-out). We instead match export shipments to firms’ ownership of suppliers and supply transactions at season level, avoiding the need to make such assumptions.

output that is of the high quality type would be seven 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, by focusing only on internal data provided to us by a single major firm. In Columns 1 and 2 of Panel B of Table 4, we repeat the analysis from Panel B of Table 3—that is, we use protein content as our measure of quality—but now at the plant rather than the firm level. This is possible for the sample of plants belonging to the fishmeal firm that shared data with us, enabling us to link the firm’s export transactions with the specific plant that produced the fishmeal. We again include plant and season fixed effects, and thus focus on changes in share VI within a production season and within a given plant. The magnitude and significance of the estimates are very similar to those in Panel B of Table 3.

6.4 Vertical integration and output quality at the plant level: IV

Organizational structure may co-vary with output quality without necessarily reflecting a causal relationship. The existing literature points to several pathways through which this could occur—such as demand for high quality output loosening credit constraints that prevent a firm from acquiring suppliers—but these generally operate at firm or firm×time period level. We have documented a positive, statistically significant, and quantitatively consistent association between Share VI and export shipments’ average quality grade at (i) firm and (ii) plant level across production seasons, and (iii) between Share VI and a directly observed output quality measure at plant×month level. Our findings are thus difficult to reconcile with explanations that do not operate at the plant (sub-firm)×month level. The remaining potential alternative to organizational structure directly affecting output quality is that plant specific shocks, for example to productivity³⁷, occur and *independently* affect the quality of a plant’s output and firm managers’ desire or ability to increase the share of the plant’s supply coming from integrated suppliers.

To address this final concern, 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 independent suppliers 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 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 and share of independent suppliers in a port as potential proxies for the relative cost of using such 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 behind this approach is simply that, at times when there happens to be an abundance of independent suppliers in a given area for exogenous reasons (e.g. weather), 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 reach out to request deliveries from independent suppliers. We thus use *leave-firm-out measures* of the presence of independent suppliers in a given port *during a given period*. In particular, our instruments for Share VI are (i) the number of independent suppliers in a given port, excluding those that supply the firm to which the plant in question belongs, and (ii) the ratio of the number of independent

³⁷Another example of a shock that may affect different plants within a firm differently is El Niño, which hit Peru in late 2009.

suppliers to the total number of suppliers in the port, again excluding any suppliers that interact with the plant/firm in question. 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. Crucially, the signs are all 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 specification are presented in columns 3 and 4 of panels A and B in Table 4. In both panels, the IV estimates are of the same sign and general magnitude as the corresponding OLS estimates, only slightly bigger. Three of the four IV estimates are also significant³⁸, suggesting that the share of a plant’s supplies coming from integrated suppliers at the time of production causally increases output quality.

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? The possibility of 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 potential concern.³⁹ However, beyond the presence of independent suppliers, there appears to be no relationship between changes in output quality across different plants within the same port. 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. It is thus clear that the composition of other plants’ suppliers does not correlate with the quality of the fishmeal produced by the plant in question through other channels than input quality, implying that our instrument’s exclusion restriction holds.

We have now established that suppliers take more quality-increasing and fewer quantity-increasing actions when they are integrated with the plant supplied, and that vertical integration consequently causally increases output quality, as our theory predicts. However, we set out not only to explore if integration and quality are related, but also to understand the reasons why quality-oriented firms are organized as they are in equilibrium. To show that quality is a meaningful part of why firms adopt an integrated organizational structure, we must show that firms vertically integrate *in order to* produce high quality output. Our strategy for doing so asks whether a firm changes its structure in response to changes in demand for quality, which requires isolating shifts in demand for high quality output. We turn to this in the next section.

³⁸Remarkably, the instrumented coefficient on Share VI is statistically significant in one of the two regressions where we use internal data from the firm that shared its data with us despite there only being 66 observations in these regressions. In column (4), where we control for installed technology, the estimated coefficient is very similar in magnitude to that found in column (3), but not significant.

³⁹It 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 C) 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.

7 Demand for Output Quality and Vertical Integration

7.1 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.⁴⁰ 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. \(2016\)](#) (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:

$$VI_{it} = \alpha + \beta_1 \text{Quality}_{it} + \gamma_i + \delta_t + \varepsilon_{it} \quad (3)$$

We control for firm and production season fixed effects and cluster the standard errors at firm level as in Section 6. In the first stage, quality grade produced is instrumented by demand shocks from specific destinations as follows:

$$\text{Quality}_{it} = \gamma_i + \delta_t + \sum_j \beta_j (I_{i2008}^j S_{-i,t}^j) + \varepsilon_{it} \quad (4)$$

where j is an export destination country, and $I_{i2008}^j S_{-i,t}^j$ are our excluded instruments. I_{i2008}^j is a dummy variable equal to one if firm i exported to destination j in 2008, the year prior to our analysis period. $S_{-i,t}^j$ 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 $S_{-i,t}^j$ and I_{i2008}^j .

⁴⁰See 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.

7.2 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.⁴¹

The OLS and the second stage IV results are reported in Table 5. 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 5 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 A5 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.⁴² 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 A6⁴³.

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 5 to the Stock-Yogo critical values⁴⁴, 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.* (2016) for a lengthier discussion of this issue in the context of “demand pull” instruments.

⁴¹The 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.

⁴²LASSO (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.

⁴³The sign of the coefficient for each instrument is broadly consistent with the relative average quality imported by each country (See Appendix Table A1).

⁴⁴Though 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 5 and the resulting integration→quality relationship shown in Section 6. 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.

8 Conclusion

Guided by Holmstrom & Milgrom (1991)’s classical ideas and subsequent theories of the firm that incorporate the multitasking aspect of suppliers’ work they emphasized, this paper identifies an overlooked motivation for and consequence of vertical integration in incomplete contracts settings: downstream firms strategically integrate to be able to produce output of high enough quality to sell to high-paying consumers abroad. Integrating existing suppliers allows manufacturing firms to incentivize quality-increasing behavior upstream 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 suppliers’ behavior should change with integration, how integration consequently should affect output quality, and how relative demand for high quality output should affect firms’ choice of organizational structure.

We test these predictions using data from the Peruvian fishmeal manufacturing industry. We show that when fishing boats get integrated with the downstream firm supplied, they change their behavior in a way consistent with an objective of delivering fresher fish—which helps firms produce higher quality fishmeal. Using within- and across-firm transaction level data and direct measures of the quality grades manufacturers produce, we then show that a vertically integrated organizational structure causally increases output quality. Finally, we show that, when firms face high *relative* demand from importers of high quality grades, they acquire more of their existing suppliers.

Overall this paper’s results demonstrate that vertical integration is a specific strategy that Peruvian manufacturing firms can and do adopt in order to upgrade the quality of their products, and why this strategy works. 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.⁴⁵ This suggests that our findings reflect an association between vertical integration and manufacturing output quality that tends to hold on average across countries and manufacturing

⁴⁵We show this in Appendix Table A7. 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 Related party share of U.S. imports $_{cpt} = \beta_{pt} + \delta_c + v_{cpt}$, 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. This regression is estimated using data from the U.S. Census Bureau.

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.

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.
- 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, G., GIBBONS, B., & MURPHY, K. J. 2011. *Relational Adaptation*. mimeo, Harvard Business School.
- BAKER, GEORGE, GIBBONS, ROBERT, & MURPHY, KEVIN J. 2001. Bringing the Market Inside the Firm? *The American Economic Review*, **91**(2), 212–218.
- 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*.
- BARRON, DANIEL, GIBBONS, ROBERT, GIL, RICARD, & MURPHY, KEVIN J. 2017. *Relational Adaptation under Reel Authority*. mimeo, Kellogg.
- BASTOS, PAULO, SILVA, JOANA, & VERHOOGEN, ERIC. 2016. *Export Destinations and Input Prices*. Unpub. paper.
- BERTRAND, MARIANNE, JOHNSON, SIMON, SAMPHANTHARAK, KRISLERT, & SCHOAR, ANTOINETTE. 2008. Mixing Family with Business: A Study of Thai Business Groups and the Families Behind Them. *Journal of Financial Economics*.
- BLOOM, NICHOLAS, EIFERT, BENN, MCKENZIE, DAVID, MAHAJAN, APRAJIT, & ROBERTS, JOHN. 2013. Does management matter: evidence from India. *Quarterly Journal of Economics*.
- 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.
- DESSEIN, WOUTER, & SANTOS, TANO. 2006. Adaptive Organizations. *Journal of Political Economy*.
- 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. 2005. Four formal(izable) theories of the firm? *Journal of Economic Behavior and Organization*, **58**(2), 200 – 245.
- GIL, R. 2009. Revenue sharing distortions and vertical integration in the movie industry. *Journal of Law, Economics, and Organization*.
- 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.
- 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*.
- HARDY, MORGAN, & MCCASLAND, JAMIE. 2015. *Are Small Firms Labor Constrained? Experimental Evidence from Ghana" with Morgan Hardy*. mimeo, University of British Columbia.
- 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 Multi-national 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.
- KHANNA, T, & YAFEH, Y. 2007. Business Groups in Emerging Markets: Paragons or Parasites? *Journal of Economic Perspectives*, XLV, 331–372.
- 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.
- MCKENZIE, D., DE MEL, S., & WOODRUFF, C. 2008. Returns to Capital: Results from a Randomized Experiment. *Quarterly Journal of Economics*.
- MCLAREN, J. 2000. Globalization and Vertical Structure. *American Economic Review*.
- MELITZ, MARC J. 2003. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, 71(6), 1695–1725.

- 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*.
- PARDES, 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.
- 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*.
- TADELIS, STEVEN. 2002. Complexity, Flexibility, and the Make-or-Buy Decision. *American Economic Review*.
- TVETERAS, SIGBJORN, PARDES, 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: SUPPLIER BEHAVIOR AND VERTICAL INTEGRATION

Panel A: Identified from all switchers (Independent 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*** (0.023)	-0.054*** (0.019)	-0.030* (0.016)
Date FEs	Yes	Yes	Yes
Supplier × Plant FEs	Yes	Yes	Yes
N	315,442	137,278	159,724
Panel B: Identified 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*** (0.027)	-0.082*** (0.026)	-0.073*** (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 3: OUTPUT QUALITY AND VERTICALLY INTEGRATED SUPPLIERS

Panel A: Output quality and number of suppliers owned						
Dep. var:	Protein content		Log(unit price)			
	(1)	(2)	(3)	(4)		
Number of suppliers owned	0.021** (0.009)	0.022** (0.010)	0.002 (0.002)	0.002 (0.002)		
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.	65.6	65.6	7.23	7.23		
N	220	220	220	220		

Panel B: Output quality and Share of inputs from VI suppliers						
Dep. var:	Protein content		Log(unit price)			
	(1)	(2)	(3)	(4)		
Share of inputs from VI suppliers	1.080*** (0.266)	1.079*** (0.267)	0.090* (0.047)	0.090* (0.047)		
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.	65.6	65.6	7.23	7.23		
N	220	220	220	220		

Panel C: Output quality and Share of inputs from VI suppliers						
Dep. var:	Protein content			Log(unit price)		
	(1)	(2)	(3)	(4)	(5)	(6)
Share of inputs from VI suppliers	1.056*** (0.335)	1.138*** (0.279)	1.106*** (0.345)	0.101** (0.048)	0.098* (0.048)	0.108** (0.049)
Share of inputs from steel boats	-0.065 (0.525)		-0.026 (0.523)	0.007 (0.044)		0.013 (0.043)
Share of inputs from boats with high capacity	0.180 (0.590)		0.137 (0.595)	-0.115 (0.101)		-0.122 (0.101)
Share of inputs from boats with cooling system	0.142 (0.919)		0.194 (0.941)	0.040 (0.072)		0.048 (0.072)
Share of industry's production		1.711 (2.217)	1.747 (2.207)		0.249 (0.167)	0.258 (0.194)
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.	65.6	65.6	65.6	7.23	7.23	7.23
N	220	220	220	220	220	220

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: OUTPUT QUALITY AND SHARE OF INPUTS FROM VERTICALLY INTEGRATED SUPPLIERS

Panel A: All Plants				
Dep. var:	High Quality Share of Production			
	OLS (1)	OLS (2)	IV (3)	IV (4)
Share of inputs from VI suppliers	0.102** (0.038)	0.064** (0.030)	0.115** (0.051)	0.091** (0.045)
Has high technology	No	Yes	No	Yes
Month FEs	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.85	0.85	0.85	0.85
N	2647	2647	2487	2487
Panel B: Plants Within a Major Firm				
Dep. var:	Protein Content			
	OLS (1)	OLS (2)	IV (3)	IV (4)
Share of inputs from VI suppliers	1.369** (0.654)	1.338** (0.656)	1.469* (0.807)	1.390 (0.918)
Has high technology	No	Yes	No	Yes
Month FEs	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
Mean of Dep. Var.	65.8	65.8	65.8	65.8
N	66	66	66	66

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

TABLE 5: VERTICAL INTEGRATION AND THE OUTPUT QUALITY NECESSARY TO MEET THE DEMAND - INSTRUMENTING OUTPUT QUALITY PRODUCED WITH FIRM-SPECIFIC DEMAND SHOCKS

Panel A				
Dep. var:	Share of inputs from VI suppliers			
	OLS (1)	OLS (2)	IV (3)	IV (4)
Protein content	0.030*** (0.010)	0.028*** (0.009)	0.128** (0.051)	0.141*** (0.054)
Log(Sales)		0.008 (0.015)		-0.042* (0.024)
Season FEs	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes
Kleibergen-Paap LM p-value (Under-id)			0.40	0.38
Kleibergen-Paap Wald F statistic (Weak inst)			11.0	15.8
Anderson-Rubin Wald test p-value			0.000	0.000
Mean of Dep. Var.	0.45	0.45	0.45	0.45
N	220	220	220	220

Panel B				
Dep. var:	Number of Boats			
	OLS (1)	OLS (2)	IV (3)	IV (4)
Protein content	0.507 (0.312)	0.337 (0.308)	2.123** (0.849)	2.366** (1.008)
Log(Sales)		0.822 (0.514)		-0.069 (0.753)
Season FEs	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes
Kleibergen-Paap LM p-value (Under-id)			0.40	0.38
Kleibergen-Paap Wald F statistic (Weak inst)			11.0	15.8
Anderson-Rubin Wald test p-value			0.000	0.000
Mean of Dep. Var.	24.5	24.5	24.5	24.5
N	220	220	220	220

Notes: One observation is a firm during a production season. *Share of inputs from VI suppliers* is the share of a firm's inputs that come from VI suppliers 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. The instruments are interactions of indicators for positive exports in 2008 to each of the top 20 destination countries with leave-firm-out share of Peru's fishmeal exports towards the destination in the relevant year. The first stage is shown in Columns 1 and 2 of Appendix Table A6. In Appendix Table A5, we both include the top 10 destinations in the first stage and use a LASSO approach to chose destinations as robustness checks. Standard errors clustered at the firm level are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE 1: PAPER OVERVIEW

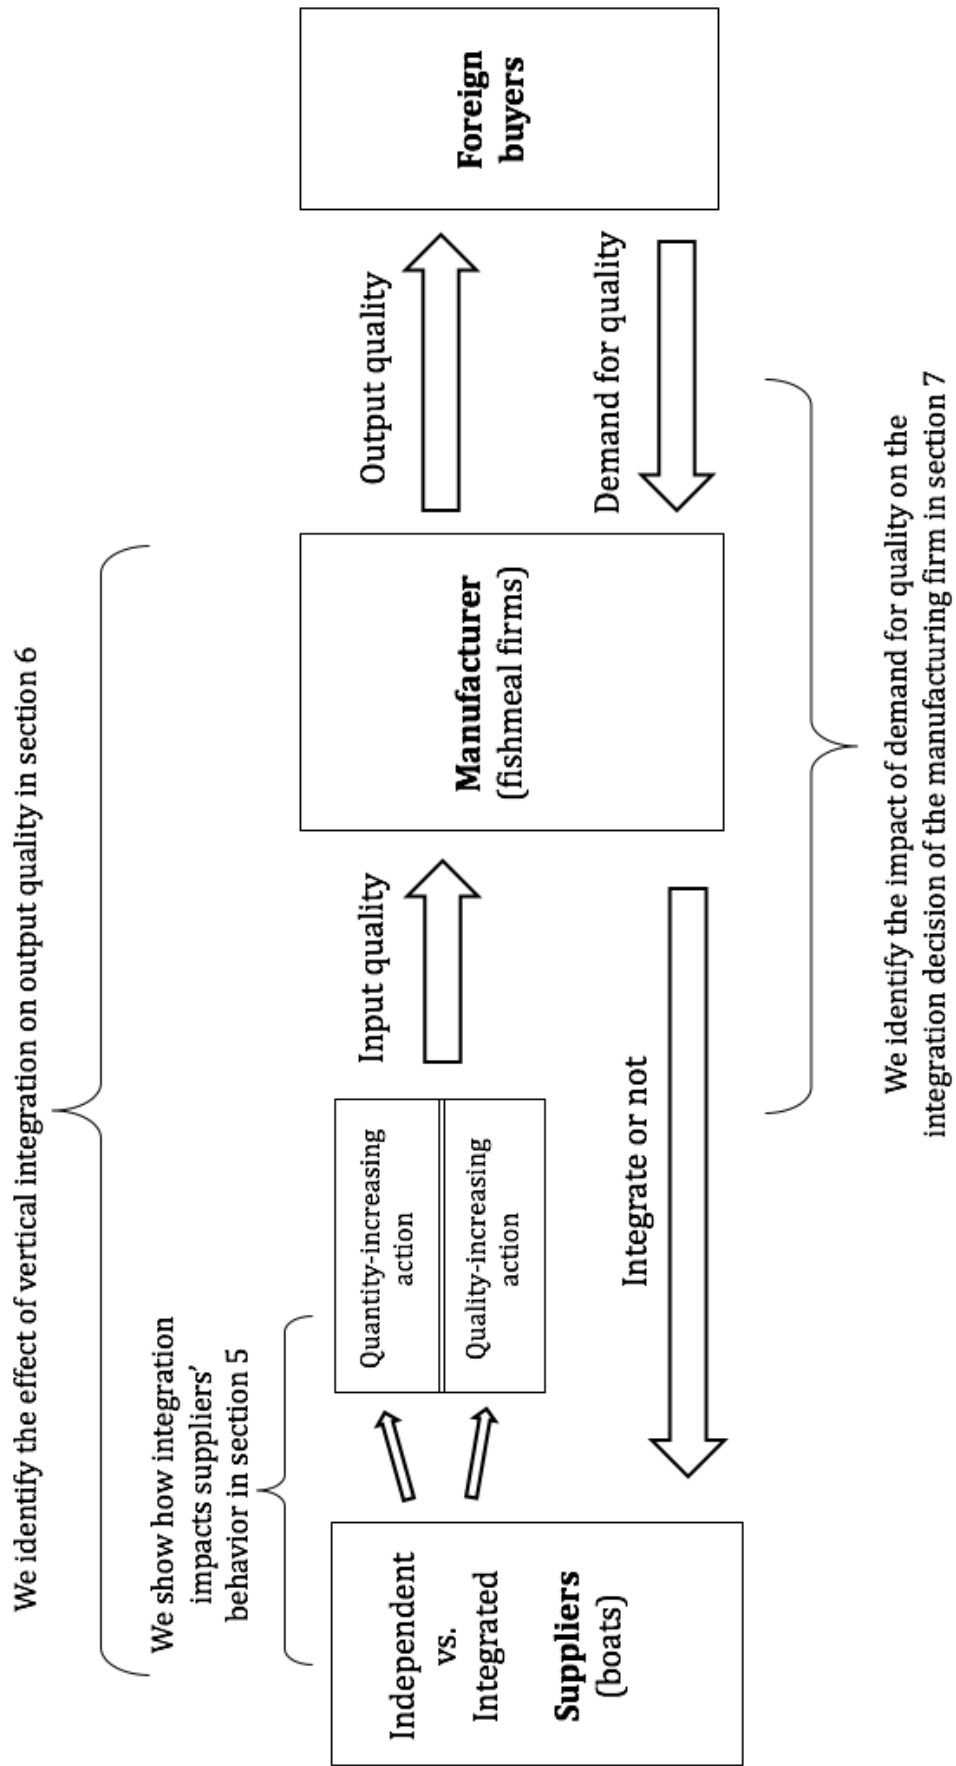
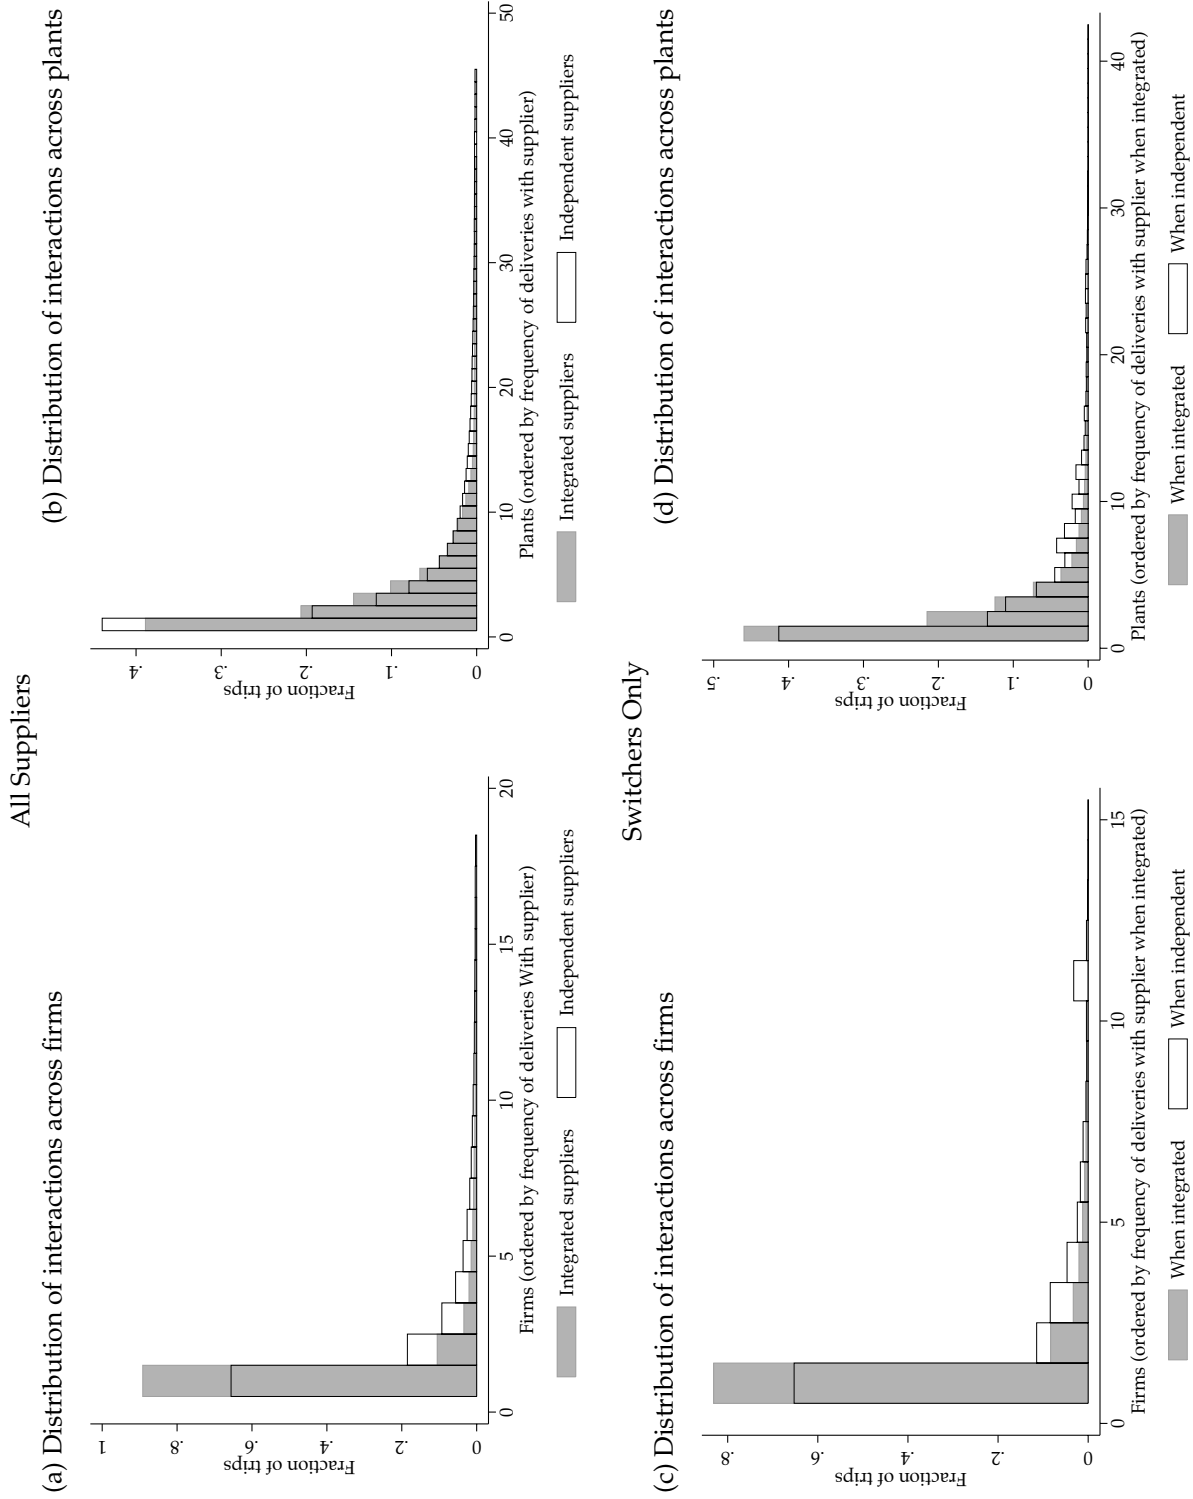
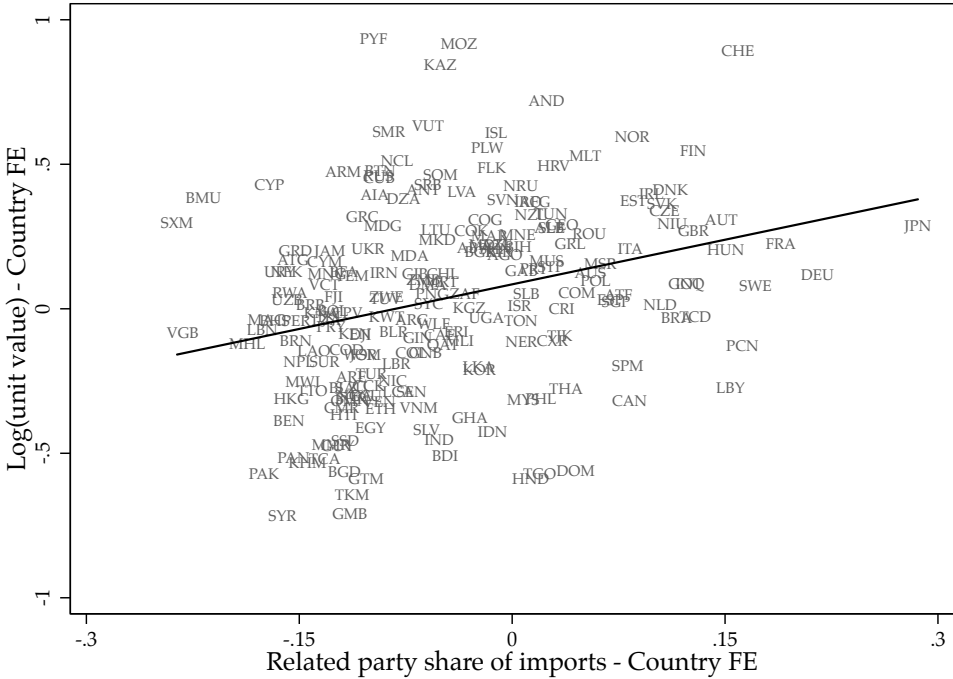


FIGURE 2: INTERACTIONS BETWEEN MANUFACTURERS AND INTEGRATED AND INDEPENDENT SUPPLIERS



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 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 $\log(\text{unit value})_{cpt} = \alpha_{pt} + \gamma_c + \varepsilon_{cpt}$, 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 \times 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 $_{cpt} = \beta_{pt} + \delta_c + u_{cpt}$, 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 \times 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: DECOMPOSITION OF THE GROWTH RATE OF SHARE OF INPUTS FROM VI SUPPLIERS

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

	Total	A	B
Growth	2.9%	2.2%	0.7%
Relative Contribution		77%	23%

Notes: The growth rate of “Share 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. VI_{*i,t*} and 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 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 VI_{*i,t*}”, Term A and Term B.

TABLE A3: 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 (16.24)	0.00 (0.06)	65.73 (27.34)	215.40 (94.78)	56.10 (7.74)
Steel - Independent	104.03 (40.77)	0.09 (0.28)	219.30 (84.35)	412.31 (189.82)	81.15 (13.43)
Steel - Switchers	148.88 (0.43)	0.25 (0.444)	301.18 (129.92)	616.30 (328.51)	92.25 (15.37)
Steel - VI	181.62 (68.13)	0.34 (0.47)	382.00 (137.11)	769.96 (352.52)	97.29 (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 A4: OUTPUT QUALITY AND SHARE OF INPUTS FROM VERTICALLY INTEGRATED SUPPLIERS - FIRST STAGE

Dep. var:	Share of inputs from VI suppliers			
	All Plants		Plants Within a Major Firm	
	(1)	(2)	(3)	(4)
Number of Independent Boats in Port	-0.001*** (0.000)	-0.001*** (0.000)	-0.000 (0.000)	-0.000 (0.000)
Share of Independent Boats in Port	-0.313*** (0.027)	-0.314*** (0.027)	-0.412** (0.200)	-0.398* (0.207)
Kleibergen-Paap LM p-value (Under-id)	0.038	0.038	0.005	0.006
Kleibergen-Paap Wald F statistic (Weak inst)	70.66	68.54	3.61	3.06
Anderson-Rubin Wald test p-value	0.04	0.02	0.24	0.31
Has High Technology	No	Yes	No	Yes
Month FEs	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes

Notes: Results from the first stage of IV specifications reported in Table 4. *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: VERTICALLY INTEGRATED SHARE OF INPUTS AND OUTPUT QUALITY PRODUCED – INSTRUMENTING WITH FIRM-SPECIFIC DEMAND SHOCKS - ROBUSTNESS CHECKS

Panel A				
Dep. var:	Share of inputs from VI suppliers			
	Top 10 Destinations		LASSO	
Protein content	0.123** (0.054)	0.138** (0.054)	0.142** (0.057)	0.150** (0.061)
Log(Sales)		-0.040 (0.029)		-0.046* (0.028)
Kleibergen-Paap LM p-value (Under-id)	0.31	0.40	0.11	0.16
Kleibergen-Paap Wald F statistic (Weak inst)	2.30	3.36	6.89	4.93
Anderson-Rubin Wald test p-value	0.072	0.079	0.406	0.398
Mean of Dep. Var.	0.45	0.45	0.45	0.45
N	220	220	220	220
Season FEs	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes

Notes: One observation is a firm during a production season. *Share of inputs from VI suppliers* is the share of a firm's inputs that come from VI suppliers 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. The instruments are interactions of indicators for positive exports in 2008 to destination countries with leave-firm-out share of fishmeal exports from Peru towards the destination in the relevant year. We both include the top 10 destinations in the first stage and use a Lasso approach to choose destinations as robustness checks. The first stage is shown in Columns 3-6 of Appendix Table A6. Standard errors clustered at the firm level are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A6: VERTICALLY INTEGRATED SHARE OF INPUTS AND OUTPUT QUALITY PRODUCED – INSTRUMENTING WITH FIRM-SPECIFIC DEMAND SHOCKS - FIRST STAGE

Dep. var:	Protein content					
	Top 20 Destinations		Top 10 Destinations		LASSO	
	(1)	(2)	(3)	(4)	(5)	(6)
Indonesia	69.972*** (26.838)	64.735** (27.919)	56.751** (26.951)	54.022** (25.406)	61.494*** (19.927)	57.915*** (18.922)
South Korea	-98.818 (61.087)	-100.589 (71.081)			-99.535* (51.153)	-103.097* (60.673)
China	-2.241 (2.141)	-2.203 (1.999)	-2.098 (2.100)	-2.017 (1.967)		
Germany	0.167 (1.726)	0.027 (1.793)	0.068 (1.711)	0.004 (1.813)		
Japan	-6.886 (7.510)	-7.401 (6.996)	-6.142 (8.252)	-6.950 (7.783)	-5.795 (6.748)	-6.110 (6.377)
Chile	-4.762* (2.603)	-3.738 (2.481)	-5.969*** (2.115)	-5.044** (2.062)	-4.371** (2.130)	-3.765* (2.216)
Vietnam	6.981 (7.471)	4.586 (7.565)	1.046 (7.566)	-0.832 (6.759)		
Taiwan	-12.450 (20.169)	-9.914 (19.243)	-7.373 (19.807)	-4.915 (18.423)		
United Kingdom	-19.492* (10.088)	-14.959 (9.879)	-10.452 (9.887)	-9.866 (7.901)	-19.308** (7.848)	-14.936* (7.839)
Turkey	-4.998 (8.260)	-3.598 (9.060)	-6.995 (6.210)	-6.016 (6.935)		
Spain	-7.679 (12.283)	-3.022 (13.107)	16.311 (15.452)	14.664 (14.445)		
Australia	0.742 (17.303)	1.246 (13.261)				
Canada	-1.905 (23.176)	-8.737 (24.443)				
France	85.177* (48.641)	54.714 (56.565)			78.745** (36.967)	54.342 (42.153)
Italy	24.686 (51.156)	26.760 (52.203)				
Bulgaria	-6.965 (47.673)	-4.527 (50.403)				
Venezuela	20.975 (102.539)	-4.654 (105.883)				
Belgium	76.085 (122.671)	53.787 (127.219)				
Philippines	53.752 (116.560)	48.170 (103.206)			37.282 (108.501)	35.118 (94.188)
India	-115.176 (75.185)	-109.410 (69.163)			-100.153** (46.060)	-101.626** (46.974)
Log(Sales)	No	Yes	No	Yes	No	Yes
Season FEs	Yes	Yes	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes

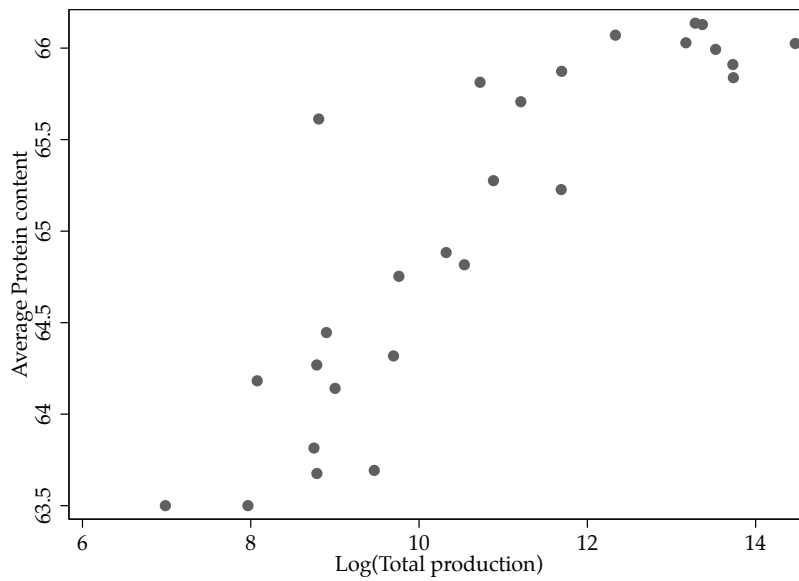
Notes: First stage results for IV specifications reported in Tables 5 and A5. One observation is a firm during a production season. *Share of inputs from VI suppliers* is the share of a firm's inputs that come from VI suppliers during a season. *Protein content* is the quantity weighted average of a measure of quality inferred with a database that provides weekly prices by quality. The instruments are interactions of indicators for positive exports in 2008 to each of the destination countries with leave-firm-out share of fishmeal exports from Peru towards the destination in the relevant year. Each instrument is labeled by the name of each destination country. Standard errors clustered at the firm level are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A7: 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

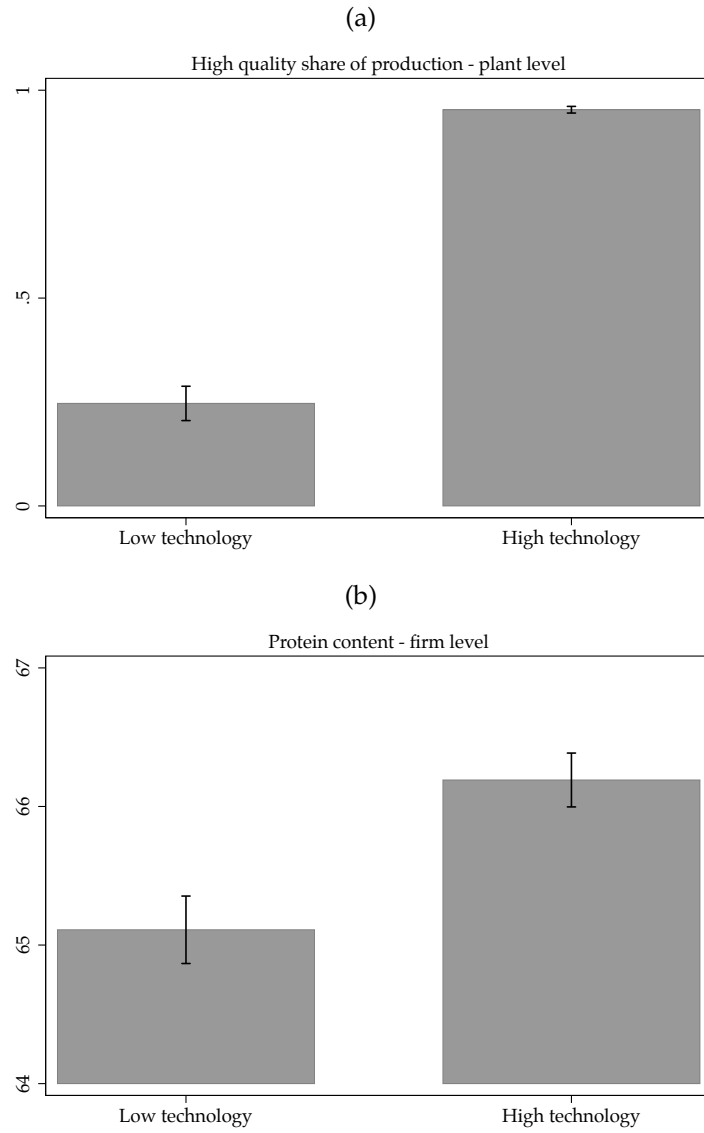
Notes: In this table, the dependent variable is ε_{cpt} from the regression $\log(\text{unit value})_{cpt} = \alpha_{pt} + \gamma_c + \varepsilon_{cpt}$, where $\log(\text{unit value})_{cpt}$ is the average log unit value of products exported from country c , of HS6 code p , in year t to the U.S.; α_{pt} 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 independent variable is v_{cpt} from the regression Related party share of U.S. imports $v_{cpt} = \beta_{pt} + \delta_c + v_{cpt}$, where Related party share of U.S. imports v_{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 v_{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



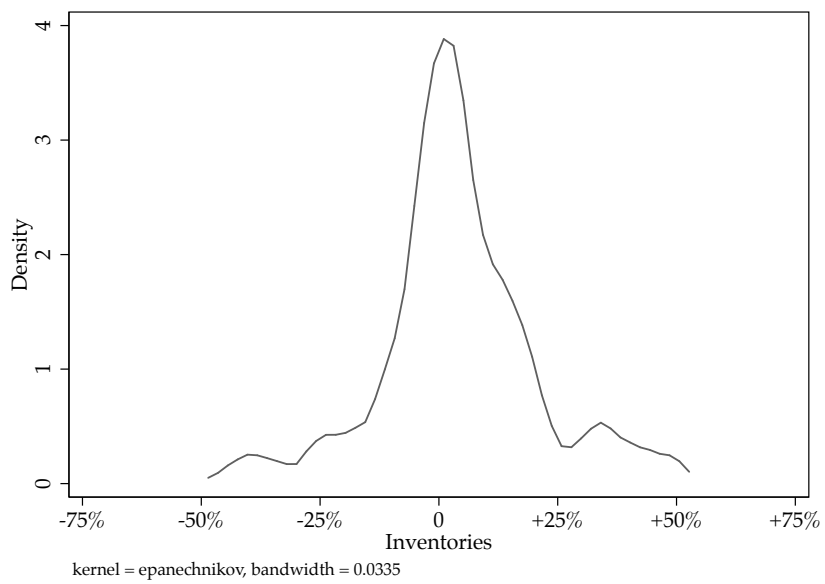
Notes: Each dot represents one fishmeal firm in our sample. Total production is the total weight of fishmeal the firm produced during our data period and average protein content is the quantity weighted average protein content of the firm's fishmeal exports.

FIGURE A2: PLANT TECHNOLOGY AND OUTPUT QUALITY



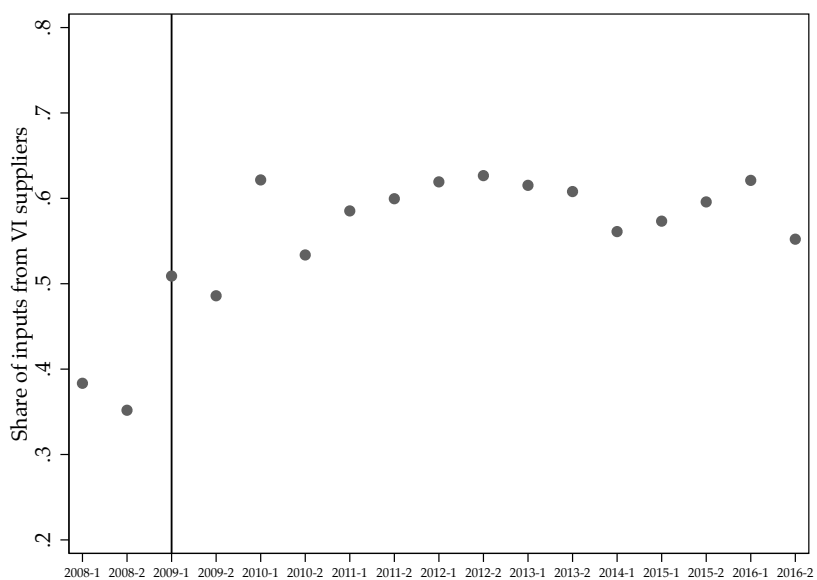
Notes: Panel (a) shows the high quality share of production for plants with high technology (as defined in Section 2) and plants that only have low technology. Panel (b) shows the average protein content (quality grade) for high and low technology firms. High technology firms are firms for which the high technology share of total capacity is above the sample median.

FIGURE A3: 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 A4: EVOLUTION OF THE VERTICALLY INTEGRATED SHARE OF INPUTS INDUSTRY-WIDE



Notes: This graph shows 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 party⁴⁶. 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.⁴⁷

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)⁴⁸. 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 party⁴⁹. 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)

⁴⁶We could also assume uncertainty over the realization of the Q surplus, but it would not change the intuition of the result below.

⁴⁷As 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.

⁴⁸In 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.

⁴⁹In 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 ⁵⁰.

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:

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

It is straightforward to see that the first best can only be achieved if the contract is of the form $b(Q(a_1, a_2)) = Q(a_1, a_2) - t$, where t is a transfer independent 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 $\{a_1^*, a_2^*\}$ ⁵¹.

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 $S^{SE} > S^{SO}$,⁵² 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

⁵⁰Alternatively, 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.

⁵¹In 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

⁵²See 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:

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

Similarly, 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:

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

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

$$S^{RE} \geq S^{SE} \quad (7)$$

(7) is actually sufficient as well as necessary, because a transfer t can always be chosen so that when (7) is satisfied, (5) and (6) are also satisfied⁵³.

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 negotiates 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:

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

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:

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

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:

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

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

⁵³For both (5) and (6) to be satisfied and the supplier to accept the contract, it must be that $Q_0 + \frac{1}{2}\frac{r}{1+r}\delta^2 \leq t \leq Q_0 + \frac{1}{2}\delta^2$

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

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 2, 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 4, 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 2 and 4 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 5, 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:

$$\begin{aligned} B_{ijt} = & \alpha + \beta_1 I[\text{VI} \times \text{supplies owner firm}]_{ijt} \times I[\text{Low Quality}]_{jt} \\ & + \beta_2 I[\text{VI} \times \text{supplies owner firm}]_{ijt} \times I[\text{High Quality}]_{jt} \\ & + \gamma_{ij} \times I[\text{High Quality}]_{jt} + \gamma_{ij} \times I[\text{Low Quality}]_{jt} + \delta_t + \varepsilon_{ijt} \end{aligned} \quad (12)$$

where $I[\text{Low Quality}]_{jt}$ 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 $I[\text{High Quality}]_{jt}$).⁵⁴ We include Supplier \times Plant \times Quality level fixed effects (that is, $\gamma_{ij} \times I[\text{High Quality}]_{jt}$ and $\gamma_{ij} \times I[\text{Low Quality}]_{jt}$) to focus on the supplier's *differential* response to the plant's input needs when integrated. The other variables are as defined in equation (1).

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.⁵⁵ 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.

⁵⁴We 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.

⁵⁵The 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 6 and 7. 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 (0.007)	0.016* (0.009)	-0.000 (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 (0.020)	0.026 (0.022)	0.002 (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 outsourcing				
Dep. var:	Protein content		Log(unit price)	
	(1)	(2)	(3)	(4)
Share of inputs from VI suppliers	1.044*** (0.342)	1.081*** (0.340)	0.088 (0.053)	0.090* (0.052)
Share of inputs from relational suppliers	-0.157 (0.505)	0.006 (0.441)	-0.008 (0.037)	0.003 (0.039)
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.	65.6	65.6	7.23	7.23
N	220	220	220	220
Panel B: Second definition of relational outsourcing				
Dep. var:	Protein content		Log(unit price)	
	(1)	(2)	(3)	(4)
Share of inputs from VI suppliers	1.072*** (0.268)	1.063*** (0.269)	0.089* (0.047)	0.088* (0.047)
Share of inputs from relational suppliers	0.208 (1.975)	0.409 (1.814)	0.018 (0.167)	0.032 (0.152)
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.	65.6	65.6	7.23	7.23
N	220	220	220	220

Notes: One observation is a firm during a production season. The regressions are similar to the ones in Table 4 but control for the share of inputs from independent suppliers under a relational contract with the downstream firm. *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. *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. *Share of inputs from relational suppliers* is the share of a firm's inputs that come from suppliers under a relational contract during a season. 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 20 times (median) with the same firm during a fishing season and so, for at least 3 consecutive fishing seasons. 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

Dep. var:	Panel A		
	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]	-0.133***	0.017	-0.013
×I[Plant producing low quality]	(0.043)	(0.047)	(0.031)
I[VI × supplies owner firm]	-0.066**	-0.067***	-0.042**
×I[Plant producing high quality]	(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 2, 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 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 concentrations⁵⁶ 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.⁵⁷

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 lower⁵⁸.

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’ movements⁵⁹.

⁵⁶We 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 \times 0.1° -longitude (roughly 10 kilometer \times 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.

⁵⁷A 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.

⁵⁸We 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).

⁵⁹Since 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.⁶⁰ 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:

$$\begin{aligned} \text{Quality}_{jt} = & \alpha + \beta_1 \text{VI}_{jt} + \beta_2 \text{Difficult conditions}_{jt} \\ & + \beta_3 \text{VI}_{jt} \times \text{Difficult conditions}_{jt} + \beta_4 \text{HighTech}_{jt} + \gamma_j + \delta_t + \varepsilon_{jt} \end{aligned} \quad (13)$$

where the firm×production season level continuous variable $\text{Difficult conditions}_{jt}$ 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.⁶¹

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.⁶² This indicates that vertically integrating

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

⁶⁰We 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.

⁶¹Greater plankton availability improves the fish's fatty acid profile, which in turn results in a fishmeal of higher protein content.

⁶²We 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 6 ⁶³.

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:

$$\begin{aligned}
 B_{ijt} = & \alpha + \beta_1 I[\text{VI} \times \text{supplies owner firm}]_{ijt} \times I[\text{Not difficult conditions}]_{ijt} \\
 & + \beta_2 I[\text{VI} \times \text{supplies owner firm}]_{ijt} \times I[\text{Difficult conditions}]_{ijt} \\
 & + \gamma_{ij} \times I[\text{Difficult conditions}]_{ijt} + \gamma_{ij} \times I[\text{Not Difficult conditions}]_{ijt} + \delta_t + \varepsilon_{ijt}
 \end{aligned} \tag{14}$$

where $I[\text{Difficult conditions}]_{ijt}$ indicates that the fishing conditions around plant j 's location are *difficult* on date t as defined above (and vice versa for $I[\text{Not difficult conditions}]_{ijt}$). Similar to the approach in Appendix B, we include Supplier×Plant×Difficult conditions fixed effects ($\gamma_{ij} \times I[\text{Difficult conditions}]_{ijt}$ and $\gamma_{ij} \times I[\text{Not Difficult conditions}]_{ijt}$) 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*⁶⁴.

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.

⁶³We also checked these results are not sensitive to how we define difficult production conditions. The corresponding tables are available from the authors.

⁶⁴These 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:	Protein content		Log(unit price)	
	(1)	(2)	(3)	(4)
Share of inputs from VI suppliers	1.313*** (0.407)	1.132*** (0.380)	0.137* (0.071)	0.115* (0.063)
Difficult conditions		-1.566 (0.942)		-0.181*** (0.061)
Share VI × Difficult conditions		1.730** (0.818)		0.237*** (0.055)
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.	65.4	65.4	65.4	7.20
N	179	179	179	179

Notes: One observation is a firm during a production season. The number of observations is lower than in Table 4 as observations after 2014 are excluded from the sample. (2015 is used to define the plankton concentration threshold at which the production conditions are considered difficult). *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 from a database that provides weekly prices by quality. *Share of inputs from VI suppliers* is the share of a firm's inputs that come from VI suppliers during a season. *High technology share of capacity* controls for the share of the firm's total processing capacity (measured in metric tons per hour and averaged across all active plants within the firm) that uses steam drying technology. *I[Difficult conditions]* is a dummy equal to 1 when the share of "good fishing locations" [$\text{Log}(\text{plankton concentration}) > 0.5$] 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). This dummy is defined at the port-day level, while the regressions are at the firm-season level, so the dummy variable is averaged by firm-season to construct *Difficult conditions*. This variable can be interpreted as the share of days when the conditions are difficult for a specific firm during a production season. The variable is normalized to a mean equal to 0 in our sample so that the first row of this table can be interpreted as the correlation between the *Share of inputs from VI suppliers* and quality. Standard errors clustered at the firm level are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

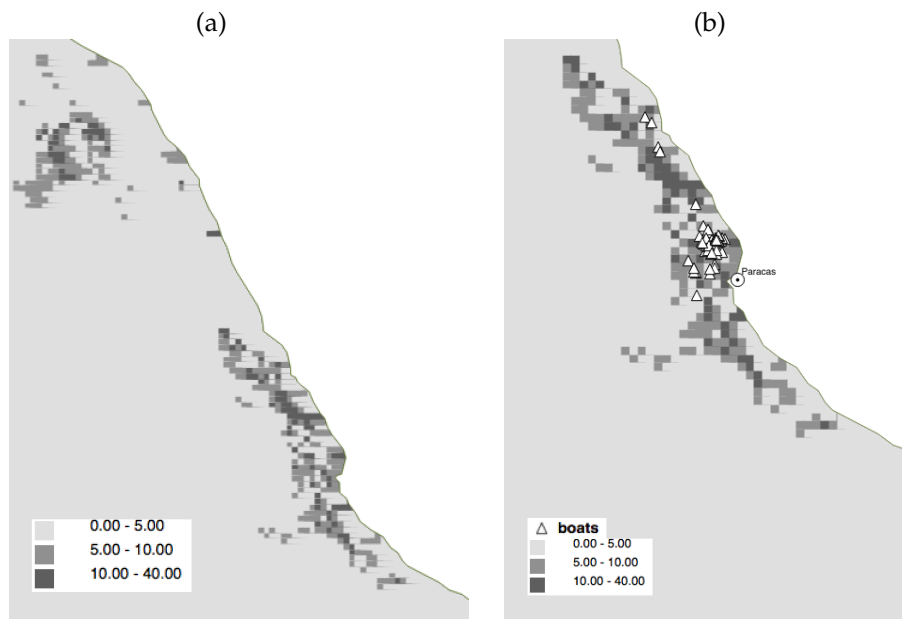
TABLE C2: SUPPLIER BEHAVIOR, VERTICAL INTEGRATION AND DIFFICULT UPSTREAM PRODUCTION CONDITIONS

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*** (0.024)	-0.039* (0.020)	-0.017 (0.017)
I[VI × supplies owner firm] × I[Difficult conditions]	-0.110 (0.154)	-0.355** (0.151)	-0.330*** (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 2 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

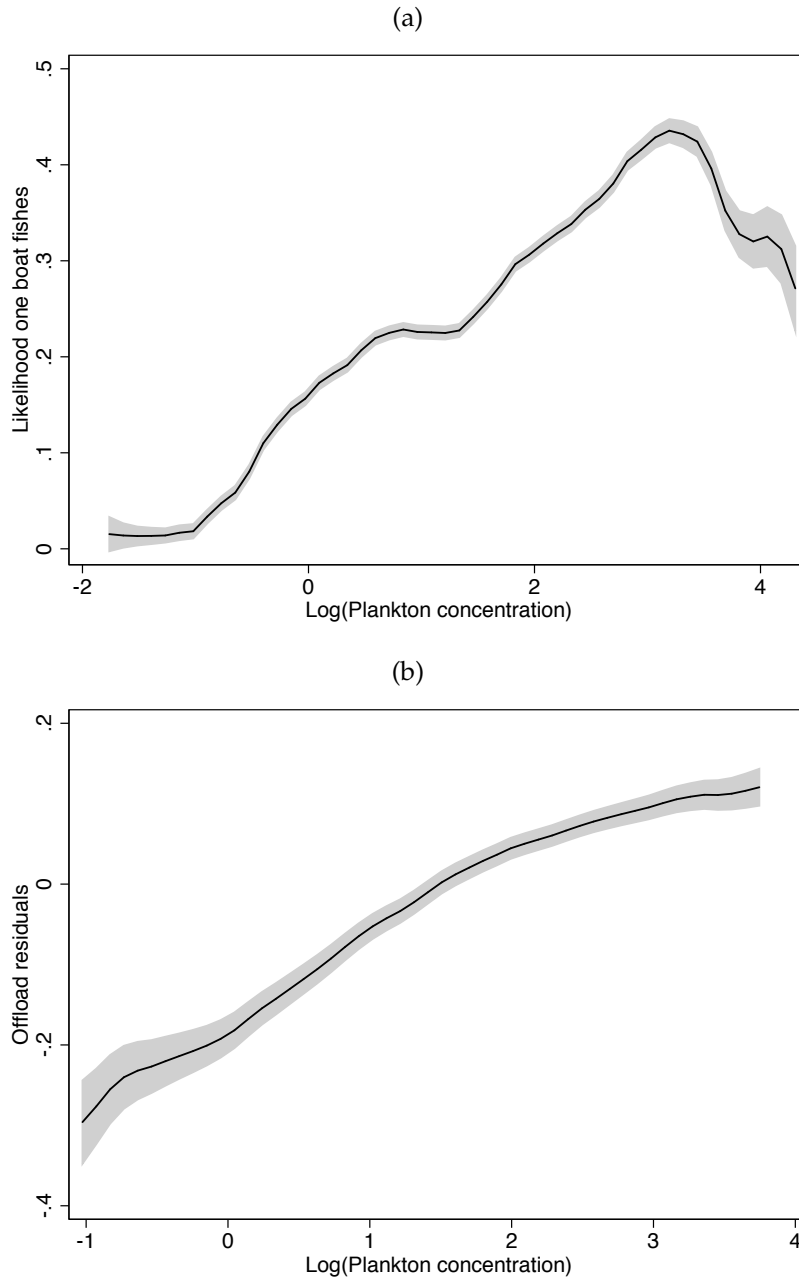
Appendix C figures

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



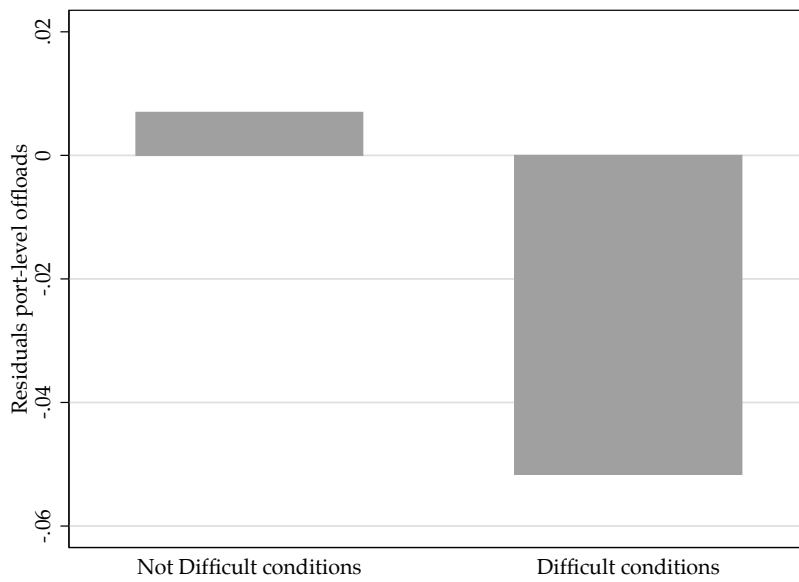
Notes: Panel (a) of this figure shows the distribution of plankton along the coast of Peru on December 10, 2012, as an example. A darker grey indicates a higher phytoplankton concentration (in mg/m^3). Panel (b) shows the same map zoomed around the port of Paracas, and the white triangles show where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

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



Notes: Panel (a) of this figure shows the likelihood that a boat fishes in a specific 0.1 degree \times 0.1 degree (roughly 10 kilometer \times 10 kilometer) grid-cell as a function of Log(phytoplankton concentration) at that location. Only locations where 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 C.

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 C.