Integration and Performance in Vertical Relations: Satellite-tracked Evidence*

Gabriel Natividad
NYU Stern†

19 May 2012

Abstract

In markets prone to stochastic supply shocks, some actions of suppliers are non-contractible; thus, vertical structures co-exist with independent supplier firms in equilibrium. Using satellite-tracked real-time data on the operations of a large vertically integrated fishing firm's own ships and its long-term supplier ships, this paper studies the causal impact of vertical integration on total factor productivity (TFP) after the acquisition of those suppliers. The results suggest an increase of 16% in TFP attributable to vertical integration. TFP growth amongst formerly-independent assets and teams appears to require the ability for more effort and operational flexibility going forward. Evidence consistent with changes in opinion about operational practices inside an integrated structure is provided.

*I am indebted to the Firm, whose identity is suppressed here, for access to its data and for unrestricted conversations and field research. I also thank the Ministry of Production of Peru for access to data. No confidential information is disclosed in this paper. John Asker encouraged me to write this paper, and his insight is much appreciated. Thanks are also due to Natarajan Balasubramanian, Nittai Bergman, Adam Brandenburger, Bob Gibbons, Ricard Gil, Avi Goldfarb, Matt Grennan, Nan Jia, Michael Katz, Mara Lederman, Jeffrey Macher, Kristina McElheran, Gonçalo Pacheco-de-Almeida, Antoinette Schoar, Rob Seamans, Brian Silverman, Olav Sorenson, Catherine Thomas, Birger Wernerfelt, Maggie Zhou, participants at the Econ-of-Strategy conference in New York, the ISNIE conference at Stanford, the AOM PDW on Firm Scope in San Antonio, the AEA Annual Meeting in Chicago, and seminars at the University of Maryland, the University of Toronto, MIT, the University of Piura, and the Peruvian Central Bank. Thanks to Miguel Ferré Trenzano for introducing me to industrial fishing in 1999. Javier Romero Haaker provided excellent research assistance. All errors are mine.

†Contact: Gabriel Natividad, NYU Stern. 40 West Fourth St., Tisch Hall 723, New York, NY 10012. Phone: (212) 998-0108. Fax: (212) 995-4235. E-mail: gnativid@stern.nyu.edu

1 Introduction

This paper studies the causal impact of vertical integration on total factor productivity (TFP). Basic as it sounds, this question has remained elusive to empirical research probably for two reasons. First, systematic information required to estimate total factor productivity within vertical structures is typically hard to find. Second, vertical integration is a firm scope decision that is difficult to justify in well-functioning markets, thus putting the burden of proof on the researcher to show what industry-wide or organizational sources of frictions might make the firm superior to the market. As a result, while prior work has shown that vertical integration impacts various dimensions of performance (Novak and Stern 2008, Lederman and Forbes 2010), and that TFP is positively associated to markets with higher prevalence of vertical integration (Hortaçsu and Syverson 2007), it has been relatively difficult to establish the direct statistical link between asset-level vertical integration and TFP, let alone the mechanisms explaining it (Lafontaine and Slade 2007, Syverson 2011). This paper introduces granular data on inputs, outputs, and operations in industrial fishing that allow for a more direct assessment of how and why vertical integration affects TFP.

From a theoretical standpoint, the fishing context seems attractive to generate new insight into the causes and consequences of vertical integration. On the one hand, because fishing companies face significant uncertainty in supply reflected in spatial and temporal variation of fish availability (Guénette et al. 2008, Fréon et al. 2008), industrial fishing offers features broadly consistent with Kranton and Minehart's (2000) characterization of why vertical structures exist alongside independent specialized firms. Though the focus of this paper is not on the drivers of vertical integration but on its consequences, it seems reasonable to assume that exogenous movements of wild fish create uncertainty in supply for reasons analogous to arguments about uncertainty in demand (Kranton and Minehart 2000), thus making vertical integration an equilibrium result for some firms. On the other hand, the operational and performance consequences of switching from independent operation to vertical integration often considered in theory can be observed in fishing both within and across firm boundaries (Gibbons 2010), thus connecting the organizational economics of production operations with market-based explanations for the nature and working of vertical structures (Bresnahan and Levin 2012).

The fishing data offer several advantages to empirically assess the impact of vertical integration on TFP. On the one hand, information on every fishing trip in the Peruvian fishing industry, the second-largest in the world, provides the baseline context to understand how ships that are part of a vertical structure (through common ownership with fish-processing plants) operate and obtain their output, fish, in comparison with independent ships. On the other hand, at a more granular level, a proprietary data set on the satellite-tracked real-time operations of a large vertically integrated fishing firm (hereafter called the Firm) in running its own ships and its relationships with long-term contractual suppliers offers unique measures of how much TFP actually changes after vertical integration and which operational practices may be conducive to that outcome. Importantly, variation in asset-level vertical ownership generated by the block acquisition of supplier fleets by the Firm that was not directly motivated by productivity enhancements helps implement a difference-in-differences approach that assuages the usual endogeneity concerns about vertical integration and outcomes.

The main findings of this paper are three. First, I find that vertical integration leads to a 16% increase in trip-level TFP with respect to long-term contractual relationships. Specifically, taking two transactions of the same ship, both with the same plant following the same route using the same level of production factors, I find that the transaction occurring when the ship is vertically linked through ownership with the Firm brings 16% more fish output than the one occurring when the ship is a friendly supplier external to the Firm. This result is largely insensitive to a variety of control variables and fixed effects, suggesting that vertical integration has a positive impact on TFP regardless of the particular model I investigate. Moreover, this difference-in-differences estimate using data on the Firm contrasts with a 'no effect' baseline estimate obtained using data on all trips of all ships in the industry, suggesting that an empirical design zeroing in on the measurement of TFP and the endogeneity of vertical integration seems warranted.

Second, I find that the effect of vertical integration on TFP is heterogeneous across different characteristics of the assets and of the asset operators that become part of a vertical structure. Specifically, larger and more powerful ships achieve higher TFP levels after integration; similarly, ships with younger captains become more productive after becoming integrated. Interestingly, the pre-integration levels of ship productivity and captain education do not matter for post-integration TFP. Overall, these patterns suggest that TFP growth amongst formerly-independent assets and teams requires the ability for more effort and more flexibility going forward, rather than a better track-record looking back.

Third, I provide suggestive evidence on the mechanisms behind the observed productivity benefits of vertical integration. Though by design these results cannot be as conclusive as the first two sets of findings, they depict a nuanced picture that appears to be internally consistent across different tests. First, I find that under a vertical integration regime, formerly-external ships travel longer distances at higher speed, increase their geographical concentration with other ships of the Firm, stick to the same extraction area before changing to the next one, reduce the meandering of travel in their return trips, and resell their catch less frequently to third parties in order to keep it instead for the Firm. Because all tests include route fixed effects, the models abstract away from the choice of fishing areas that ship operators may make differently within an integrated structure. Second, I correlate these operational changes with changes in ship-level TFP; in other words, I ask whether the observed differences in ship operations enacted by vertical integration may matter for TFP growth, and I find that this is so but only for a smaller set of operational changes. Specifically, rather than being associated with longer distance or higher speed inside a vertical structure (an "effort" story), TFP appears to grow with the geographical concentration with other ships of the Firm, the persistence in extraction areas, and diminished reselling to third parties. Thus, I view the impact of integration on TFP as a "changes in opinion" story because it does not seem to originate in the hard assets of a firm but in the different beliefs and expectations about what is the best way to find fish and become more productive.

This paper seeks to contribute to the ongoing debate on how vertical structures operate and why it matters. First, the study is among the first to assess the causal impact of vertical integration on TFP, thus adding a new component to productivity models that are widely used in evaluating the state of the broad economy (Syverson 2011), while also highlighting the link between managerial actions and productivity (Lieberman and Dhawan 2005). Second, the results of this paper suggest that hard measures of asset operations change significantly after vertical integration, possibly due to soft, intangible mechanisms that are enacted inside firm boundaries. After much empirical work in manufacturing industries, Atalay, Hortaçsu, and Syverson (2012) have recently proposed that the channel through which vertical integration affects productivity may be the transfer of intangibles, rather than physical goods. However, this conjecture does not squarely fit Pérez-González's (2005) finding that capital investment is a key mechanism linking ownership and TFP. My study suggests that the connection is less based on physical investments, as modeled by property-rights theories (Grossman and Hart 1986, Hart and Moore 1990), and possibly more interpersonal, as proposed by Van den Steen (2010). While the mechanism that I propose can only remain exploratory here, the magnitude of the potential effect warrants renewed attention in the literature going forward (Gibbons 2010, Bresnahan and Levin 2012).

2 Institutional Background

In the context of a vast literature in fishery economics and management, very little is known about the vertical relations between ships and plants (see van Putten et al. (2012) for an extensive review of fishing behavior research). Studying vertical integration in fishing bears resemblance to work on oil well drilling (Kellogg 2011), gasoline (Houde 2011), and airlines (Lederman and Forbes 2010) that has explored ownership, geographical, and time variation in relatively homogenous environments. In fishing, all these dimensions vary substantially, so it is illustrative to understand its fundamentals from the outset.

2.1 The activity

Industrial fishing consists in the capture of wild fish by large-scale vessels operated by professional crews. In the Peruvian fishery, the second most abundant in the world and the focus of this study, ships are technologically capable of extracting fish, not transforming it.[1] The mission of an extraction company is to capture fish from the Pacific ocean, bring it to the coast, and sell it or transfer it vertically to plants located nearby; by contrast, a processing company seeks to acquire or source raw fish to transform it into fishmeal (a flour-like paste used as an animal feed) or seafood at its plants. Transactions between extraction and processing parties happen over the market or within firm boundaries if they are integrated through ownership. When extracting fish for fishmeal, ships carry their catch without refrigeration or stringent sanitary previsions, and trips are typically short (e.g., less than 24 hours); by contrast, when extracting fish for seafood, ships must activate their refrigeration system and trips are longer and more cautious with respect to how the raw fish is handled. This study is centered on the extraction side of the industry and includes both types of fishing.

The sea's reserves are both mobile and imperfectly observable — wild fish are alive and instinctively do not want to be captured. Through technological developments in oceanographic detection, fishing companies have become more accurate over the years in determining fish locations, though the most effective way to search remains more art than science. Depending on sea conditions and the procurement needs of firms, ships conduct their extraction operations in seasons determined by government; thus, ships may fish more or less intensively within a given season.

There are no territorial partitions of fishing grounds — by law, any extraction company with a license can operate anywhere on the Peruvian ocean.[2] Though in principle ships could move around throughout the coast to search for and capture fish wherever they want, the combinatorial complexity of planning routes on a large surface appears daunting. In practice, the search choices of extracting

[1] See Bertrand et al. (2007) and Bertrand et al. (2008) for oceanographic descriptions of Peruvian fishing.

[2] Though I cannot demonstrate that all ships always have access to all fishing grounds, a simple inspection of landing data for the population of trips in the industry suggests that ship locations are sufficiently dispersed as to assuage concerns of unlawful foreclosure.

companies are facilitated by the different likelihoods to fish abundance across different latitudes and by the geographical deployment of other ships and plants along the coastline. These and other factors affecting fishing operations are incorporated into the empirical analysis.

2.2 The existence and nature of vertical relations in fishing

Because this study finds a positive effect of vertical integration on productivity, it is important to explain (i) why not every ship in the industry is vertically integrated, and (ii) whether the vertical ownership of ships by plants is in fact a way to remove some kind of friction, thus making the firm superior to the market. Although this paper does not seek to test for market frictions or other drivers of vertical integration, progress in the organizational literature can be more meaningful if these points are clarified (Bresnahan and Levin 2012), even if only qualitatively.

First, the nature and economic structure of fishing may help explain why both integrated firms and independent firms co-exist in equilibrium. On the one hand, fishing is an activity with stochastic geographic supply shocks that make the fixed locations of plants unequally attractive to ships at a given point in time. Thus, the spatial and temporal variation with which fish appears on the ocean can be effectively smoothed out by independent ships selling fish to third parties rather than being owned by a firm that may not use its ships because its plants are too far from the extracting area. As a result, the flexibility of being independent pays off under certain states of nature. On the other hand, government regulation prohibits vertical integration in some cases (i.e., wooden ships[3]), and the investment required to own a processing plant appears prohibitively expensive for some extraction companies that may not have good financial backing, thus remaining subject to contracting problems (e.g., Acemoglu, Johnson, and Mitton 2009). Hence, in this empirical setting, it is not surprising that some firms find it optimal to remain independent, developing (largely informal) supply relationships of different degrees with processing companies, or simply selling their output on the spot market.

Second, to understand the sources of frictions giving an edge to vertical integration, some stylized industry patterns summarized below can help to interpret changes in vertical integration as consistent or inconsistent with well-known theoretical mechanisms that make market transactions (e.g., contracts) less attractive than ownership. My understanding of these stylized patterns comes from field research spanning over a decade, and while some of these patterns can only be phrased verbally here for brevity, they provide a useful backdrop for the empirical design.

[3] All ships of the Firm and its long-term supplier firms studied here are steel ships.

- *Ownership vs. operation.*- Individuals operating ships are not ship owners, as those who can afford a ship with a license prefer to hire somebody else to run it. Ships are owned by firms, which may or may not be vertically integrated into fish processing.

- *Effort and employment relations.*- Neither the captains nor any member of a ship's crew receives a salary or any fixed compensation. Their sole monetary payment is a share of the catch market value, disbursed in cash the week after each trip. This proportional share is uniform across all ship crews and across organizational forms in the industry. Crew members' effort within a trip is reflected in their time aboard the ship, their promptness to cast the purse seine net deftly once the captains fix the fishing location in a matter of minutes, as well as other tasks; given their pay structure, fishing teams are incentivized to catch as much fish as possible during a trip.

- *No evidence of asymmetric information.*- Ship and plant operators communicate constantly and fluidly in real time in regard to their procurement needs. Firm managers observe the operations of their affiliated and contractual ships in real time. Ship captains communicate within and across firm boundaries using sophisticated telecommunication systems; though the quality of this information flow cannot be verified instantly, the repeated nature of fishing operations suggests that truthful revelation is plausible. Both vertically integrated and independent extraction firms seem to face the same challenges to find fish.[4]

- *Relational contracting.*- Fish-processing firms maintain several types of formal and informal relationships with fish suppliers. Across the industry, it is quite common that processing firms choose not to write trip-by-trip contracts, possibly due to the unobservability of catches in real time and to the relatively underdeveloped judicial system of the country.[5] Instead, they prefer to use unwritten non-binding agreements (hereafter called contracts for simplicity) of different strength. The strongest agreement is that in which the supplier firm shares information (e.g., satellite-tracking) with the client firm, and this is the focus of the main empirical analysis here. In fact, these informal agreements appear to be important. In the case of independent extraction companies, 70% of landings are at the processing firm's facilities where the independent extraction firm most frequently lands at. Given that ports rarely have fewer than three plants, this suggests that ships tend to favor some firms over others. Interestingly, plants do not pay markedly different

[4]Beyond this qualitative evidence, which I find credible based on my understanding of the industry, no test of an asymmetric information mechanism in vertical integration will be provided in this paper; no credible proxy for information that affects vertical ships differently from independent ships exists in my data. See Table 2 for models of horizontal integration that may be indirectly interpreted as considering this mechanism.

[5]See Acemoglu, Johnson, and Mitton (2009) for the relationship between cost of contracting and vertical integration across countries. According to the World Bank's Development Indicators, the procedures to enforce a contract in Peru are more numerous (i.e., more costly) than in the other top three fishing countries, China and Chile.

prices to ships in the same plant-week. Instead, these landing patterns are suggestive of the existence of relationships.

- *Price incentives*.- Processing firms pay the market price of raw fish to any of their suppliers regardless of the strength of their relationship. Market prices are determined weekly or even less frequently through supply and demand. Vertically integrated firms use the market price to compensate its ships' crews in the same way as independent firms do.

- *Asset specificity and bilateral dependency*.- Both ships and plants are generically oriented to commodity inputs and outputs. Unloading fish does not entail differential investments across client firms, and skippers can switch across firms. There are no single-plant ports throughout the coastline, and over 50% of all ships in the industry are not vertically integrated.

- *Hold-up risk*.- Across the industry, queuing time can be a source of hold-up power for plants. By contrast, in terms of queueing time, the Firm treats its long-term suppliers equitably on par with its ow ships; no statistical evidence of different queuing times for owned or contractual ships is found in the case of the Firm.

- *Deep pockets to finance search*.- Processing firms do not finance the fish search of third parties, as they pay for the output seven days after it is brought to land. While this pattern may suggest that vertical integration brings about a deep-pocket effect to search, the independent firms of interest here are large and could accumulate past profits to finance their own search.

Again, the aforementioned points speak to several potentially important rationales for why vertical integration matters in general, even if they will not appear to be satisfying explanations in this empirical setting in particular. Moreover, my upfront discussion of these channels does not seek to pick a single mechanism as the only plausible one. The main goal of this paper is to provide causal evidence on the direct statistical link between vertical integration and TFP.

In light of this qualitative characterization, a couple of reasons appear particularly likely to give vertical ownership an edge over contracts: (i) the widespread occurrence of stochastic geographic supply shocks that makes asset movements non-contractible, and (ii) the speculative nature of fishing practices that creates differences of opinion between ship operators and the managers of a vertically integrated firm. First, the stochastic nature of fish movements along the coastline can be harmful to independent ships, for example, through creating an abundant supply of fish that cannot be unloaded effectively at third-party plants;[6] this downside of independence can therefore be offset by vertical integration for some firms. Second, independent extraction companies develop their own practices and styles to execute fishing tactics, a realm of decision-making in which third-party plant managers typically cannot interfere. Qualitative evidence gathered through interviews across the industry suggests that when a processor really wants to wholly control the actions of an external supplier, the chosen path is vertical integration.

2.3 Empirical patterns of vertical integration and productivity across the industry

I supplement the above descriptions with an empirical characterization of vertical integration in the industry. To do this, I make use of data on the population of fishing operations in the industry between 2003 and 2008, and employ regression models relating ship output and vertical integration at the level of each fishing trip.

Panel I of Table 1 describes the variables employed for this characterization. Data on 380,982 trips include the identity of the transacting parties and their vertical or independent nature, the date of the trip, the output in tons of fish caught, the ship's hold capacity in cubic meters, and the geographical location of the destination port, as well as other market characteristics mentioned at the bottom of Panel II for the case of one firm but analogously available for the whole industry.

To measure the relationship between fish output (i.e., weight of fish caught) and vertical integration in the case of ship j during trip i ending on date t, consider

$$y_{ijt} = \beta_1 * VI_{jt} + \beta_2 * k_j + \beta_3 * X_{ijt} + \epsilon_{ijt} \quad (1)$$

where y is the logarithm of output, k is the logarithm of the capital equipment of a ship, that is, its hold capacity in cubic meters, X is a vector of control variables, and ϵ is an error term.

Table 2 presents regression results of two variants of equation (1). The first and second model implement a cross-sectional design in which $\epsilon_{ijt} = \tau_{t \times g} + \mu_{ijt}$, that is, day-port interaction fixed effects are included but no ship fixed effects are included. The first column of Table 2 displays a negative and significant coefficient on the dummy for vertically integrated ships, amounting to a 3.6% decrease in output. When expanding the specification to include a dummy for horizontal integration (i.e., the common ownership of the ship with other ships in the industry), still the coefficient on β_1 is negative and significant. Thus, in the cross-section of trips in the industry, it appears to be the case that independent

[6] The argument that natural shocks can be hedged or contracted upon has been shown to be empirically weak even in sophisticated markets such as the United States (e.g., Garmaise and Moskowitz 2009).

assets have a statistically robust, though economically small, premium over vertically integrated ships.

The last two columns of Table 2 introduce ship fixed effects and year-week dummies to specification (1). Specifically, by defining $\epsilon_{ijt} = \delta_j + \tau_{y.week(t)} + \mu_{ijt}$, the coefficient β_2 on logged capital cannot be separately identified any more from the ship-fixed effect, as ship capacity k_j is time-invariant by regulation in this industry. The results in the third model of Table 2 indicate that, when accounting for ship-fixed heterogeneity through ship fixed effects, the apparent connection between vertical integration and output disappears; vertical integration is also insignificant in the fourth model, which uses horizontal ship integration as an additional explanatory variable.

A first reading of Table 2 suggests that there are no clear benefits to vertical integration in terms of trip-level output across the industry. But there can be different interpretations of this result. On the one hand, the lack of significance of the vertical integration variable may point to market equilibrium forces obviating vertical scope; on the other hand, the result may also raise concerns that the statistical measurement is too coarse to make a clear statement about TFP in this setting and that the endogeneity of vertical integration choices muddles the analysis. These issues are the point of departure for my empirical investigation of vertical integration in relation to TFP.

3 Vertical Integration and Total Factor Productivity

My focus is on the relationship between upstream asset ownership by a vertical structure and TFP. I thus conceptualize vertical integration as each firm's choice of whether to own productive assets both in the upstream and downstream sectors, or instead exist as an independent specialist transacting on the market. This definition of vertical integration differs from others in the literature because it is centered on variation in asset-level ownership within a firm. By contrast, Hortaçsu and Syverson (2007) focus on vertical integration patterns *across* firms within markets, and Lederman and Forbes (2010) study the choice of *using* vertically integrated assets instead of third-party contractors, even if the choice of ownership had already been made in the past. As detailed in Section 2.2, industrial fishing offers substantial variation in asset ownership and vertical relations not only across firms but also within firms; this double feature of vertical integration has received little attention in the literature, probably because of data availability, and will be exploited here in the empirical analysis.

An equally important dimension of this study is the micro analysis of TFP. Though productivity is a central construct in the literature, definitions vary significantly from study to study, probably because of their within-industry nature, thus providing only indirect background to understand TFP more broadly.[7] Of course, an advantage of within-industry case studies is that they help control for the influence of other factors on productivity heterogeneity, and may serve to clarify the impact of the variable of interest. My study of industrial fishing seeks to provide new insight into how TFP varies at the transaction level.

Yet better measurement may not be enough to address concers about the endogeneity of scope and performance (Lafontaine and Slade 2007, Bresnahan and Levin 2012). Specifically, the match of firm scope and foreseeable outcomes such as productivity is typically non-random. For example, it has been argued that successful firms tend to be vertically integrated, a scope decision that may correlate well with higher TFP and success (Atalay, Hortaçsu, and Syverson 2012). Thus, the endogeneity of firm choices must be addressed when assessing the impact of vertical integration on TFP.

To do this, I employ data on a single firm in connection with its suppliers. The Firm that is at the core of this study obtains fish for processing from four sources. First, its own vertical ships; second, its long-term contractual supplier firms with whom it fully shares operating information and from whom it enjoys a special priority (with no obligation nor a guarantee to buy or sell on either part) to receive fish catches; third, weaker relationship suppliers without shared information or priority; and fourth, spot transactions. While the total number of ships providing the Firm with fish surpasses 1,000, this main sample for this study focuses on the first two groups, accounting for 53% of the Firm's fish before the latest corporate acquisitions and totaling between 50 and 100 ships, with the exact number omitted for confidentiality.[8]

The long-term suppliers of the Firm are firms with large fleets (i.e., numerous ships) that only have a long-term agreement with the Firm and not with other processing firms. Though an exact characterization of the relations of the Firm with all its suppliers is not possible (for example, no satellite data on weak-links or spot suppliers is available), qualitative evidence suggests that the Firm prefers to source fish from its own ships and next from its long-term suppliers, and only below these, from weaker links.

Focusing on the Firm and its long-term suppliers provides several advantages. First, TFP can be measured with precision at the trip level, as the inputs for production are revealed by the Firm to the researcher, thus going beyond what can be observed in the population of trips in the industry.

[7] For example, Ichniowski, Shaw, and Prennushi (1997) define productivity as production-line uptime in steel manufacturing, and Bandiera, Barankay, and Rasul (2005) and Mas and Moretti (2009) define productivity only in terms of one production factor, labor.

[8] Supplementary evidence on the complete procurement of fish by the Firm, including fish provided by weak-link relationships and spot suppliers, is discussed below.

(Compare, for example, the detail of Panel I and Panel II of Table 1). Second, the Firm pursued corporate acquisitions through which it acquired the entire fleets of its long-term contractual suppliers; as will be explained in Section 5.2, the reasons leading to these acquisitions were not directly related to micro productivity enhancements, thus allowing for a difference-in-differences analysis of the effect of vertical integration on TFP. Finally, the Firm provided me with satellite-tracked real-time data on the operations its own ships and the ships of its long-term contractual suppliers both before and after they were acquired. Given the general patterns of vertical integration in the industry displayed in Table 2, a micro-data case study offers an opportunity for better identification.

4 Data

A contribution of this paper is to bring satellite-tracked data to the study of vertical integration and productivity. Per the General Fishing Bill of Peru,[9] a system of satellite surveillance is mandated for all firms in the industry, with the goal of overseeing good fishing practices and controlling extraction of fish, a national good. Following this general law, the specific regulation for the satellite system[10] details that the functioning of the transmission equipment on board is mandatory for all industrial ships and for all periods regardless of whether the ships are fishing or not. The information transmitted by the approved device includes date and time, ship ID, longitude and latitude (with a precision of +/-100 meters), speed, and direction. The transmitter is required to allow for reprogramming the default interval between signals to any value in the range between 15 minutes and 24 hours at the regulator's request. Disconnecting the equipment from the system for whatever reason requires government authorization. The providers of the satellite system are private firms also authorized by the government.

An advantage of the satellite information is that it is available for both the Firm's own fleet and the fleets of its long-term suppliers before and after integration. The period covered is between 1 January 2003 and 20 April 2010, totaling over 3 million records.

This satellite information on ship movements becomes richer when combined with the hourly internal records of the Firm's plants and with daily economic information on all ships and plants in the industry. This information is supplemented with hand-collected ownership records and other regulatory information (e.g., ban announcements), and with detailed prices and quantitites for all the Firm's transactions. Because the Firm does not routinely merge its satellite information with its economic information, I employ simple algorithms to determine when the ship is actively fishing, resulting in a

[9]Decreto Supremo 012-2001-PE of 13 March 2001.
[10]Decreto Supremo 026-2003-PRODUCE of 12 September 2003.

total of 13,536 trips with 689,084 satellite snapshots, an average of 51 snapshots per trip. The exact location of each plant of the Firm, each point in the trajectory of a ship, and points on the coastline are specified using coordinates from the NOAA National Geophysical Data Center website. Distances between two points in the trajectory of a ship, or between the exact location of a ship and a plant, are calculated using the haversine formula, which accounts for the curvature of the Earth and performs particularly well for the numerical computation of small distances.[11]

Table 1 defines the main variables of the study and provides summary statistics for the Firm's own ships and its long-term contractual suppliers. Each trip's catch carries a mean (median) output of 203 (178) tons of fish; catch weight is highly dispersed, as the standard deviation is almost as large as the mean. Several exogenous factors such as bans and ban announcements are also summarized.

How are the satellite data exploited in the analysis of vertical integration and productivity? Recall that the economic value of fishing operations is uncertain before starting a fishing trip, so that real-time marine operations are crucial to profitability. By knowing the time and location of each ship during a fishing trip, a number of unique operational constructs can be gauged. To illustrate, Figure 1 shows satellite information on two trips of the same ship before and after the acquisition by the Firm. The conditions of the trips are similar, but the trajectories are quite different. The goal of the empirical analysis is to assess the impact of vertical integration after controlling for other factors that may influence the behavior and performance of a ship on a given trip.

5 Empirical Design

I seek to understand the impact of vertical integration on TFP at the level of each transaction, that is, a fishing trip. I first start describing a standard multiplicative production function and then extend it to consider vertical integration. Then I detail the identification strategy.

5.1 Specification

For the case of ship j during trip i ending on date t, consider the multiplicative form

$$Y_{ijt} = L_{ijt}^{\alpha} * K_{ijt}^{\rho} * M_{ijt}^{\gamma} * X_{ijt}^{\theta} * \epsilon_{ijt} \tag{2}$$

where K_{ijt} is the use of capital during the trip; L_{ijt} is the use of labor; M_{ijt} is the use of fuel and other materials, X_{ijt} are control variables, and ϵ_{ijt} is an error. Importantly, note that all variables are

[11] The haversine formula is available at `http://www.movable-type.co.uk/scripts/latlong.html`.

in physical quantities, not in monetary values, thus providing some methodological advantages (Foster, Haltiwanger, and Syverson 2008). A key property of fishing trips is that the use of capital, labor, and inputs can be parameterized as a direct function of the duration of the trip, labeled here as U_{ijt}, thus leaving

$$Y_{ijt} = (L_j * U_{ijt})^{\alpha} * (K_j * U_{ijt})^{\rho} * (M_j * U_{ijt})^{\gamma} * X_{ijt}^{\theta} * \epsilon_{ijt} \tag{3}$$

$$= U_{ijt}^{\phi} * (L_j)^{\alpha} * (K_j)^{\rho} * (M_j)^{\gamma} * X_{ijt}^{\theta} * \epsilon_{ijt} \tag{4}$$

Transforming (4) taking logs leaves the general productivity specification

$$y_{ijt} = \phi * u_{ijt} + \alpha * l_j + \rho * k_j + \gamma * m_j + \theta * x_{ijt} + \nu_{ijt} \tag{5}$$

and by extension, the impact of vertical integration on total factor productivity can be captured by

$$y_{ijt} = \beta * VI_{jt} + \phi * u_{ijt} + \alpha * l_j + \rho * k_j + \gamma * m_j + \theta * x_{ijt} + \nu_{ijt} \tag{6}$$

where vertical integration is an asset-level, time-variant indicator variable. An advantage of the data is that the use of inputs u_{ijt} is observed with precision through the satellite tracking system (e.g., Mullowney and Dawe 2009), thus capturing much variation from the residual to focus on the marginal effect of vertical integration β for the estimation. While this cross-sectional specification narrowly compares vertically integrated assets with third-party ships in their trip TFP, the preferred specification improves upon (6) to estimate

$$y_{ijt} = \beta * VI_{jt} + \phi * u_{ijt} + \theta * x_{ijt} + \delta_j + \tau_{y.week(t)} + \kappa_{(rw)\times j} + \zeta_{(rw)^2\times j} + \lambda_{route} + \nu_{ijt} \tag{7}$$

where δ_j is a ship fixed effect that subsumes the ship-fixed factors l_j, k_j, and m_j; linear and quadratic relative week trends κ_{rw} and $\zeta_{(rw)^2}$ with respect to the date in which each asset becomes vertically integrated are interacted with ship fixed effects; and λ is a route fixed effect. The first key advantage of specification (7) consists in capturing asset-level time-invariant heterogeneity. (The identification assumptions required to give this fixed-effects approach a causal interpretation are discussed in the next subsection). A second benefit of equation (7) is the inclusion of route fixed effects λ, defined in the case of these fishing trips as a joint fixed effect for the triad *(plant of departure, plant nearest the zone of fishing, plant of arrival).* This route fixed effect captures much of the variation that could be attributed to the fact that ships move to latitudes[12] that are more productive; instead, by fixing routes through these triple interaction dummies, the effect of vertical integration is assessed within a given route for

[12]The Peruvian coastline runs roughly from North to South; thus, the plants used for these fixed effects mostly capture North-South movements in various directions.

a given ship. Third, including linear and quadratic trends interacted with ship fixed effects assuages concerns that other shifters moving simultaneously with the adoption of vertical integration are driving the results. Because trips of the same ship may be correlated, standard errors are clustered at the level of each ship.

Moreover, the inclusion of exogenous control variables strengthens specification (7). The control variables in x_{ijt} are largely exogenous to the Firm's actions. Fuel prices capture exogenous cost shocks. Bans on anchovy or white anchovy are dummies reflecting whether government has imposed a ban on fishing these species for fishmeal at the exact location of the landing port at the moment of arrival; these variables control for potential changes in the behavior of ships trying to get into (or out of) a particular zone to avoid penalties. Announced bans on anchovy or white anchovy are dummies for the period immediately preceding (but after the announcement of) the enactment of such bans, thus capturing an increased intensity in fishing behavior. In addition, fish for fishmeal is a dummy that captures technical differences across fishing purposes described in Section 2. This variable becomes fixed immediately after the trip starts, as mandated by regulation. Overall, the inclusion of these control variables helps distinguish the marginal effect of vertical integration from other drivers of TFP.

5.2 Estimation and Identification

Identification is facilitated by the observation of changes in ownership brought about by corporate acquisition events. Specifically, at different points in time the Firm decided to more fully integrate backwards, ceasing to become a buyer of its long-term suppliers. Focusing on *asset* (i.e., ship) level changes is justified by the fact that operational decisions are made by asset operators on the sea in conjunction with the Firm's managers on land.

Consider two ships, each supplying fish to the downstream plants of the Firm. One of the ships is a long-term external supplier while the other is owned by the Firm. Neither ship is owner-operated. At some point, the external ship's operator learns that its owner has changed, becoming the Firm. Under the null that the operation and performance of the ships are random, the ship operators do not care, so the acquisition will have no effect. However, if ownership by the Firm enacts policies affecting the way controlled ships operate, the acquisition will lead to changes in the functioning of the external ship. Because the acquisition events shift vertical integration status only for formerly independent ships, leaving already owned ships unchanged, a difference-in-differences test would remove biases due to omitted variables or endogeneity concerns.

This empirical design requires that long-term contractual firms are acquired for reasons somewhat unrelated both to the way ships operate on the sea, and also unrelated to anticipated reasons for these ships to change their fishing behavior after the acquisitions. Understanding the cause of the acquisitions is important because, if there are unobservables that are correlated with the Firm's decision to acquire its long-term suppliers and are also correlated with the operations and performance of the acquired ships, the tests would be problematic. Note that I do not need to assume that the acquisitions were randomly determined. Rather, I need to inquire whether acquisition-related unobservables are mechanically determining the behavior and performance of ships.

Several pieces of qualitative evidence suggest that operational changes were not anticipated by those involved in the acquisition deals. First, in interviews with shareholders and managers of the Firm, I learned about the conditions leading to the acquisitions: the need of larger corporate size to face a more competitive environment, the need of a larger balance sheet to raise more banking funds, and a largely-exogenous request of more collateral for a syndicate of banks as the basis for more sophisticated financing operations in the future. Because none of these aspects is linked to any differential behavior of newly-acquired vs. already owned ships, the assumption of uncorrelated unobservables seems plausible.

Second, I also met with external participants in the acquisition decision: the law firm and the investment bank advising the Firm. When asked whether changes in fishing behavior and productivity were foreseen at the moment of negotiations, the principals of these advising firms told me that that was not the case. The driving force behind the acquisitions was largely financial rather than operational.

There exist other advantages of the empirical setting further strengthening the use of corporate acquisitions as a valid shifter in vertical integration: the exogenous variation of fish behavior and the quasi-experimental features of the acquisitions. On the one hand, the biological factors that affect ship productivity are largely random and therefore, even if they are unobservable, they are unlikely to be correlated with changes in the vertical integration status of a ship. On the other hand, if the industry patterns described qualitatively in Section 2.2 hold, specification (7) is likely to gauge the main statistical effect of vertical integration, and to a large extent keep several plausible theoretical mechanisms relatively flat. For example, changes in ship crew composition and economic incentives, relation-specific investments, or the threat of hold-up are themes that run through the literature as important mechanisms linking vertical integration and performance, but they do not seem to bind in the dealings of the Firm and its long-term suppliers. Although the Firm hypothetically could have decided to change these dimensions of vertical relations after the acquisitions, the institutional details suggest that unilaterally changing these dimensions would have been detrimental, as the norm of all other firms reveals equilibrium behavior. The interpretation of why vertical integration affects TFP is conducted in an exploratory manner in Section 7 after presenting the main results.

6 Main Results

6.1 Vertical Integration and TFP

Panel A of Figure 2 presents a standard difference-in-differences graph analyzing the output of catches before and after vertical integration. Raw monthly data and unconditional means are shown, without removing any micro variation. The overall pattern in Panel A of Figure 2 is that integration leads to a substantial increase in output among formerly contractual ships, leaving already owned ships relatively unaffected. However, the scattered points show much variation that must be addressed in the regression analysis.

Table 3 presents the main regression results. The sample for these tests are trips of the Firm's ships and its long-term contractual firms' ships. The first model estimates equation (6), that is, a cross-sectional regression that does not include ship fixed effects or year-week dummies. Instead, in order to make the comparison across ships meaningful, interaction fixed effects based on the triad *(date of arrival, departure port, and arrival port)* are employed, as well as hour-of-departure dummies. The results suggest that seemingly identical transactions are more productive for vertical integration ships than for independent ships in the cross-section, with a 7% differential accruing to vertical integration. As expected, the coefficients on capital, labor, and use are positive and significant. The economic magnitude of the gains due to vertical integration in these cross-sectional estimates appears somewhat small, suggesting that ship operators may not be fully aware in real time of the benefits of integration across ships with different ownership affiliation.

The second column of Table 3 introduces ship fixed effects. The coefficient on vertical integration is statistically significant and economically larger than in the cross-sectional estimation, indicating a gain of 14.4% for formerly-external contractual ships that become part of the Firm after vertical integration. As detailed in Section 5.1, production inputs are uniquely captured by the use variable, modeled as the duration of the trip in hours as measured by the satellite. This specification also allows for the introduction of a number of exogenous controls that vary for each day or each week.

The results so far may be subject to the criticism that other shifters of productivity vary at the same time as does vertical integration status. To assuage this concern, the third model of Table

3 implements specification (7) using linear and quadratic trends measured as relative distance to the moment in which assets became vertically integrated (or zero if they never changed vertical status), interacting these linear and quadratic weekly trends with ship fixed effects. The coefficient on vertical integration is still positive and economically significant, indicating a 17.3% TFP gain attributable to the fact that the ship is vertically linked to the Firm.

While the difference-in-differences design focuses on changes happening among those ships formerly independent that become integrated, the Firm also obtains fish from its own ships as well as hundreds of other weak-relationship and spot ships. The positive effect of vertical integration on TFP among formerly long-term contractual ships could be argued as being jointly determined by other policies of the Firm with respect to its sources of fish. For example, the Firm could start using its own ships *less* than before the acquisition events, thus displacing them from the ocean in order to give preference to the newly integrated firms. Moreover, the Firm might be changing its sourcing strategies with respect to other third party ships.

To address these arguments empirically, the fourth model of Table 3 includes control variables capturing the micro dynamics of competition and procurement. Specifically, a control for the number of ships in the fleet to which the ship in question belongs helps address a displacement channel by which productivity would increase for newly integrated ships. Moreover, a control for the number of weak-relationship and spot ships unloading fish at that port during that week captures local competition and sourcing strategies implemented by the Firm that may by conceptualized as an alternative driver of the focal ship's productivity (Syverson 2004). Finally, a third additional control variable uses the number of fishing infractions committed by third-party ships (i.e., not the Firm's or its long term suppliers') in the specific geographic area of the destination port during the week in question, thus controlling for any turbulence that might affect productivity. After introducing these controls, the effect of vertical integration on TFP is statistically and economically similar, reflecting a 16.1% gain to vertical firms.

Panel B of Figure 2 displays kernel densities of TFP residuals for formerly contractual ships and owned ships both before and after integration. The differential shift to the right for treated ships is roughly comparable to the average treatment effect captured by the regression coefficients in Table 3.

What is, then, the interpretation for how vertical integration improves TFP in this setting? Because the main production inputs are accounted for in the specification, the answer to how firms become more productive is likely to be found in their increased chances to find fish. On the one hand, at the moment when integration happens, not all assets or asset operators may be equally ready to achieve increments in TFP. On the other hands, after integration, the channels leading to increased chances to find fish require a careful examination of fishing operations. I turn to this analysis next.

6.2 Heterogeneity in the effect of vertical integration on TFP

The results thus far have captured the global effect of vertical integration on the whole set of newly integrated assets. In a new set of tests, I ask whether the effect of vertical integration on TFP might differ depending on the characteristics of the asset or the asset operators that become part of a vertical structure.[13] To conduct these tests, I use continuous ship-level variables for the pre-integration period on the ships that are to become vertically integrated, and create an above-median dummy that interacts the main variable of interest, *vertically integrated*, to explain TFP. For these tests to be effective, it is important to focus exclusively on conditions at the moment when integration starts. Specifically, I use information on the physical characteristics of the acquired ships — productivity, size, engine power — and of their captains — age, education. Though limited in nature, these interaction variables do provide a transparent characterization of where the benefits of integration may be more pronounced.

The results of these tests are reported in Table 4. Larger and more powerful ships achieve higher TFP levels after integration; similarly, ships with younger captains become more productive after becoming integrated. Interestingly, the pre-integration levels of ship productivity and captain education do not matter for post-integration TFP. Overall, these patterns suggest that TFP growth amongst formerly-independent assets and teams requires the ability for more effort and more flexibility going forward (e.g., Leamer 1999), rather than a better track-record looking back (e.g., Rosen 1982). The specific operational practices brought about by integration are analyzed next.

7 Discussion of Causal Mechanisms

How do operations change with vertical integration, and how does that matter for TFP? Using the satellite data, I construct variables reflecting operational changes that may explain the observed productivity gains.[14] I proceed in two steps. First, I employ specification (7) using the operational metrics as dependent variables at the trip level. (All these variables are defined in Table 1, and they are

[13] A different dimension of heterogeneity is the timing of TFP changes after integration. In untabulated tests, I find that the gains in TFP for newly integrated ships are quite immediate, accruing in the first year after integration, and staying high for subsequent years.

[14] Though ideally I would have liked to use standard variables from the fishing literature, there is really no prior art in modeling micro fishing dynamics with satellite data (Putten et al. (2012)). Thus I construct measures based on my understanding of fishing practices.

modeled in logarithms in the regressions to interpret the regression coefficients as percentage changes.) Next, to assess which of these mechanisms may matter most for the newly integrated ships to increase TFP, I conduct a ship-period analysis of TFP and operational practices.

Table 5 displays the regression results. The first and second columns show that under integration, formerly-external ships travel longer distances and conduct their trips at higher speed. Because the specification controls for the duration of the trip (i.e., use), an increase in distance traveled is tantamount to an increase in speed, as revealed by the commensurate point estimates on vertical integration in the regressions. Moreover, as displayed in the third column of Table 5, vertical integration increases the variance in the speed of ships between different points of a trip. Because all tests include route fixed effects, these mobility patterns reveal more thoroughness in covering an area rather than a different set of geographic choices to search for fish.

The next two columns of Table 5 model whether the geographical concentration of ships changed substantially in relation to other ships. The fourth column shows no significant increase in the average distance with respect to the nearest neighbor ship of the fleet to which the ship is associated; not surprisingly, because this distance is not changing but fleet ownership is, the fifth model of Table 5 shows a negative effect of vertical integration on separation with respect to the nearest Firm's ship. After integration, the Firm does not seem to be spreading out its ships more widely.

Does vertical integration lead to more effort in the intensive margin within a trip?[15] The sixth column of Table 5 indicates that this is not so, as ships do not make more tries ("fishing sets") throwing their nets to catch fish. This result suggests that, conditional on trip conditions, ship operators both before and after integration are equally incentivized to catch as much fish as possible; effort in casting the large purse seiner net to catch fish does not vary under the vertical structure.

The last three columns of Table 5 capture other operational decisions that may affect a ship's chances to obtain fish. In the seventh column, the regression uses the distance with respect to the prior trip's extraction area as the dependent variable, showing a negative and significant coefficient on vertical integration. Thus, newly-integrated ships become more persistent in their attempts to find fish. The eight column of Table 5 employs the ratio of the distance traveled in the return trip divided by the straight line between the return point and the arrival port; this "meander ratio" variable is significantly reduced by vertical integration, suggesting that ships are more promptly returning to their plants without meandering unnecessarily during the return trip. Upon arrival, moreover, newly integrated

[15] A different dimension in which effort might change is the extensive margin, i.e., whether ships started fishing more often after integration. In untabulated tests, I find no significant effect of vertical integration on the frequency of trips.

ships show a stronger tendency to transfer their goods to the Firm rather than to third parties, as displayed in the last model of Table 5. Overall, the results suggest that integration with the Firm brings a strong focus to be persistent in some fishing areas, make more direct return trips, and keep the increased productivity gains exclusively through vertical transfers rather than market transactions. Plausibly, these practices are consistent with finding more fish or reducing travel costs.

Which of the several mechanisms explored in Table 5 may matter the most for gains in TFP among ships that become integrated? I construct a TFP residual at the ship level for each of these ships in the periods both before and and after integration, and then correlate the increment in TFP residuals with the changes in operational practices before and after integration. Figure 3 (with details in its caption section) displays fractional polynomial graphs for the variables that appeared to be significantly affected by vertical integration in Table 5. The results of this analysis suggest that distance to prior area, return meandering, and resale are mechanisms pointing in a consistent direction in the regression analysis of Table 5 and in the graphical analysis of Figure 3. For example, formerly-external ships that used to distance themselves the most from their prior trip's fishing area are the ones that benefit the most in terms of TFP after integration.

Overall, the patterns found in the data are not unilaterally pointing to greater effort. Rather, the findings suggest a change in opinion with respect to how best to conduct fishing operations. Moreover, because the compensation structure of ship operators remains purely variable (i.e., based on the quantity of fish output per trip) after integration, the fact that these actions were not pursued before integration even if they were in the best interest of the long-term supplier companies and their ship operators suggests that external ship operators thought such policies would not work. Thus, the gray area of differences in opinion may have plausibly, though not exclusively, changed through vertical integration, leading to substantial productivity gains.

How were these operational changes enacted inside the vertical structure after integration? Qualitative evidence suggests that the role of the Firm's managers on land is crucial to bring ship operators in line with the procurement needs of the plants, as well as to agree on search, deployment, and operational practices. For example, coordination meetings at the beginning of each fishing season served as a vehicle to communicate fishing tactics. Importantly, before integration, the captains of long-term supplying ships were invited to these meetings with the Firm but did not really attend, as they had their own planning meetings for their fleets.

Do TFP gains represent higher profitability? Given the longer distances traveled by integrated ships, it may be argued that ultimately the replacement cost of assets could swamp productivity gains. I collected additional evidence to assess whether newly integrated ships become more profitable. First, although cost data are not available for external ships before integration, I conducted out-of-sample predictions to fill in this cost series using available operational data; based on this imputed cost series, profitability seemed to significantly increase (in dollar terms) for newly integrated ships after accounting for all the monetary expenses of vertical operations. Second, just by focusing on the post-acquisition period, I found a very high correlation (e.g., above 92%) between TFP residuals of the models presented here and expanded-factor TFP models using very detailed information on material expenses such as fuel, third-party maintenance services, and so on. All in all, the evidence is consistent with a significant increase in profitability attributable to vertical integration.

Finally, it is important to assess whether these productivity gains matter for the broader procurement strategies of the Firm. On the one hand, the identification of TFP gains in the subset of owned and long-term contractual ships provides insight into a value-enhancing dimension of corporate scope within a given set of relationships, even after controlling for other procurement shifts in all models. On the other hand, as detailed in Section 3, the Firm sources fish from weaker-relationship ships, as well, so the broader impact of the TFP growth achieved post-integration is not measured by my specification. In untabulated tests, I find that the share of total fish procured from weaker-relationship and spot ships decreases after integration; it is therefore likely that the Firm used its productivity gains for substitution of the market rather than for increased market power.

8 Conclusion

In this paper, I use satellite-tracked real-time data on the movements of a large vertically integrated fishing firm's own ships and its long-term supplier ships to assess the causal impact of vertical integration on total factor productivity after the acquisition of those suppliers. The results indicate an increase of 16% in TFP attributable to vertical integration. Productivity gains require the ability for more effort and more flexibility going forward, rather than a better track-record looking back. Moreover, these gains appear to be driven by a change in opinion about the best way to conduct productive operations that is enacted within firm boundaries through vertical integration.

Recent research on the nature and working of vertically integrated firms has proposed that the gains through which vertical integration affects productivity may be the transfer of intangibles, rather than physical goods. The results of this paper give a new perspective to this conjecture by finding that hard, physical measures of asset operations change substantially after vertical integration, possibly due to soft, intangible mechanisms that are better enacted inside firm boundaries.

More broadly, different scholarly traditions have contributed to our current understanding of firm scope and performance. While the market and the firm are sometimes viewed as opposite ends of the spectrum across different paradigms, integrating these views of contracting offers a promising avenue for future work in strategy.

References

Acemoglu, Daron, Simon Johnson, and Todd Mitton, 2009, Determinants of vertical integration: Financial development and contracting costs, *Journal of Finance* LXIV, 1251–1290.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson, 2012, Why do firms own production chains?, University of Chicago mimeo.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul, 2005, Social preferences and the response to incentives: Evidence from personnel data, *Quarterly Journal of Economics* 120, 917–962.

Bertrand, Sophie, Arnaud Bertrand, Renato Guevara, and François Gerlotto, 2007, Scale-invariant movements of fishermen: The same foraging strategy as natural predators, *Ecological Applications* 17, 331–337.

Bertrand, Sophie, Erich Díaz, and Matthieu Lengaigne, 2008, Patterns in the spatial distribution of Peruvian anchovy (engraulis ringens) revealed by spatially explicit fishing data, *Progress in Oceanography* 79, 379–389.

Bresnahan, Timothy, and Jonathan Levin, 2012, Vertical integration and market structure, in Robert Gibbons, and John Roberts, ed.: *Handbook of Organizational Economics* . , vol. forthcoming (Princeton University Press).

Foster, Lucia, John Haltiwanger, and Chad Syverson, 2008, Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?, *American Economic Review* 98, 394–425.

Fréon, P., M. Bouchon, C. Mullon, C. García, and M. Ñiquen, 2008, Interdecadal variability of anchoveta abundance and overcapacity of the fishery in Peru, *Progress in Oceanography* 79, 401–412.

Garmaise, Mark J., and Tobias J. Moskowitz, 2009, Catastrophic risk and credit markets, *Journal of Finance* 64, 657–707.

Gibbons, Robert, 2010, Inside organizations: Pricing, politics, and path dependence, *Annual Review of Economics* forthcoming.

Grossman, Sanford J., and Oliver D. Hart, 1986, The costs and benefits of ownership: A theory of vertical and lateral integration, *Journal of Political Economy* 94, 691–719.

Guénette, Sylvie, Villie Christensen, and Daniel Pauly, 2008, Trophic modelling of the Peruvian upwelling ecosystem: Towards reconciliation of multiple datasets, *Progress in Oceanography* 79, 326–335.

Hart, Oliver D., and John Moore, 1990, Property rights and the nature of the firm, *Journal of Political Economy* 98, 1119–1158.

Hortaçsu, Ali, and Chad Syverson, 2007, Cementing relationships: Vertical integration, foreclosure, productivity, and prices, *Journal of Political Economy* 115, 250–301.

Houde, Jean-Francois, 2011, Spatial differentiation in retail markets for gasoline, *American Economic Review* forthcoming.

Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi, 1997, The effects of human resource management practices on productivity: A study of steel finishing lines, *American Economic Review* 87, 291–313.

Kellogg, Ryan, 2011, Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch, *Quarterly Journal of Economics* 126, 1961–2004.

Kranton, Rachel E., and Deborah F. Minehart, 2000, Networks versus vertical integration, *RAND Journal of Economics* 31, 570–601.

Lafontaine, Francine, and Margaret Slade, 2007, Vertical integration and firm boundaries: The evidence, *Journal of Economic Literature* 45, 629–685.

Leamer, Edward E., 1999, Effort, wages, and the international division of labor, *Journal of Political Economy* 107, 1127–1163.

Lederman, Mara, and Silke J. Forbes, 2010, Does vertical integration affect firm performance? Evidence from the airline industry, *RAND Journal of Economics* 41, 765–790.

Lieberman, Marvin B., and Rajeev Dhawan, 2005, Assessing the resource base of Japanese and U.S. auto producers: A stochastic frontier production function approach, *Management Science* 51, 1060–1075.

Mas, Alexandre, and Enrico Moretti, 2009, Peers at work, *American Economic Review* 99, 112–145.

Mullowney, D. R., and E. G. Dawe, 2009, Development of performance indices for the Newfoundland and Labrador snow crab (Chionoecetes opilio) fishery using data from a vessel monitoring system, *Fisheries Research* 100, 248–254.

Novak, Sharon, and Scott Stern, 2008, How does outsourcing affect performance dynamics? Evidence from the automobile industry, *Management Science* 54, 1963–1979.

Pérez-González, Francisco, 2005, The impact of acquiring control on productivity, Manuscript.

Rosen, Sherwin, 1982, Authority, control, and the distribution of earnings, *The Bell Journal of Economics* 13, 311–323.

Syverson, Chad, 2004, Market Structure and Productivity: A Concrete Example, *Journal of Political Economy* 112, 1181–1222.

——— , 2011, What Determines Productivity?, *Journal of Economic Literature* 49, 326–365.

Van den Steen, Eric, 2010, Interpersonal authority in a theory of the firm, *American Economic Review* 100, 466–490.

van Putten, I.E., S. Kulmala, O. Thébaud, N. Dowling, K.G. Hamon, T. Hutton, and S. Pascoe, 2012, Theories and behavioural drivers underlying fleet dynamics models, *Fish and Fisheries* 13, 216–235.

Figure 1: Satellite Snapshots of Ship Trajectories

The Cartesian maps show two trips of the same ship, before and after being acquired by the Firm, both times fishing for fishmeal and obtaining similar output (low). The solid squares are points in each trajectory, labeled with the ship's speed in km/hour rounded to integers (bottom map). The small dots represent the coastline, East of which there is land.

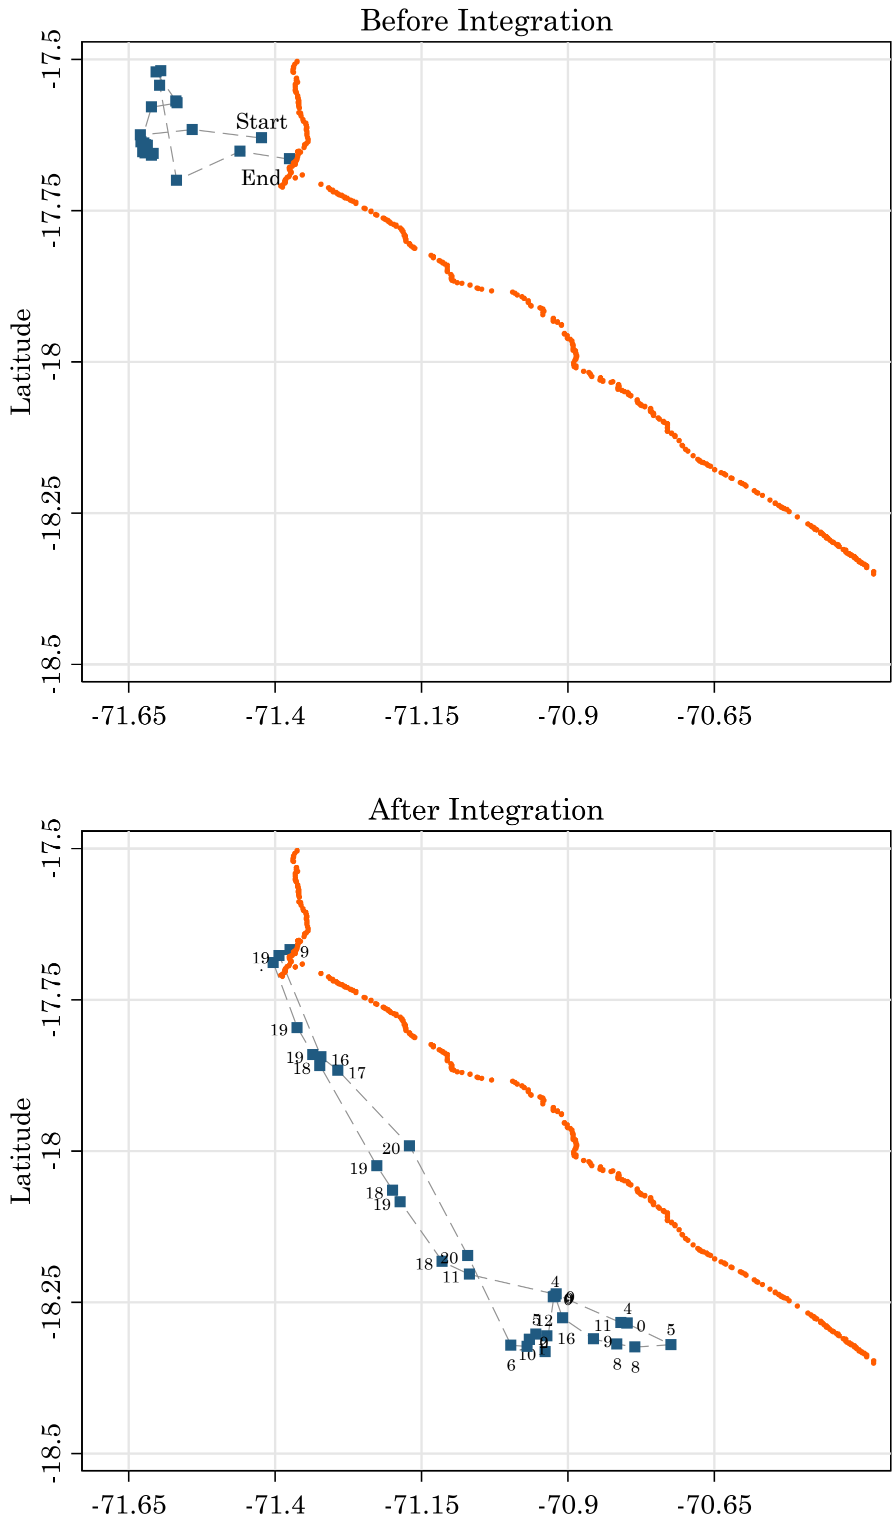

Figure 2: Output and Productivity by Treatment Before and After Vertical Integration

This figure shows output data (Panel A) and productivity residuals (Panel B) on owned and contractual ships before and after vertical integration into the Firm. Panel A summarizes the weight of catches at the monthly level for all trips of owned ships (circles) and long-term supplier ships that become integrated after being acquired by the Firm (squares). The pre- and post-integration sample averages of the raw monthly data are displayed a thick line for ships that become integrated and as a thin line for ships that are always owned. Panel B displays the kernel densities of residuals of a full trip-level productivity regression (Model 3.4) grouping the observations into four juxtaposed curves: owned by the Firm before integration of long-term contractual suppliers (thin solid line), owned after integration (thin dashed line), long-term contractual ships before they become integrated (thick solid line), and long-term contractual ships after they become integrated (thick dashed line); the four kernel density curves are shown only for the range [-1,1] for ease of reading.

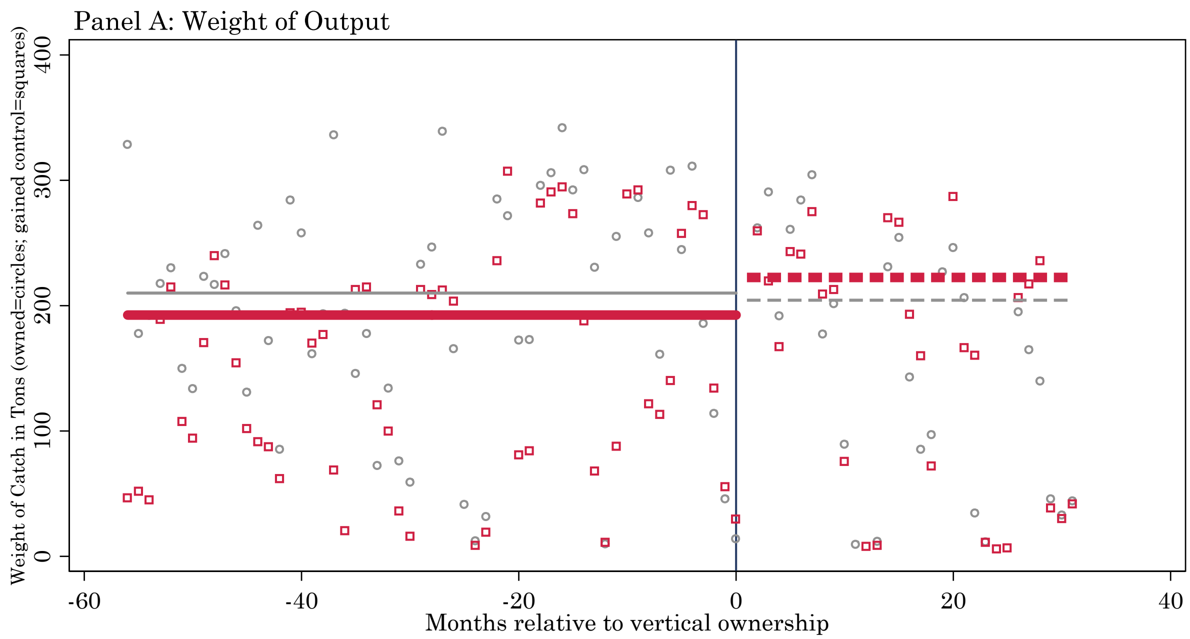

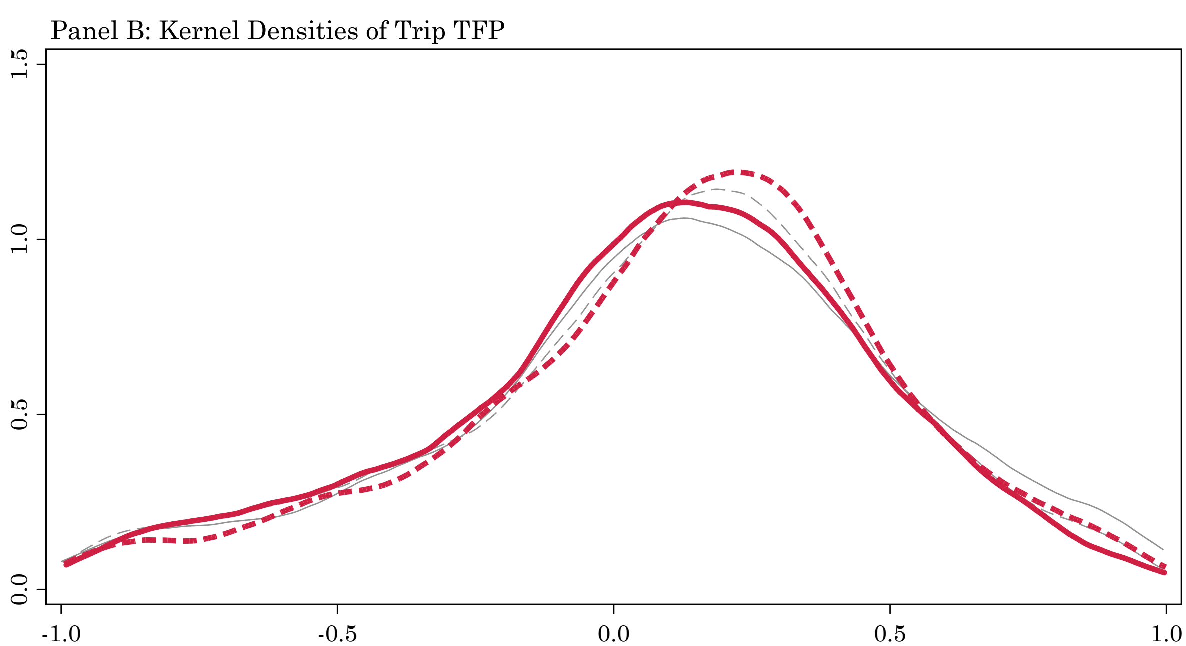

Figure 3: Relating Changes in Ship TFP to Changes in Trip Operations

The sample for the graphs consists of long-term contractual ships before and after they become integrated into the Firm. Each graph presents a fractional polynomial fit of changes in ship-level TFP on changes in ship-level trip operating variables, with 95%-level confidence intervals shown in gray shades. The estimation summarizes all inputs, outputs, and operations for these ships before and after vertical integration, thus obtaining two observations per each ship i to estimate the regression

$$y_{it} = \beta_1 * l_{it} + \beta_2 * k_{it} + \beta_3 * u_{it} + \alpha * 1(t = post) + \epsilon_{it}$$

For each ship-period observation, ship TFP is defined as the residual $\widehat{e_{i,t}} = y_{it} - \widehat{y_{i,t}}$ using the above regression coefficient estimates. For each ship, the increment of this productivity residual is defined as Δ Ship TFP$_i = \widehat{e_{i,post}} - \widehat{e_{i,pre}}$ and plotted in the vertical axis of each graph. The horizontal axis measures the increment in an operating variable in the period after integration with respect to the period before integration for each ship.

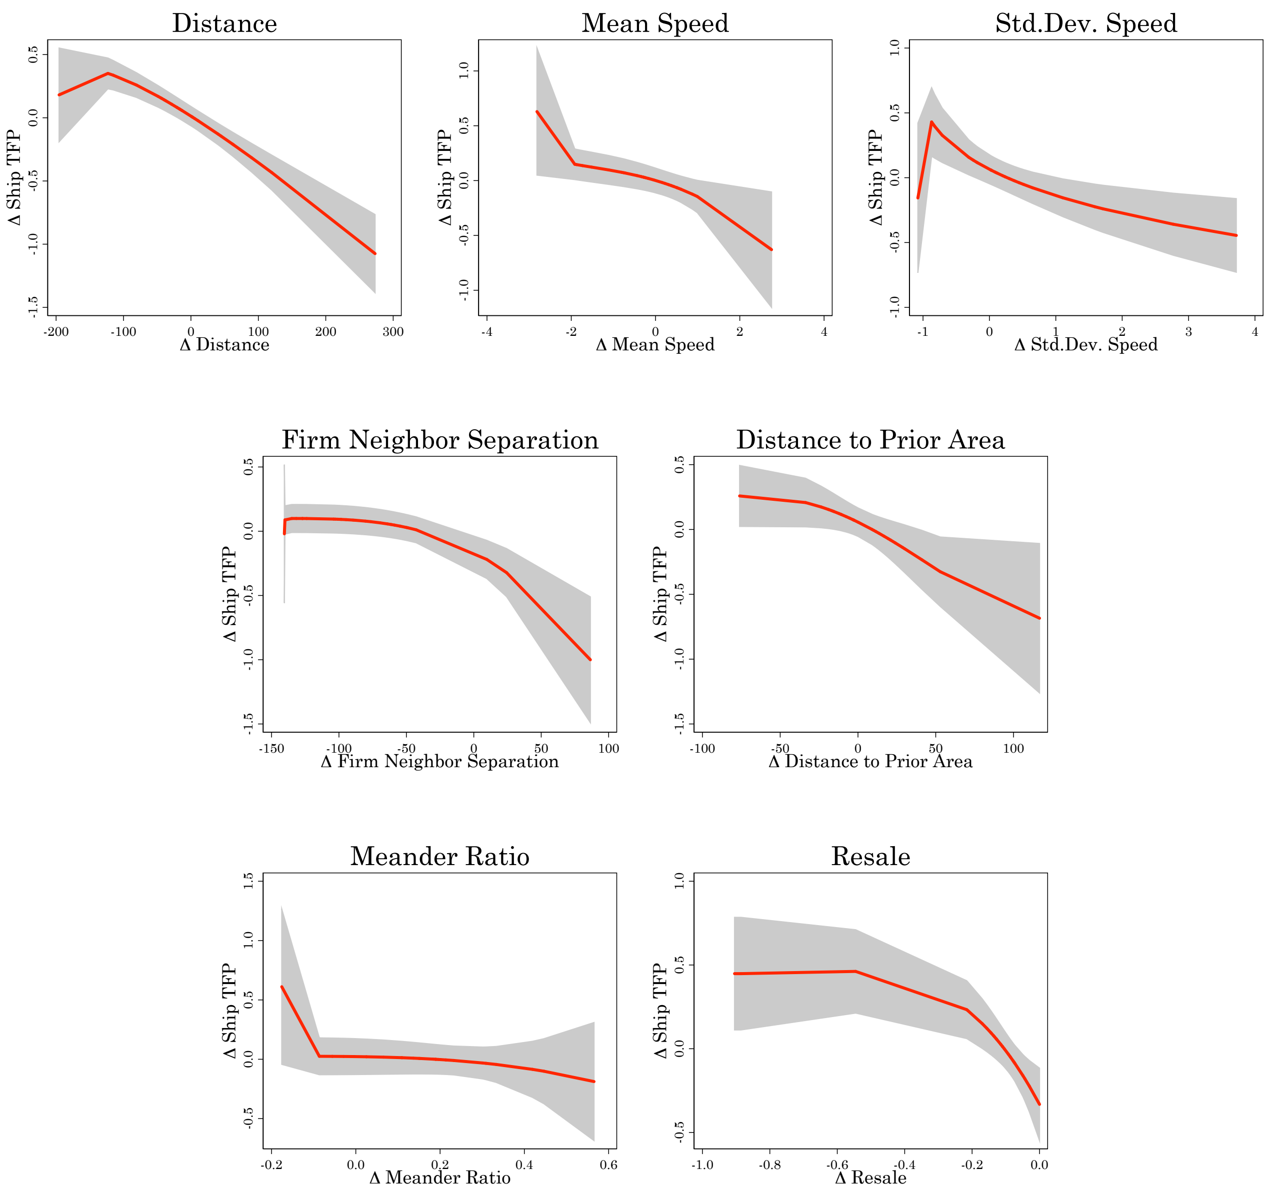

Table 1: Summary Statistics

This table defines trip-level variables on the population of ships in the industry (Panel I), trip-level variables on ships owned and contracted through long-term agreements by the Firm (Panel II), and information on the Firm's ships with exact details suppressed for confidentiality (Panel III). Vertical integration is a dummy for whether the ship is owned by a firm that also owns fish processing plants; in the case of Panel II, this is the Firm. Output is in tons (000 of kilos) of fish caught by the ship on the trip. Capital is the ship's hold capacity in cubic meters. Labor is the number of crew members on the ship. Use is the duration of the trip in hours. Distance is the exact geographical movement of the ship throughout its trajectory as measured by the satellite. Average speed is distance divided by use. The standard deviation of speed on a given trip uses the speed between consecutive locations measured by the satellite to calculate the trip-level standard deviation of speed. Neighbor separation is the average distance to the nearest ship calculated at each point of the trajectory of the ship with respect to any ship of the Firm or its long-term contractual suppliers. Firm neighbor seperation is an analogous measure but calculated only with respect to the Firm's owned ships. Number of tries is the number of times that the ship is detected to be moving at less than 2.9 kilometers per hour, proxying for an attempt to fish. Distance to prior fishing area is the distance of the current trip's maximum point off the coast and the same ship's maximum point off the coast on the immediately preceding trip, in kilometers. Meander ratio is the distance travelled by the ship between its maximum point off the coast and its destination plant divided by the straight line between those two points. Resale to third parties is a dummy for whether the ship sells its catch to a firm different from the Firm.

Fuel price is a daily series of price per gallon of diesel expressed in dollars after national taxes. Ban is a dummy for whether there is a ban on fishing anchovy for fishmeal that day at the location of landing. Announced ban is a dummy for the days between government's announcement of a ban and the day of enactment of the ban. Fish for fishmeal is a dummy equal to one when the purpose of the trip is to catch fish for fishmeal, and zero when the purpose is to catch fish for seafood; by regulation, the purpose of a trip is always determined in advance so it is not contingent upon within-trip decisions.

I. Industry Trip-level population (n=380,982)	Mean	Median	Std. dev.	Min.	Max.
Vertically integrated	0.26	0.00		0.00	1.00
Output	89.21	51.89	97.10	0.01	810.05
Capital	187.12	128.00	151.99	2.78	868.27
II. The Firm's Trip-level sample (n=13,536)					
Vertically integrated	0.51	1.00		0.00	1.00
Output	203.45	177.60	150.24	0.05	624.97
Capital	367.52	365.00	163.96	72.84	614.05
Labor	15.39	16.00	2.61	9.00	18.00
Use	32.61	23.21	28.28	2.00	397.00
Distance	304.04	197.01	352.12	8.12	2528.65
Mean speed	8.57	8.18	3.52	0.31	32.76
Std.dev. speed	6.24	6.22	2.45	0.00	46.94
Neighbor separation	33.98	9.45	82.86	0.02	1433.01
Firm neighbor separation	91.80	40.27	150.73	0.00	1444.59
Number of tries	15.19	11.00	18.20	0.00	451.00
Distance to prior fishing area	107.72	47.72	173.20	0.00	1413.19
Meander ratio	1.86	1.16	2.17	1.00	72.42
Resale to third parties	0.07	0.00	0.25	0.00	1.00
Fuel price	2.37	2.42	0.36	0.48	3.16
Ban on anchovy	0.22	0.00		0.00	1.00
Ban on white anchovy	0.19	0.00		0.00	1.00
Announced ban on anchovy	0.09	0.00		0.00	1.00
Announced ban on white anchovy	0.08	0.00		0.00	1.00
Fishing for fishmeal	0.56	1.00		0.00	1.00

III. The Firm's sample	Count
Total number of ships	$50 < N < 100$
Always owned, no long-term contract	$\approx 0.37N$
First long-term contract, then owned	$\approx 0.60N$
Always long-term contract	$< 0.03N$

Table 2: Vertical Integration and Trip Productivity across the Industry

This table reports estimates of the relationship between vertical integration and output across the industry for the period of interest. The unit of observation is a fishing trip, and the sample is all trips of all ships of the industry. All variables are defined in Table 1. The first and second models implement a cross-sectional approach using date × arrival port interaction dummies but no ship fixed effects. The third and fourth models use year-week dummies and ship fixed effects. Robust clustered standard errors are reported in parentheses.

Dependent Variable:	**Log(Output)**			
	(2.1)	(2.2)	(2.3)	(2.4)
Vertically integrated	−0.036***	−0.041***	−0.021	−0.018
	(0.00)	(0.00)	(0.02)	(0.02)
Horizontally integrated		0.012***		−0.007
		(0.00)		(0.02)
Log (Capital)	0.927***	0.927***		
	(0.01)	(0.01)		
Fuel price			−0.160	−0.160
			(0.14)	(0.14)
Ban on anchovy			−0.060***	−0.060***
			(0.01)	(0.01)
Ban on white anchovy			−0.016	−0.015
			(0.02)	(0.02)
Announced ban on anchovy			−0.096***	−0.096***
			(0.01)	(0.01)
Announced ban on white anchovy			0.097***	0.097***
			(0.01)	(0.01)
Fishing for fishmeal	0.820***	0.815***	0.809***	0.809***
	(0.03)	(0.03)	(0.07)	(0.07)
Fixed effects:				
Ship	No	No	Yes	Yes
Year-Week	No	No	Yes	Yes
Date × Arrival port	Yes	Yes	No	No
Arrival port	No	No	Yes	Yes
R^2	0.77	0.77	0.71	0.71
Sample size	380982	380982	380982	380982
Number of clusters (Date × Arrival port)	16241	16241		
Number of clusters (ships)			1643	1643

***, **,* significant at the 1%, 5% and 10% level. Standard errors are heteroskedasticity-robust and clustered as indicated.

Table 3: Vertical Integration and Trip TFP

This table reports estimates of the effect of vertical integration on total factor productivity (TFP). The observations are at the trip level and include all trips of the Firm's ships or its long-term contractual suppliers. The variables are described in Table 1. The first model implements a cross-sectional estimation with detailed within-day controls but no ship fixed effects. The second, third, and fourth models include year-week dummies and ship fixed effects. The fourth model includes three additional control variables. The number of ships in the fleet fishing this week considers each firm's fleet (the Firm or its long term suppliers) as distinct before integration but pools over all of them as a single firm after integration. The number of third party ships at this port are counts of all distinct ships that are not the Firm's or its long term suppliers' and provide fish this week to the Firm's plant to which the ship arrives during this trip. The number of third party infractions in this area is based on a registry of ship-trip level infractions for each firm in the industry, and excludes any infraction of ships of the Firm or its long term suppliers. Robust clustered standard errors are reported in parentheses.

Dependent Variable:	**Log(Output)**			
	(3.1)	(3.2)	(3.3)	(3.4)
Vertically integrated	0.070**	0.144***	0.173***	0.161***
	(0.03)	(0.05)	(0.06)	(0.05)
Log (Capital)	0.861***			
	(0.04)			
Log (Labor)	0.569***			
	(0.18)			
Log (Use)	0.060**	−0.008	−0.008	−0.006
	(0.03)	(0.02)	(0.02)	(0.02)
Fuel price		−0.234	−0.254	−0.280
		(0.34)	(0.33)	(0.33)
Ban on anchovy		−0.102**	−0.107**	−0.102**
		(0.04)	(0.04)	(0.05)
Ban on white anchovy		0.001	0.007	0.000
		(0.06)	(0.06)	(0.06)
Announced ban on anchovy		0.043	0.034	0.037
		(0.06)	(0.06)	(0.06)
Announced ban on white anchovy		−0.031	−0.027	−0.029
		(0.06)	(0.07)	(0.07)
Fishing for fishmeal	0.023	0.004	0.008	0.007
	(0.02)	(0.02)	(0.02)	(0.02)
Number of ships in fleet fishing this week				0.031***
				(0.01)
Number of third party ships at this port, this week				−0.001**
				(0.00)
Number of third party infractions in this area, this week				0.000
				(0.00)
Fixed effects:				
Ship	No	Yes	Yes	Yes
Year-week	No	Yes	Yes	Yes
Date × Departure port × Arrival port	Yes	No	No	No
Hour of departure dummies	Yes	No	No	No
Departure port × Fishing port × Arrival port	No	Yes	Yes	Yes
Linear week trend × Ship Fixed Effect	No	No	Yes	Yes
Quadratic week trend × Ship Fixed Effect	No	No	Yes	Yes
R^2	0.88	0.72	0.72	0.72
Sample size	13536	13536	13536	13536
Number of clusters (Date × Dep.port × Arr.port)	4697			
Number of clusters (ships)		*N*	*N*	*N*

***, **,* significant at the 1%, 5% and 10% level. Robust standard errors are clustered by date-dep.date-arr.port or by ship.

Table 4: Heterogeneity in the Impact of Vertical Integration on TFP

This table reports estimates of the heterogeneous effect of vertical integration on TFP depending on ship and captain characteristics. The specification is exactly as in Table 3 with the additional inclusion of an interaction term for vertical integration with Z. Term Z is a greater-than-median dummy based on information on or before the date when integration is enacted; while the underlying data are granular, the greater-than-median dummy is based on ship-level comparisons using exclusively data on ships that become integrated into the Firm. High pre-productivity is based on the residual of a productivity regression. Large size is based on each ship's hold capacity. High engine power is constructed using the horse-power specification of each ship's engine. Old age captains is based on the age of the captain and chief mate at the moment of integration. Highly educated is based on the schooling of the captain and chief mate of the ship at the moment of integration. Robust standard errors clustered at the level of each ship are in parentheses.

	Dependent Variable: **Log(Output)**				
Z interaction term:	High Pre-Productivity	Large Size	High Engine Power	Old Age Captains	Highly Educated
	(4.1)	(4.2)	(4.3)	(4.4)	(4.5)
Vertically integrated	0.152**	0.104	0.090	0.203***	0.172***
	(0.06)	(0.07)	(0.07)	(0.06)	(0.06)
Vertically integrated × Z	0.022	0.093*	0.110*	−0.091*	−0.030
	(0.04)	(0.05)	(0.06)	(0.05)	(0.05)
Log (Use)	−0.006	−0.006	−0.006	−0.005	−0.006
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Fuel price	−0.283	−0.272	−0.270	−0.273	−0.281
	(0.33)	(0.33)	(0.33)	(0.33)	(0.33)
Ban on anchovy	−0.102**	−0.104**	−0.106**	−0.102**	−0.102**
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Ban on white anchovy	0.001	0.003	0.004	0.001	0.000
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
Announced ban on anchovy	0.036	0.039	0.038	0.037	0.036
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
Announced ban on white anchovy	−0.029	−0.031	−0.030	−0.030	−0.029
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)
Fishing for fishmeal	0.007	0.007	0.007	0.007	0.007
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
N.ships in fleet this week	0.031***	0.032***	0.033***	0.034***	0.031***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
N.3rd.party ships this port, week	−0.001**	−0.001**	−0.001**	−0.001**	−0.001**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
N.infractions this area, this week	0.000	0.000	0.000	0.000	0.000
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Fixed effects:					
Ship	Yes	Yes	Yes	Yes	Yes
Year-week	Yes	Yes	Yes	Yes	Yes
Dep.port × Fish.port × Arr.port	Yes	Yes	Yes	Yes	Yes
Linear trend × Ship Fixed Effect	Yes	Yes	Yes	Yes	Yes
Quad. trend × Ship Fixed Effect	Yes	Yes	Yes	Yes	Yes
R^2	0.72	0.72	0.72	0.72	0.72
Sample size	13536	13536	13536	13536	13536
Number of clusters (ships)	N	N	N	N	N

***, **,* significant at the 1%, 5% and 10% level. Robust standard errors are clustered by date-dep.date-arr.port or by ship.

Table 5: Vertical Integration and Trip Operations

The observations are at the trip level and include all trips of the Firm's ships or its long-term contractual suppliers. All dependent variables are in logarithms, and described in Table 1. All explanatory variables are as in the fourth model of Table 3 but not displayed here for brevity. Robust standard errors clustered by ship are in parentheses.

Dependent Variable:	**Distance**	**Mean Speed**	**Std.Dev. Speed**	**Neighbor Separation**	**Firm Neighbor Separation**	**N.Tries**	**Dist. to Prior Area**	**Meander Ratio**	**Resale**
	(5.1)	(5.2)	(5.3)	(5.4)	(5.5)	(5.6)	(5.7)	(5.8)	(5.9)
Vertically integrated	0.094^{**}	0.094^{**}	0.082^{*}	0.054	-0.365^{*}	−0.038	-0.173^{**}	-0.081^{**}	-0.062^{***}
	(0.04)	(0.04)	(0.05)	(0.10)	(0.21)	(0.03)	(0.07)	(0.04)	(0.02)
Other controls as in Model (3.4)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects:									
Ship	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-week	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Dep.port × Fish.port × Arr.port	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear trend × Ship Fixed Effect	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quad. trend × Ship Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.88	0.53	0.47	0.56	0.63	0.74	0.31	0.25	0.29
Sample size	13536	13536	13536	13533	13511	13536	13536-N	13492	13536
Number of clusters (ships)	N	N	N	N	N	N	N	N	N

***, **,* significant at the 1%, 5% and 10% level. Standard errors are heteroskedasticity-robust and clustered by ship.