

Supplier Search and Re- matching in Global Sourcing – Theory and Evidence from China

Fabrice Defever, Christian Fischer, Jens Suedekum

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Supplier Search and Re-matching in Global Sourcing – Theory and Evidence from China

Abstract

In this paper, we consider a dynamic search-and-matching problem of a firm with its intermediate input supplier. In our model, a headquarter currently matched with a supplier, has an interest to find and collaborate with a more efficient partner. However, supplier switching through search and re-matching is costly. Given this trade-off between the fixed costs and the expected gains from continued search, the process will stop whenever the headquarter has found a sufficiently efficient supplier. Using firm-product-level data of fresh Chinese exporters to the United States, we obtain empirical evidence in line with the predictions of our theory. In particular, we find that the share of short-term collaborations is higher in industries with more supplier-cost dispersion, an indication of higher expected search opportunities.

JEL-Codes: F230, D230, L230.

Keywords: input sourcing, relational contracts, supplier search.

Fabrice Defever
University of Nottingham / UK
fabrice.defever@nottingham.ac.uk

Christian Fischer
University of Bayreuth / Germany
christian.fischer@uni-bayreuth.de

Jens Suedekum
Düsseldorf Institute for Competition Economics
(DICE), Heinrich-Heine-University
Düsseldorf / Germany
suedekum@dice.hhu.de

November 1, 2017

The authors thank Pol Antràs, Wilhelm Kohler, Erdal Yalcin, and the participants of the CESifo Summer Institute 2015 (“New developments in Global Sourcing”) in Venice for helpful comments and suggestions. All errors are our responsibility. Suedekum acknowledges financial support by the German National Science Foundation (DFG), grant SU-413/2-1.

1 Introduction

Numerous studies from the management literature argue that headquarter corporations have a vital interest in establishing and maintaining long-term collaborations with their intermediate input suppliers.¹ But before such an enduring relationship can be formed, the headquarters first needs to find a supplier who is technologically capable of producing the desired component in appropriate quantity and quality and at low costs. Finding such an efficient and reliable partner is not easy, and can be thought of as a complicated search-and-matching problem that may involve a time-consuming trial-and-error process before the firm is finally satisfied.

Several authors have recently incorporated elements of search and matching into models of international trade. Rauch (1999) and Antràs and Costinot (2011) study the search of a domestic producer for a foreign sales agent or distributor. Once the firm has found a suitable agent and a long-term collaboration is formed, the partners also start investing more into the relationship. This has been shown by Araujo, Mion and Ornelas (2016) and Aeberhardt, Buono and Fadinger (2014). In their frameworks, exports are increased at the intensive margin as the firm builds up trust that the current foreign distributor does not behave opportunistically, cheat on the exporting firm, engage in haggling and re-negotiation, and so on.

When it comes to global sourcing (i.e., the search-and-matching decision of the domestic *importing* firm with respect to their foreign input suppliers), there is consistent empirical evidence showing that firms indeed engage in a trial-and-error search. Rauch and Watson (2003) and Besedes (2008) find that domestic firms initially place small test orders when they first deal with a new foreign supplier. If they are dissatisfied, they terminate the collaboration, while they increase order volumes when they are satisfied. There are also recent theoretical models that explicitly consider the search of domestic or multinational enterprises for foreign input suppliers, see Rauch and Watson (2003), Carballo, Ottaviano and Volpe-Martincus (2013) and Bernard, Moxnes and Ulltveit-Moe (2014). However, until very recently, global sourcing models had a hard time coming to grips with those empirical patterns about trial-and-error search.

A starting point of the global sourcing literature is the canonical incomplete contracts model by Antràs (2003) and Antràs and Helpman (2004), which places the seminal property-rights theory of Grossman and Hart (1986) and Hart and Moore (1990) in an international trade context. It features hold-up and underinvestment problems as the key elements of the buyer-supplier relationship, which can be fine-tuned by the firm's decision how to organize the production process and the ownership structure along the value chain. As it is well known, these baseline models by Antràs (2003) and Antràs and Helpman (2004) are static models of a one-shot interaction and, therefore, not suitable to study long-term collaborations. However, several authors have recently embedded this model into a dynamic setting with repeated interactions; see Baker, Gibbons and Murphy (2002). In particular, Kukharsky (2015), Kamal and Tang (2015) and Kukharsky and Pflüger (2011) show that the inefficiencies caused by contract incompleteness can be overcome by *relational contracts*, that is, by informal, non-binding and purely trust-based agreements to cooperate in an enduring relationship. These papers study

¹Early contributions come from Dwyer, Schurr and Oh (1987) or Kalwani and Narayandas (1995).

the crucial ownership choice of integration versus outsourcing in such a dynamic global sourcing model.² However, they consider a setting with a fixed partner, and therefore cannot capture the evidence by Rauch and Watson (2003) and Besedes (2008) on trial-and-error search for efficient and reliable component manufacturers.

In our own related research, see Defever, Fischer and Suedekum (2016), we have introduced supplier search and re-matching into a dynamic property-rights model of global sourcing with relational contracts. In our repeated-game setup, the domestic-headquarter firm observes the efficiency of the current foreign supplier, and decides in every period whether to stick with that partner or to engage in a costly search for another, potentially better match. Moreover, in each round the firm can promise the current supplier a relational contract (RC) via an ex post bonus if the supplier does not behave opportunistically, but obeys to their ex ante agreement. Much as in Kukharsky (2015) or in Kamal and Tang (2015), we find that such an efficient RC can be implemented in equilibrium if agents are patient enough, and this RC then induces first-best effort levels and overcomes the underinvestment problem. Going beyond this insight, our model then makes clear that long-term collaborations (LTCs) and relational contracts (RCs) are two different concepts that sharply need to be distinguished. If agents are impatient, the firm may collaborate non-cooperatively with a supplier on a long-term basis, but they are never able to establish an efficient informal agreement. Moreover, our model shows that some firm only start a RC once they are satisfied with the efficiency of its current supplier. In other words, the endogenous decision to stop searching for a better partner and, thus, to launch a LTC *causes* the emergence of a RC. As this comes with increasing investment levels of both the headquarters and the supplier, our model is therefore fully consistent with the evidence by Rauch and Watson (2003) and Besedes (2008). More generally, our model's message is that the decision of *with whom* to interact is crucially important for the interrelated question of *how* the contractual nature of this interaction will look.

The aim of this chapter is to complement the analysis by Defever et al. (2016) in some important dimensions. In particular, we start from a stripped-down version of our model and completely assume away contract incompleteness. The investment levels of the two partners are fully verifiable and contractible, and hence, relational contracts are not needed to induce efficiency in this environment. This simplification lets us focus exclusively on the firm's search and matching problem. The trade-off that we study in this chapter is simple: The headquarters is currently matched with a supplier of a certain type, more specifically, with a certain unit cost level that determines the supplier's efficiency. The firm has an interest to collaborate with a supplier that is as efficient as possible, because this also raises the firm's own profits. But supplier switching is subject to a fixed cost for search and re-matching. The firm then essentially weighs this fixed costs against the expected gains from continued search, and will stop searching if it has – at some point – found a supplier that is efficient enough, i.e., with unit costs below a certain threshold level.

²There is also a more broadly related literature which has studied theoretically and empirically how relational agreements can overcome hold-up problems in the context of firm organization and input sourcing, see Macchiavello and Morjaria (2014), Board (2011) and Corts and Singh (2004).

While this basic formalization of the search-and-matching process is similar as in Defever et al. (2016), we consider two important generalizations in this chapter. First, rather than just assuming two supplier types (with high and low unit costs) as in that paper, we work here with a general function $g(c)$ that describes how unit costs are distributed across the mass of potential suppliers in the economy. Second, and even more importantly, we elicit the role of uncertainty when the firm searches for a new partner, and encounters a candidate supplier in its random search. The firm may or may not be able to infer the efficiency of this candidate before definitively giving up its previous match; if it is able to do so, it can separate the processes of search and re-matching, and decide to stick with its old partner if the encountered candidate is not more efficient than its current match. When the candidate's efficiency is only revealed after the first round of interaction, however, the firm will always re-match when it decides to search, and this corresponds to Rauch's notion of trial-and-error search for component suppliers. It turns out that both model versions, with search and re-matching being separable or non-separable, give rise to one specific prediction that we establish in Propositions 2 and 2' below. Namely, we show that the firm in our model tends to search more if the distribution $g(c)$ is more dispersed.³ The intuition is that, with more dispersion, more highly efficient matches are out there in the population of candidate suppliers, which in turn makes searching more attractive.

In the final section of this chapter, we then aim to address the empirical relevance of this novel theoretical prediction. The main challenge we are facing here is data availability. Our theory is about particular matches of domestic buyers and foreign suppliers, but current data only very rarely allows observing such matches. A notable exception is the paper by Eaton, Eslava, Jinkins, Krizan and Tybout (2014) who are able to construct pairs of Colombian exporting firms and their importers in the United States. They show that most exporters contract only with a single importing partner, thus suggesting that most trade is indeed relationship-specific. Moreover, they find considerable variation in relationship durations in the data, pointing to a co-existence of long-term and one-shot collaborations in that market. More recent studies which also use proprietary Census data on matches of US importers and foreign exporters include Monarch (2015) and Kamal and Tang (2015).

In this paper, we exploit Chinese customs data that encompass all export transactions to the United States, Japan, the United Kingdom and Germany between 2000 and 2006. We build a sample of firms that start exporting a particular product (6-digit industry level) to one of those destinations, but unfortunately we cannot identify the buyer. Still, following these fresh Chinese exporters over time, we observe which firms still export the same product to the same destination after a few years, as opposed to those that have terminated that exporting activity in the meantime. This allows us to build a proxy for whether a particular export transaction can be thought of as a long-term collaboration or rather a one-shot deal, and to construct the share of short-term relationships across Chinese HS-6 industries. We correlate this share with an industry-wide measure of cost dispersion across Chinese firms, which we construct from unit-value data.

³More specifically, a mean-preserving spread of $g(c)$ will increase the critical search cost level below which the firm decides to engage in supplier switching.

Our model tells us that this correlation between cost dispersion and the share of short-term collaboration should be positive, because headquarter firms tend to search more, and will thus exhibit more supplier turnover and one-shot deals in a given time frame. And indeed, the empirical evidence from our Chinese data is firmly in line with this theoretical prediction. More precisely, by aggregating the Chinese export transactions at the HS-6 product level, we consistently find a positive correlation between the share of one-shot transactions and the variance in the log of unit-value within an industry. This result is robust for exports across different destination markets, for independent Chinese firms as well as for Chinese foreign affiliates, as well as for different types of export transactions. Moreover, apart from that, this chapter also establishes some interesting stylized facts about the resilience of Chinese exporters that may be useful independently from the empirical test of our model.

The rest of this chapter is organized as follows. Section 2 describes our model, and Section 3 introduces the data and present the empirical results. Section 4 provides a brief conclusion.

2 The Model

2.1 Environment

Consider an industry where a single headquarter firm (labeled H) sources a customized intermediate input to produce a final good. There are many suppliers in that industry who could potentially manufacture this component, and those suppliers differ in their efficiency. In particular, we assume that there is a continuum of potential suppliers, which are all identical except for the constant marginal costs c that they would incur in the production of the input. Costs are distributed across potential suppliers according to some probability density function $g(c)$ with support $[\underline{c}, \bar{c}]$, where $\underline{c} > 0$. The corresponding cumulative density function $G(c)$ is continuously differentiable in c .

The headquarters initially gets matched with one supplier M^0 with unit costs c^0 randomly drawn from $g(c)$; this initial matching is costless. From the production and sale of the final output, the firm can appropriate a flow payoff $\pi(c)$ per period that is positive, differentiable and monotonically increasing in the supplier's efficiency, i.e., decreasing in marginal production costs: $\frac{\partial \pi}{\partial c} < 0$. Production is repeated ad infinitum, and the discrete time periods are indexed by $t = 0, 1, 2, \dots$, and $\delta > 0$ denotes the time discount factor.

In this model, the firm has an incentive to collaborate with a supplier that is as efficient as possible (with the lowest unit costs c), and this raises the static payoff per period as well as the discounted stream of future profits. Yet we assume that a switch of suppliers involves a fixed cost. More specifically, we start our analysis with the case where the firm can perfectly infer the efficiency of candidate suppliers, which allows it to separate the processes of search and re-matching. Every stage of our infinitely repeated interaction then has the following two consecutive steps:

1. **Transaction stage:** The current supplier delivers the input, final output is produced

and sold, and the headquarter firm receives the payoff $\pi(c)$, with $\partial\pi/\partial c < 0$.

2. **Re-matching stage:** The firm can decide to search for a new supplier. When deciding to search, it incurs a fixed cost $F > 0$. Let c^t be the unit cost level of its current supplier, and c^{t+1} the unit cost of the new supplier that it has encountered during its search. The cost c^{t+1} is randomly drawn from the distribution function $G(c)$ which is i.i.d. over periods. The firm can perfectly observe c^{t+1} before deciding to re-match. If the cost draw is such that $c^{t+1} < c^t$, the headquarters re-matches and continues the game with the new supplier. If $c^{t+1} \geq c^t$ it keeps its previous supplier.

Some comments about this simple setting are in order before we continue with the analysis. First, notice that our model is a stripped-down version of Defever, Fischer and Suedekum (2016). We abstract entirely from contractual frictions and focus solely on the firm's search and matching problem in an environment where, within a given match, investment choices are perfectly observable and contractible. To keep the exposition as parsimonious as possible, we furthermore stick to a reduced-form model that does not explicitly specify market demand, investment decisions, and revenue sharing. Those elements could be easily introduced, but they are not necessary to derive the main theoretical predictions that we later take to the data.

Furthermore, our model focuses entirely on the decisions of the headquarter firm H but it assigns a passive role to the suppliers. The firm, which cannot be in a relationship with more than one supplier at the same time, seeks for the best possible partner but the supplier itself (regardless of c) has no incentive to break up the collaboration but stays inside the relationship forever unless it is fired by the headquarters. We make this assumption, which could be justified by a zero (or sufficiently small) outside option for all potential suppliers, for convenience in order to focus on the re-matching incentives of the headquarter firm. In addition, this assumption allows us to abstract from the strategic incentives of the supplier. For this reason, we also do not need to specify the payoff for the supplier, but we simply assume that it is positive and, thus, better than the outside option. We also abstract from break-up costs, but one may interpret the search costs F at least partly as a compensation payment for the old supplier.

Finally, in the baseline version of our model studied in Section 2.2, we assume that the firm has the option to stick to its old supplier if it is not satisfied with the candidate supplier that it encounters during the search in any given period. On the one hand, this is realistic since suppliers have small (zero) outside options and would therefore always prefer to stay inside the relationship, even if the firm actively seeks for a better partner. On the other hand, there are also arguments in favor of an alternative assumption, namely that the firm *must* re-match if it decides to engage in searching. Specifically, this alternative scenario can be understood in terms of the information that is available to the firm about the candidate that it encounters in the search process. In the baseline version, the firm can perfectly observe its cost-efficiency before actually re-matching to it. This allows the firm to *separate* the steps of search and re-matching into two distinct decisions. In the second scenario, however, we may assume that the firm can only observe the costs of the newly drawn supplier *ex post*, i.e., only after the first period of

collaboration. When the candidate's costs are uncertain, the firm will always re-match when engaging in search (otherwise it would not pay F), and the search and re-matching decisions then become *non-separable*. We study this alternative setting below in Section 2.3.

2.2 Separable search and re-matching

Starting with the first scenario, denote by \tilde{c} the cost realization of the firm's current supplier, and by $\tilde{\pi}$ the firm's associated payoff per period. In order to see whether the firm engages in supplier search, given that $c^0 = \tilde{c}$, the expected payoff when engaging in supplier search, denoted by $E[\pi^S | c^0 = \tilde{c}]$, has to be larger than the continued payoffs with the current match, which is given by the infinite geometric series of the per-period payoff $\tilde{\pi}$. This can be expressed formally by the following supplier search incentive constraint:

$$E[\pi^S | c^0 = \tilde{c}] > \frac{\tilde{\pi}}{1 - \delta} \quad (\text{IC-Search})$$

Suppose it is profitable for the headquarters to search in period 0, i.e. (IC-Search) is fulfilled. It then pays F and with probability $1 - G(\tilde{c})$ draws a supplier that does not give it a payoff improvement when compared to the initial match. If that happens, it will stick to the initial supplier and search again at the end of the following round, because search must be optimal at the end of round 1 if it was optimal at the end of round 0 in this infinitely repeated game. With probability $G(\tilde{c})$, however, it finds a supplier with lower costs and, hence, re-matches. In the period after re-matching, the firm re-evaluates (IC-Search) with the cost-realization of the new supplier and thereby determines whether it will engage in further search for an even more efficient partner, and so on.

The expected search payoff $E[\pi^S | c^0 = \tilde{c}]$ can be formally expressed as the solution to the following program:

$$\begin{aligned} V_0 &= \tilde{\pi} - F + \delta V_1 \\ V_t &= Pr(c < \tilde{c}) \frac{E[\pi | c < \tilde{c}]}{1 - \delta} + Pr(c \geq \tilde{c}) (\tilde{\pi} - F + \delta V_{t+1}), \quad t = 1, 2, \dots \end{aligned}$$

The first equation captures the value of the relationship at time 0, denoted V_0 , which includes the option value of searching in period 1. This search, formalized by the second equation, leads to a payoff improvement (and, thus, to an actual re-matching) only with a certain probability. With the respective counter-probability, however, the search is unsuccessful in period 1, in which case it continues in period 2, and so on.

It is important to understand that the expression (IC-Search) is conditional on the current cost realization, i.e., a new version of it has to be formulated for every supplier that the firm encounters in the sequence of re-matches that it performs. Then, observing that the decision problem is the same in every round where the headquarters is still matched with \tilde{c} and hence is on the search for a better partner, we can set $V_t = V_{t+1}$ for $t = 1, 2, \dots$ and simplify the program

to:

$$\begin{aligned} V_0 &= \tilde{\pi} - F + \delta V_1 \\ V_1 &= \frac{1}{1 - \delta(1 - G(\tilde{c}))} \left[(1 - G(\tilde{c}))(\tilde{\pi} - F) + \frac{G(\tilde{c})E[\pi | c < \tilde{c}]}{1 - \delta} \right] \end{aligned} \quad (1)$$

Using (1) in (IC-Search), we can then solve for F and obtain a search cost threshold $\bar{F}(\tilde{c})$ that determines whether or not the headquarter firm will engage in supplier search:

$$F < \frac{\delta}{1 - \delta} G(\tilde{c}) [E[\pi | c < \tilde{c}] - \tilde{\pi}] \equiv \bar{F}(\tilde{c}) \quad (2)$$

Expression (2) demonstrates that the fixed costs F must be low enough in order for search to occur. Notice that the critical search cost level $\bar{F}(\tilde{c})$ depends positively on δ , i.e., search is more attractive the more patient agents are. The intuition is that future payoffs matter more, and hence it becomes more important to find an efficient supplier with low costs. Moreover, $\bar{F}(\tilde{c})$ depends positively on $G(\tilde{c})$ and on the term $[E[\pi | c < \tilde{c}] - \tilde{\pi}]$. This shows that search is more likely the higher the chance to find a more efficient supplier, and the larger the headquarters' expected per-period payoff difference when comparing a better with the current supplier.

We summarize the result in the following proposition:

Proposition 1. *Suppose the headquarters is currently matched with a supplier with marginal input production costs \tilde{c} . Then the headquarters will engage in the search for a better supplier whenever $F < \bar{F}(\tilde{c})$ and stick to the current match otherwise.*

Note that whenever (IC-Search) is fulfilled, the interaction of the firm with its current supplier is short-term in the sense that it will end in the first period where the firm draws a supplier with $c' < \tilde{c}$. To be precise, this means that, while being engaged in search, the firm will be matched with the current supplier for only a finite number of periods. The relationship with a particular supplier will only be long-term (i.e. durable in all future periods) when the probability of finding a better partner together with the size of possible payoff improvements, does not justify the per-period search expenses, which is the case precisely when (IC-Search) does not hold anymore.

Essentially, the search decision is driven by the comparison of search costs and expected gains from search. The latter of which hinge crucially on the properties of the distribution $g(c)$, and on where in the distribution the current supplier is located. In the following, we derive an important result that relates the firm's propensity to search with the cost distribution of potential suppliers:

Proposition 2. *A mean-preserving spread (MPS) on the distribution function $G(c)$ of potential suppliers increases the firm's propensity to search for a more cost-efficient supplier by increasing the critical search cost level $\bar{F}(\tilde{c})$ for all possible values of \tilde{c} .*

To prove Proposition 2 formally, let $\hat{G}(c)$ be a MPS of $G(c)$, and let $\hat{g}(c)$ be its associated probability density function. Denote by \hat{F} the search cost threshold resulting under $\hat{G}(c)$ and by \bar{F} that for $G(c)$. Then for any $c^0 = \tilde{c}$ it must hold that

$$\hat{F}(\tilde{c}) > \bar{F}(\tilde{c}). \quad (3)$$

Substituting (2) into (3) and simplifying gives

$$\hat{G}(\tilde{c}) (E_{\hat{G}}[\pi | c < \tilde{c}] - \tilde{\pi}) > G(\tilde{c}) (E_G[\pi | c < \tilde{c}] - \tilde{\pi}).$$

Plugging in the conditional expected values and simplifying gives

$$\int_{\underline{c}}^{\tilde{c}} \pi(c) \hat{g}(c) dc - \hat{G}(\tilde{c}) \tilde{\pi} > \int_{\underline{c}}^{\tilde{c}} \pi(c) g(c) dc - G(\tilde{c}) \tilde{\pi}.$$

Integrating by parts allows us to simplify further:

$$\int_{\underline{c}}^{\tilde{c}} \frac{\partial \pi}{\partial c} G(c) dc > \int_{\underline{c}}^{\tilde{c}} \frac{\partial \pi}{\partial c} \hat{G}(c) dc. \quad (4)$$

Because $\frac{\partial \pi}{\partial c} < 0$, it follows from the definition of the MPS that expression (4) is true for any value of \tilde{c} . From this we can conclude that a more dispersed cost distribution in an industry increases the headquarters' search and re-matching propensity.

Intuitively, a MPS can be interpreted as a shift of probability mass towards the ends of the probability density function $g(c)$. With relatively more probability mass close to the most efficient supplier at \underline{c} , both $G(\tilde{c})$ as well as $E[\pi | c < \tilde{c}]$ increase which in turn shifts $\bar{F}(\tilde{c})$ upwards. This makes it more attractive to search for a partner with lower costs, essentially because the firm only re-matches when it can improve its per-period payoff. An MPS adds efficient suppliers to the tail of the distribution and generates such a payoff increase.

2.3 Non-separable search and re-matching

In the previous subsection, we have assumed that the headquarters can perfectly observe the cost-efficiency of the candidate supplier before actually starting a relationship with it. Obviously, this assumption is a strong one. We now study the alternative scenario where the firm only learns a supplier's cost-efficiency once it is matched to it (while still having perfect information on the distribution $G(c)$). Whenever the headquarters cannot observe the cost-level of the suppliers that it encounters during the search, it will re-match to the new supplier whenever it decides to engage in search. Otherwise the firm could refrain from searching in the first place, and save the fixed costs F . We reformulate step 2 of the stage game as follows:

- 2'. **Re-matching stage:** The firm decides whether to re-match with a new supplier. When deciding to re-match, it incurs a fixed cost $F > 0$. Let c^t be the unit cost level of its

current supplier, and c^{t+1} the unit cost of the new supplier that it has encountered during its search. The cost c^{t+1} is randomly drawn from the distribution function $G(c)$ which is i.i.d. over periods. The firm cannot observe c^{t+1} before the start of a transaction relationship and hence re-matches whenever it searches.

We will first assess the robustness of Proposition 1 for this case where search and re-matching cannot be separated. Again, we denote by \tilde{c} and $\tilde{\pi}$ the current supplier's costs and the firm's associated payoffs, respectively. Equivalently to (IC-Search), the firm trades off the expected payoffs when engaging in re-matching with the continuation payoff resulting from its current relationship. Formally:

$$E[\pi^{RM} | c^0 = \tilde{c}] > \frac{\tilde{\pi}}{1 - \delta} \quad (\text{IC-ReMatch})$$

$E[\pi^{RM} | c^0 = \tilde{c}]$ follows the same logic as $E[\pi^S | c^0 = \tilde{c}]$ from above, but differs in one respect. Suppose that, after paying the search costs F , the firm draws a supplier with $c > \tilde{c}$, which happens with probability $1 - G(\tilde{c})$. The firm cannot observe these costs directly, but only when arriving at the transaction stage of the following period; thus, from an ex-ante perspective, the firm will receive the expected payoff with this worse supplier, $E[\pi | c \geq \tilde{c}]$, for one period and then continue searching. Formally, then, we can then express $E[\pi^{RM} | c^0 = \tilde{c}]$ as the solution to the following system:

$$\begin{aligned} V_0 &= \tilde{\pi} - F + \delta V_1 \\ V_t &= Pr(c < \tilde{c}) \frac{E[\pi | c < \tilde{c}]}{1 - \delta} + Pr(c \geq \tilde{c}) (E[\pi | c \geq \tilde{c}] - F + \delta V_t), \quad t = 1, 2, \dots \end{aligned}$$

We can bring the program into the non-recursive form:

$$\begin{aligned} V_0 &= \tilde{\pi} - F + \delta V_1 \\ V_1 &= \frac{1}{1 - \delta(1 - G(\tilde{c}))} \left[(1 - G(\tilde{c})) (E[\pi | c \geq \tilde{c}] - F) + \frac{G(\tilde{c}) E[\pi | c < \tilde{c}]}{1 - \delta} \right] \end{aligned}$$

By a method analogous to what is above, we can now solve (IC-ReMatch) for F and obtain a threshold level $\tilde{F}(\tilde{c})$ that determines whether the firm will engage in supplier re-matching or not:

$$F < \frac{\delta}{1 - \delta} \left[\left(E[\pi] - \tilde{\pi} \right) - (\delta(1 - G(\tilde{c})) \left(E[\pi | c > \tilde{c}] - \tilde{\pi} \right)) \right] \equiv \tilde{F}(\tilde{c}) \quad (5)$$

Expression (5) demonstrates that search costs F have to be low enough in order for re-matching to occur. Again, $\tilde{F}(\tilde{c})$ depends positively on the patience level δ . Notice that because $(E[\pi | c > \tilde{c}] - \tilde{\pi})$ is always negative, $(E[\pi] > \tilde{\pi})$ is a sufficient condition for $\tilde{F}(\tilde{c}) > 0$; that is, having a below-average supplier is always enough for re-matching to occur if search costs are small enough.

We summarize the result in an adjusted version of Proposition 1:

Proposition 1’. *Suppose the headquarters is currently matched with a supplier with marginal input production costs \tilde{c} . When the headquarters can infer the cost-efficiency of a candidate supplier only ex-post to contracting, it will engage in re-matching whenever $F < \tilde{F}(\tilde{c})$ and stick to the current match otherwise.*

Note that, whenever (IC-ReMatch) is fulfilled, the relationship with the current supplier is short-term as it was also the case in the previous subsection when search and re-matching decisions could be separated by the firm. In the present setting, however, the relationship will definitely exist for only one period of the game, while in the previous case, the current firm-supplier pairing would last until a more efficient supplier was found. Finally, when encountering a better supplier in the re-matching sequence with $c' < \tilde{c}$ for which $F \geq \tilde{F}(c')$ turns out to hold, the re-matching process stops and the firm will engage in a long-term relationship with this supplier.

Turning to Proposition 2, we now study again the effects of an increase in dispersion of the distribution $g(c)$ on the firm’s propensity to re-match:

Proposition 2’. *Consider the case where the candidate supplier’s cost-efficiency is only observable ex-post to the firm. A mean-preserving spread $\dot{G}(c)$ on the distribution function $G(c)$ increases the firm’s propensity to engage in re-matching by increasing the critical search cost level $\tilde{F}(\tilde{c})$ for all possible values of \tilde{c} .*

The proof of Proposition 2’ goes along the same lines as in the previous case. Let $\dot{G}(c)$ be an MPS of $G(c)$, and $\dot{g}(c)$ its associated probability density function. Denote by \dot{F} the search cost threshold resulting under $\dot{G}(c)$ and by \tilde{F} the search cost threshold for $G(c)$. Then for Proposition 2’ to be true, it must hold for any $c^0 = \tilde{c}$ that

$$\dot{F}(\tilde{c}) > \tilde{F}(\tilde{c}). \quad (6)$$

Substituting (5) into (6) and simplifying gives:

$$(1 - \dot{G}(\tilde{c}))(E_{\dot{G}}[\pi | c > \tilde{c}] - \tilde{\pi}) < (1 - G(\tilde{c}))(E_G[\pi | c > \tilde{c}] - \tilde{\pi})$$

We now plug in the conditional expected values and simplify:

$$\int_{\tilde{c}}^{\tilde{c}} \pi(c)\dot{g}(c) dc - (1 - \dot{G}(\tilde{c}))\tilde{\pi} < \int_{\tilde{c}}^{\tilde{c}} \pi(c)g(c) dc - (1 - G(\tilde{c}))\tilde{\pi}$$

Integrating by parts we arrive at:

$$\int_{\tilde{c}}^{\tilde{c}} \frac{\partial \pi}{\partial c} \dot{G}(c) dc > \int_{\tilde{c}}^{\tilde{c}} \frac{\partial \pi}{\partial c} G(c) dc \quad (7)$$

Because $\dot{G}(c)$ is a MPS of $G(c)$, and $\frac{\partial \pi}{\partial c} < 0$, it follows that expression (7) is true for all values of \tilde{c} . This concludes the proof.

2.4 Summary of theoretical results and main prediction

Summing up, the two model versions – with search and re-matching being separable or non-separable – yield qualitatively similar results and predictions. In order for search/re-matching to occur, the fixed costs F must be below a threshold level. This cutoff is stricter (the threshold for F is lower), the lower is the unit cost of the current supplier. That is, the better the current supplier is already, the lower is the chance that the firm will continue looking for a better match. Moreover – turning to the key prediction of our model – both model versions show that more cost dispersion among the potential suppliers in the industry lowers this search threshold, and thus leads to more search/re-matching, because it leads to an expected gain from searching as better potential matches are now out there in the industry.

Figure 1 gives a graphical summary. Making the additional structural assumption that $\frac{\partial^2 \pi}{\partial \tilde{c}^2} < 0$, we can draw \bar{F} as a function that is increasing and convex in \tilde{c} . Then, the model predicts that for all (\tilde{c}, F) -combinations that lie below this function the firm will engage in search for a better supplier and hence be in a short-term collaboration (STC) with its current partner (blue area of the Figure). Otherwise, for all (\tilde{c}, F) -combinations above that convex curve, the firm will not search, but will establish a long-term collaboration (LTC) with its current match. Performing an MPS on the cost distribution $g(c)$ shifts the \bar{F} -function upwards to \hat{F} , and thereby enlarges the parameter range for which the firm will engage in supplier search / re-matching.




Figure 1: Illustration of search threshold and MPS

Our results do not differ qualitatively between the two cases, which underscores that our predictions are robust with respect to the quality with which information on supplier costs is transmitted to the headquarters. In particular, it indicates that the prediction that a larger cost dispersion within an industry leads to a higher search and re-matching propensity is not

crucially influenced by how quickly firms learn about the supplier’s efficiency. This gives us confidence in our industry-level prediction of a positive relationship between the cost-dispersion of potential suppliers and firms’ propensity to engage in supplier search and re-matching.

3 Empirical Analysis

In this section, we describe how we construct proxies for the two variables that are required to empirically test the key prediction of our theoretical model: i) whether a buyer-supplier match is an STC as opposed to an LTC, and ii) the industry-wide cost dispersion of potential suppliers.

3.1 Data and variables

Export transactions of fresh Chinese suppliers: The empirical evaluation of the model exploits transaction-level customs data from the Chinese General Administration of Customs for the period 2000-2006. For each year, the data allow us to identify the manufacturing export transactions of Chinese firms, where an export transaction is defined as a firm-product-destination combination. We drop all transactions with zero or negative values as well as export transactions with destination China. Moreover, to ensure consistency of the product categorization over time, we use the conversion table from UN Comtrade and convert the product code used for the years 2000 and 2001 into the HS 2002 codes.

In our theoretical framework, the headquarter deals with a new supplier when it decides to re-match. To reflect this, we only consider transactions of *fresh* Chinese exporters that establish their first exporting activity, so that foreign importing firms cannot infer any quality signals from their previous exporting experiences. Moreover, we restrict our sample to new transactions realized in the United States, Japan, United Kingdom and Germany, which are the top four main destinations for Chinese exports besides Hong-Kong and Macao. Specifically, our sample is composed of Chinese firms i that export a product j to a country c in year t , where $t = \{2001, 2002, 2003\}$, but that have not exported *anything* to *any* destination outside China during the previous years since 2000. We exclude all industries j with fewer than ten active firms, and those where control variables are not available.

An important observation, as shown in the upper panel of table 1, is that quite a few of the new Chinese exporters in our data are owned by multinational corporations. This includes Chinese foreign affiliates that are dependent subsidiaries or part of a joint venture. Altogether, the final sample includes 28,370 fresh Chinese exporters (15,017 domestic firms and 13,353 foreign affiliates), starting a total of 166,032 new export transactions, and spans 1,192 different HS-6-digit product categories. In our baseline estimations we include all these observations, but we also conduct robustness checks that focus only on independent Chinese exporters.

Regarding the destination markets, while our full sample pools transaction with the United States, Japan, the United Kingdom and Germany, we also present results for each market independently, and the lower panels of table 1 show the number of Chinese exporters, exported

products and export transactions in those four markets. In addition, for our baseline estimations we keep all types of goods, and in the robustness checks we limit this sample to complex goods with a relatively high degree of relationship-specificity, as this might better reflect the context of our theoretical model.

Table 1: Firms, products and destinations

	All destinations	
	Domestic firms	Foreign affiliates
# firms	15017	13353
# products	1192	1166
# transactions	122707	43325
	All firms	
	US	Japan
# firms	16150	13683
# products	1004	861
# transactions	63564	48877
	All firms	
	UK	Germany
# firms	5996	6606
# products	452	478
# transactions	16531	16427

Supplier turnover: Following the fresh Chinese exporters over time, we observe which of the firms that started an exporting activity in country c and product j in year t , terminate their export activity in year $t + 3$, as opposed to those that still export the *same* product j to the *same* country c after three years. Ideally, we would like to observe whether the Chinese exporter still deals with the same importer, but the data do not allow us to do so. Nevertheless, looking specifically at the United States as the country of destination in table 3, our duration measure provides a pattern very similar to the one observed by Monarch (2015), who is able to use confidential US Customs data on US import transactions from China, where firms on both sides are uniquely identified. He finds that 45% of US importers change their Chinese partner from one year to the next. This number is quite close to the one obtained in our data (43.4%).

Aggregating to the HS-6 product level, we compute the share of transactions in industry j that got started in $t = \{2001, 2002, 2003\}$ and that have stopped in year $(t + 3)$. This gives us our measure on the supplier turnover in industry j , which we use as our main explanatory variable: the share of short-term transactions.

Industry-wide cost dispersion: To construct an empirical proxy for the industry-wide cost dispersion, we approximate the transaction-specific marginal cost of the respective supplier by the unit-value of each transaction. The cost dispersion measure for industry j is then given by the standard deviation of the log of unit-values within the respective HS-6 industry during our period of observation. As usual, marginal costs are difficult to observe directly. Thus, our

underlying assumption, follows standard practise and postulates that Chinese suppliers engage in fixed-markup pricing, so that the variation in unit-values captures cost dispersion. The recent findings by Monarch (2015) support this approach. Using confidential customs data, he shows that US importers paying the highest prices to their Chinese suppliers (also proxied by the unit-value in his paper) are most likely to change their partner in the future. This is consistent with Proposition 1 of our model, where re-matching is more likely the higher the marginal costs of the current supplier.

Manova and Zhang (2012) argue that differences in unit-values could also reflect quality variations. In this case, our empirical proxy would be capturing heterogeneity in suppliers' quality. However, under iso-elastic demand and monopolistic competition, marginal costs and quality differences are isomorphic in the sense that they enter equilibrium firm revenue in exactly the same way (see Melitz and Redding, 2014). As a result, while a lower quality would have the opposite effect on re-matching compared to lower cost, decreasing the variance of either quality or cost would have the same impact on the decrease of supplier turnover. Hence, the interpretation of our results would be robust to the way we interpret differences in unit-values.

3.2 Descriptive evidence

Table 2 reports the duration of the transactions for independent firms and for foreign affiliates. It is important to take these ownership structures into account, as Chinese foreign affiliates may differ systematically in their economic behavior from independent contractors. While the ownership dimension does not feature in our theoretical model, the key trade-off formalized by our model is – in principle – not limited to arm's-length outsourcing relationships, but can also arise with a similar intuition in the context of intra-firm trade.

Looking at table 2, of the 122,707 transactions made by domestically-owned firms, it turns out that 59,361 (48.4%) are terminated after less than one year, while 34,175 (27.9%) of them endure for more than three years. While transactions involving foreign affiliates tend to be more resilient, still 28.2% of the transactions stop after the first year. Digging deeper into the different countries of destinations (the United States, Japan, the United Kingdom and Germany), table 3 shows that the share of one-year transactions within these four subgroups is rather similar across destinations (from 40.4% for Japan to 45.8% for the United Kingdom).

Table 2: Duration of export transactions by ownership

Duration of transactions	Domestic firms		Foreign affiliates	
	# transactions	Percentage	# transactions	Percentage
< 1 year	59,361	48.4%	12,201	28.2%
> 1 year	17,970	14.6%	5,814	13.4%
> 2 years	11,201	9.1%	4,512	10.4%
> 3 years	34,175	27.9%	20,798	48.0%
Total	122,707	100.0%	43,325	100.0%

Table 3: Duration of export transactions by export destination

Duration of transactions	Exports to the US		Exports to Japan	
	# transactions	Percentage	# transactions	Percentage
< 1 year	27,572	43.4%	19,764	40.4%
> 1 year	9,498	14.9%	6,963	14.2%
> 2 years	6,108	9.6%	4,759	9.7%
> 3 years	20,402	32.1%	17,399	35.6%
Total	63,580	100.0%	48,885	100.0%

Duration of transactions	Exports to the UK		Exports to Germany	
	# transactions	Percentage	# transactions	Percentage
< 1 year	7,577	45.8%	7,068	43.0%
> 1 year	2,361	14.3%	2,231	13.6%
> 2 years	1,477	8.9%	1,534	9.3%
> 3 years	5,118	31.0%	5,594	34.1%
Total	16,533	100.0%	16,427	100.0%

Ordinary least squares (OLS) regressions. In table 4, we present some OLS results at the industry level. To do this, we first aggregate the data to the HS-6 product-destination level. Specifically, we use all new transactions by fresh Chinese exporting firms i in industry j to country c . We create a dummy variable indicating whether the transaction is a short-term transaction (value 1) or if it is still observed in $t+3$ (value 0), and may thus be classified as LTC. We then obtain our dependent variable by aggregating the data at the product-destination level by computing the share of short-term transactions among all the transactions within an HS-6 product. All our baseline regressions include a set of destination fixed effects δ_c .

$$\text{Share short-term transactions}_{jc} = a_0 + a_1 \text{Cost Dispersion}_{jc} + \text{controls}_j + \delta_c + \epsilon_{jc}$$

In column 1, we regress our dependent variable on the within-product standard deviation of the log of unit prices, while controlling for the ownership status of the Chinese firm. We find a positive correlation between the share of short-term transactions and the within-industry spread in unit value, as expected from Proposition 2 of our theoretical model. In column 2, we introduce a set of additional industry-level control variables, namely the capital-labor ratio, the human capital intensity, and the R&D-sales ratio, which are standard measures for headquarter-intensity. Moreover, we include the relationship-specificity index by Nunn (2007), and add demand elasticities from Broda and Weinstein (2006).⁴ Adding these control variable leaves the main coefficient unchanged.

In Columns 3 and 4 of table 4, we exclude all the export transactions of foreign affiliates before aggregating the data at the product level. This leads to a decreasing number of industries

⁴As data are disaggregated at the HS 10 digit level, we use the median value for each HS-6 product.

as we still exclude all products with fewer than ten observations. The main coefficient remains very similar and remains significant at the 1% level. Overall, this indicates that our results remain robust also in a smaller sample consisting only of independent Chinese-owned component manufacturers.

Table 4: Relational-contracts and long-term collaborations: Transaction-level

Dependent variable: Share of one-shot transactions	Full sample		Excluding foreign affiliates	
Ln unit value SD	0.050 ^a (0.009)	0.072 ^a (0.010)	0.065 ^a (0.010)	0.065 ^a (0.011)
Foreign ownership	-0.207 ^a (0.024)	-0.199 ^a (0.023)		
ln capital/worker		0.001 (0.009)		-0.021 ^b (0.010)
ln skilled/worker		0.062 ^a (0.021)		-0.010 (0.026)
ln R&D sales		-0.009 ^b (0.004)		-0.003 (0.005)
Nunn measure		0.034 (0.042)		0.065 (0.056)
ln demand elasticity		0.007 (0.007)		-0.000 (0.008)
Country fixed effects	Yes	Yes	Yes	Yes
Observations	3,456	3,456	1,029	1,0290
R^2	0.058	0.074	0.042	0.053

Notes: Dependent variable: Share of export transactions which are one-shot on the overall number of transactions within an HS-6 product. Robust standard-errors in brackets using clustered standard errors at the HS-6 level.

Next, in table 5 we report results separately for transactions realized in each of the top-four export destinations, that is the United States, Japan, the United Kingdom and Germany. Our main coefficient of interest remains stable in all specifications, and is always significant at the 1% level, regardless of whether we include or exclude foreign affiliates or whether we consider the additional industry-level control variables.

We find that, taken together, the results in tables 4 and 5 are firmly in line with the predictions of our theoretical model— more specifically, with Proposition 2, which states that higher cost dispersion among the suppliers leads to more search and re-matching in that industry, and thus negatively affects observed match durations.

3.3 Robustness checks and discussion

Table 6 provides a battery of robustness checks. The first one consists of dropping all transactions that include processing trade. As shown in Defever, Fischer and Suedekum (2016), processing transactions are usually associated with high buyer-supplier interrelationship and

Table 5: Relational-contracts and long-term collaborations: Industry-level

Dependent variable: Share of one-shot transactions	Country of destination			
	US		Japan	
ln unit value SD	0.050 ^a (0.009)	0.065 ^a (0.010)	0.049 ^a (0.010)	0.088 ^a (0.011)
Foreign ownership	-0.238 ^a (0.034)	-0.235 ^a (0.033)	-0.151 ^a (0.029)	-0.194 ^a (0.028)
ln capital/worker		-0.004 (0.010)		-0.043 ^a (0.011)
ln skilled/worker		0.030 (0.022)		0.117 ^a (0.029)
ln R&D sales		-0.008 ^c (0.004)		-0.010 ^c (0.005)
Nunn measure		-0.047 (0.058)		0.119 ^b (0.058)
ln demand elasticity		0.007 (0.008)		0.010 (0.008)
Observations	1,004	1,004	861	861
R^2	0.063	0.077	0.051	0.129

Dependent variable: Share of one-shot transactions	Country of destination			
	UK		Germany	
ln unit value SD	0.047 ^b (0.018)	0.058 ^a (0.021)	0.036 ^b (0.016)	0.062 ^a (0.018)
Foreign ownership	-0.258 ^a (0.045)	-0.232 ^a (0.047)	-0.187 ^a (0.049)	-0.159 ^a (0.048)
ln capital/worker		0.035 ^b (0.017)		0.012 (0.014)
ln skilled/worker		0.018 (0.033)		0.118 ^a (0.033)
ln R&D sales		-0.008 (0.007)		0.002 (0.007)
Nunn measure		-0.151 (0.098)		0.005 (0.106)
ln demand elasticity		0.013 (0.017)		0.004 (0.012)
Observations	452	452	478	478
R^2	0.075	0.098	0.036	0.064

Notes: Dependent variable: Share of export transactions which are one-shot on the overall number of transactions within an HS-6 product. Robust standard-errors in brackets using clustered standard errors at the HS-6 level.

relational contracts. Column 1 shows that the main coefficient for the within-sector price dispersion remains extremely similar to our baseline results.

In column 2 we only keep industries where the Nunn-measure is above the median, and in column 3 we limit the sample even further to industries in the upper quartile of specificity. In fact, our baseline estimates are for all 1,004 product categories, some of which are rather

Table 6: Relational-contracts and long-term collaborations: Transaction-level

Dependent variable: Share of one-shot transactions	Drop Processing	Nunn measure > median	Nunn measure > 75 percentile	Main sector
One shot dummy	0.069 ^a (0.010)	0.067 ^a (0.012)	0.045 ^a (0.015)	0.068 ^a (0.013)
Foreign ownership	-0.123 ^a (0.037)	-0.253 ^a (0.041)	-0.263 ^a (0.056)	-0.062 ^b (0.029)
ln capital/worker	-0.023 ^b (0.009)	0.016 (0.015)	0.010 (0.021)	-0.011 (0.012)
ln skilled/worker	0.024 (0.025)	0.048 (0.031)	0.018 (0.040)	0.022 (0.031)
ln R&D sales	-0.014 ^a (0.005)	-0.012 ^b (0.006)	-0.021 ^b (0.010)	-0.022 ^a (0.006)
Nunn measure	0.058 (0.050)	2.586 ^a (0.678)	4.727 ^a (1.071)	0.064 (0.059)
ln demand elasticity	0.011 (0.007)	0.016 ^c (0.009)	0.017 (0.011)	0.025 ^a (0.009)
Observations	1,101	602	380	904
R^2	0.064	0.143	0.142	0.068

Notes: Dependent variable: Share of export transactions which are one-shot on the overall number of transactions within an HS-6 product. Robust standard-errors in brackets using clustered standard errors at the HS-6 level.

standardized goods or final products. Since our model features the duration of a relationship between buyer and supplier when producing a specific intermediate input, we now limit the sample and – rather than just adding the Nunn measure as a control variable – focus on product categories with a high degree of relationship-specificity. As can be seen, all results remain qualitatively unchanged, even if the point estimates become slightly smaller in Column 3. This, even when focussing on highly relationship-specific goods, we find that the within-sectoral cost-spread increases the proportion of short-term transactions.

In column 4, since many firms export multiple products, we exclude all the transactions that are not part of the firm’s main sector of activity which we define by the 2-digit HS product for which we observe most exported HS-6 products. We only keep the transactions that belong to this main 2-digit sector of each firm, and again we find that the coefficients are very similar to our baseline estimates.

We have also conducted all these robustness checks separately for the different destination markets, along the lines of table 5, and we have found that the pattern remains very stable in all these specifications. That is, in line with Proposition 2 of our model, we find a robust positive correlation between the industry-wide measure of cost dispersion and the share of short-term export collaborations in that industry, indicating that headquarters tend to engage in more supplier search when there is more cost dispersion among potential suppliers.

4 Conclusion

In this paper, we develop a dynamic search and matching model of global sourcing where a domestic headquarters seeks to obtain an intermediate input from a foreign input supplier. We show that the option to search and re-match with a new supplier crucially depends on an industry's characteristics. More precisely, the industry-wide cost dispersion of potential suppliers affects the expected search payoff. In industries with a large cost dispersion among suppliers, headquarters tend to search more, and will thus exhibit more supplier turnover and one-shot deals in a given time frame. This is true both in the case where search and re-matching are separable, and when they are non-separable.

This key prediction of our model is empirically supported in the context of export transactions of fresh Chinese suppliers to the United States, Japan, the United Kingdom and Germany, the top four main destinations for Chinese exports besides Hong-Kong and Macao. By aggregating the Chinese export transactions at the HS-6 product level, we show that the share of short-term transactions increase with an within-industry cost dispersion, which we construct from unit-value data. This result is robust for exports across different destination markets, for independent Chinese firms, as well as for Chinese foreign affiliates.

More generally, this chapter opens several promising avenues for future research that have not been explored yet, in particular by documenting novel stylized facts about the resilience of product-destination export transactions.

References

- [1] AEBERHARDT, R., BUONO, I., AND FADINGER, H. Learning, Incomplete Contracts and Export Dynamics: Theory and Evidence from French Firms. *European Economic Review* 68 (2014), 219–249.
- [2] ANTRÀS, P. Firms, Contracts, and Trade Structure. *The Quarterly Journal of Economics* 118, 4 (2003), 31–77.
- [3] ANTRÀS, P. *Global Production: Firms, Contracts and Trade Structure*. Princeton University Press, 2015.
- [4] ANTRÀS, P., AND COSTINOT, A. Intermediated Trade. *The Quarterly Journal of Economics* 126 (2011), 1319–1374.
- [5] ANTRÀS, P., AND HELPMAN, E. Global Sourcing. *Journal of Political Economy* 112, 3 (2004), 552–580.
- [6] ARAUJO, L., MION, G., AND ORNELAS, E. Institutions and Export Dynamics. *Journal of International Economics* 98 (2016), 2–20.
- [7] BAKER, G., GIBBONS, R., AND MURPHY, K. J. Relational Contracts and the Theory of the Firm. *The Quarterly Journal of Economics* 117, 1 (2002), 39–84.

- [8] BERNARD, A. B., MOXNES, A., AND ULLTVEIT-MOE, K. H. Two-Sided Heterogeneity and Trade. Working Paper 20136, National Bureau of Economic Research, May 2014.
- [9] BESEDEŠ, T. A Search Cost Perspective on Formation and Duration of Trade. *Review of International Economics* 16, 5 (2008), 835–849.
- [10] BOARD, S. Relational Contracts and the Value of Loyalty. *American Economic Review* 101, 7 (2011), 3349–67.
- [11] BRODA, C., AND WEINSTEIN, D. E. Globalization and the Gains from Variety. *The Quarterly Journal of Economics* 121, 2 (2006), 541–585.
- [12] CARBALLO, J., OTTAVIANO, G., AND VOLPE-MARTINCUS, C. The Buyer Margins of Firms’ Exports. Discussion Paper 1234, CEP, 2013.
- [13] CORTS, K. S., AND SINGH, J. The Effect of Repeated Interaction on Contract Choice: Evidence from Offshore Drilling. *Journal of Law, Economics, and Organization* 20, 1 (2004), 230–260.
- [14] DEFEVER, F., FISCHER, C., AND SUEDEKUM, J. Relational contracts and supplier turnover in the global economy. *Journal of International Economics* 103 (2016), 147–165.
- [15] DWYER, R. F., SCHURR, P. H., AND OH, S. Developing Buyer-Seller Relationships. *The Journal of Marketing* 51, 2 (1987), 11–27.
- [16] EATON, J., ESLAVA, M., JINKINS, D., KRIZAN, C. J., AND TYBOUT, J. A Search and Learning Model of Export Dynamics. Working Paper, 2014.
- [17] GROSSMAN, S. J., AND HART, O. D. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *The Journal of Political Economy* 94, 4 (1986), 691–719.
- [18] HART, O. D., AND MOORE, J. Property Rights and the Nature of the Firm. *Journal of Political Economy* 98, 6 (1990), 1119–58.
- [19] KALWANI, M. U., AND NARAYANDAS, N. Long-Term Manufacturer-Supplier Relationships: Do They Pay off for Supplier Firms? *The Journal of Marketing* 59, 1 (1995), 1–16.
- [20] KAMAL, F., AND TANG, H. Relational Trade: Evidence from the United States. Mimeo, 2015.
- [21] KUKHARSKYY, B. Quality of Institutions and Global Sourcing. Mimeo, 2015.
- [22] KUKHARSKYY, B., AND PFLÜGER, M. Relational Contracts and the Economic Well-Being of Nations. Discussion Paper, BGPE, 2011.
- [23] MACCHIAVELLO, R., AND MORJARIA, A. Competition and Relational Contracts: Evidence from Rwanda’s Coffee Mills. Working Paper, 2014.

- [24] MANOVA, K., AND ZHANG, Z. Export Prices Across Firms and Destinations. *The Quarterly Journal of Economics* 127 (2012), 379–436.
- [25] MELITZ, M. J., AND REDDING, S. J. Heterogeneous Firms and Trade. *Handbook of International Economics* 4 (2014), 1–54.
- [26] MONARCH, R. It's Not You, It's Me: Breakups in U.S.-China Trade Relationships. Working paper, 2015.
- [27] NUNN, N. Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade. *The Quarterly Journal of Economics* 122, 2 (2007), 569–600.
- [28] RAUCH, J. E. Networks versus Markets in International Trade. *Journal of International Economics* 48, 1 (1999), 7–35.
- [29] RAUCH, J. E., AND WATSON, J. Starting Small in an Unfamiliar Environment. *International Journal of Industrial Organization* 21, 7 (2003), 1021–1042.