Relational Contracts and Supplier Turnover in the Global Economy

Fabrice Defever* Christian Fischer† Jens Suedekum‡

August 2015

Abstract

Headquarters and their specialized component suppliers have a vital interest in establishing long-term collaborations. When formal contracts are not enforceable, such efficiency-enhancing cooperations can be established via informal agreements, but relational contracts have been largely ignored in the literature on the international organization of value chains. In this paper, we develop a dynamic property rights model of global sourcing. A domestic headquarter collaborates with a foreign input supplier and makes two decisions in every period: i) whether to engage in a costly search for a better partner, and ii) whether to make a non-binding offer to overcome hold-up problems. Our key result is that the possibility to switch partners crucially affects the contractual nature of buyer-supplier relationships. In particular, some patient firms do not immediately establish a relational contract, but only when they decide to stop searching and thus launch a long-term collaboration with their supplier. From our model, we develop an instrumental variable estimation strategy that we apply using transaction-level data of fresh Chinese exporters to the US. We obtain empirical evidence in line with the theoretical prediction of a positive causal effect of match durations on relational contracting.

Keywords: Firm organization, input sourcing, relational contracts, supplier search, processing trade, China.

JEL-class.: D23, L23, F23

*University of Nottingham, GEP, CESifo and CEP. Email: `fabrice.defever@nottingham.ac.uk`

†Düsseldorf Institute for Competition Economics (DICE), Heinrich-Heine-Universität Düsseldorf, Germany. Email: `fischer@dice.hhu.de`

‡Düsseldorf Institute for Competition Economics (DICE), Heinrich-Heine-Universität Düsseldorf, Germany, and CEPR. Email: `suedekum@dice.hhu.de`

The authors thank Pol Antràs, Jonathan Eaton, Hartmut Egger, Gabriel Felbermayr, Willi Kohler, Kalina Manova, Michael Pflüger, Andrei Potlogea, John Ries, Frederic Robert-Nicoud, Esteban Rossi-Hansberg, Deborah Swenson, Olga Timoshenko, Jo van Biesebroeck, Frank Verboven and various seminar audiences 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

Intermediate inputs account for a substantial share of global trade. A large chunk of this trade involves a *buyer* – often a large and powerful headquarter firm located in a high-income country – who imports components from a foreign *supplier* for further processing (Johnson and Noguera, 2012). These suppliers are often located in low-wage countries with relatively unfamiliar market conditions and weak legal institutions, and it is well documented that economic exchange in such environments is hampered by two main problems: search and contractual frictions. First, the buyer needs to find a supplier who is technologically capable of producing the desired input in appropriate quantity and quality and at low costs. Second, it wants to make sure that this supplier does not behave opportunistically, frequently renege on the agreed terms, engage in haggling, and so on. Finding such an efficient and reliable partner is costly and can involve a time consuming trial-and-error process before the firm is finally satisfied with the match.

How can hold-up problems and the resulting inefficiencies be alleviated when formal contracts are not enforceable? The theory of the firm, especially the work by Baker, Gibbons and Murphy (2002), suggests that *relational contracts* (RC) play a very important role. These are purely informal and trust-based agreements to cooperate in an enduring relationship. And indeed, various management studies show that headquarter corporations and specialized component manufacturers have a vital interest in establishing and maintaining long-term collaborations in order to create and to share relational rents.[1] The previous economics literature on the international organization of value chains has so far largely ignored relational agreements, however, and how they might resolve underinvestment problems in the context of input sourcing. Our model builds on the seminal property rights theory along the lines of Grossman and Hart (1986) and Hart and Moore (1990), which has been embodied in an international trade context by Antràs (2003) and Antràs and Helpman (2004), and places this into a dynamic setting with repeated interactions. Moreover, our framework allows for supplier re-matching subject to a search friction, since firms in practice can terminate collaborations if they are not satisfied. We endogenously determine *with whom* and *how* the buyer interacts, and show that the possibility to switch partners crucially affects the nature of buyer-supplier relationships.

Theory: In our dynamic model, there is a domestic headquarter firm that collaborates with a foreign supplier in an environment of incomplete contracts. The firm makes two decisions in every period. First, observing the current supplier's efficiency, she decides whether to engage in a costly search for another, potentially better match. Second, she can offer a RC to her supplier, by promising him an ex post bonus payment if he reliably provides the input as stipulated in their agreement. Neither this bonus, nor the relationship-specific investments of the two parties are contractible. Yet, if agents are patient enough, an efficient long-term RC can be implemented in equilibrium that induces first-best effort levels and thus overcomes hold-up and underinvestment problems.

Our model brings out two main results: First, we make clear that long-term collaborations (LTCs) and relational contracts are two different concepts that sharply need to be

[1] Early studies include Dwyer, Schurr and Oh (1987) or Kalwani and Narayandas (1995).

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. By contrast, patient agents want to establish an efficient relational agreement in principle, but they only form this RC once the firm is satisfied with the efficiency of the current supplier and decides to stop searching for a better partner. In other words, and turning to our second main result, our model shows that there is a positive *causal* relationship between match durations (LTCs) and the stability of RCs: some firms will not immediately engage in relational contracting, but once they decide to stop searching and to launch a LTC, this will then *cause* the emergence of a RC.

These insights add a novel perspective to the theory of the firm. Baker, Gibbons and Murphy (2002) have shown that sufficiently patient agents may overcome inefficiencies that bad legal institutions and weak contract enforcement may create. This also applies in the context of input sourcing. Our theory then goes beyond this point, and explicates that the firm's decision *with whom* to interact is crucially important for the interrelated question *how* this interaction will look like.

Empirical Evaluation: Turning to the empirical part of this paper, we estimate the key prediction of our model using Chinese customs data on export transactions in the US (2000-2006). This context is well suited for our purpose, because the paradigm of incomplete contracts appears to be quite plausible when it comes to US input sourcing from China (Antràs 2015; Manova and Zhang 2012). Of course, putting the full micro-structure of our model to a test is challenging, because several aspects are unobservable to the researcher. Our theory is about particular matches of domestic buyers and foreign suppliers, but current data only very rarely allows observing such matches.[2] Moreover, even if they were known, we cannot observe the detailed explicit and implicit match-specific arrangements between the buyer and the supplier on which our model makes sharp predictions.[3] Still, despite those limitations, we exploit the unique features of the Chinese customs data to construct empirical proxies for the two main elements of our theory, the duration and the contractual nature of a buyer-supplier relationship.

In particular, our main aim is to quantify the positive effect of LTCs on RCs that our model predicts. To do so, we build a sample of firms that start exporting a particular product (6-digit industry level) to the US. Following these fresh Chinese exporters over time, we observe which firms still export the same product to the US after a few years, as opposed to those which have terminated that exporting activity in the meantime. This allows us to observe whether a particular export transaction can be thought of as a one-shot deal or a LTC. Regarding the prevalence of RCs, our dataset provides some relevant information whether the Chinese supplier operates under an ordinary or a processing

[2] 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 US. 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 at a co-existence of one-shot and long-term collaborations in that market. Our model is consistent with both empirical features. 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).

[3] Recent work by Macchiavello and Morjaria (2014) is even able to observe, to some extent, the contractual details of certain buyer-supplier arrangements, yet only for one very specific market (coffee beans in Rwanda).

trade arrangement with his US partner. We argue below that particular types of these processing trade arrangements capture the essence of RCs fairly well, for example, if the US headquarter has provided its Chinese supplier with specifically designed equipment. As such arrangements signal a rather strong and interlinked relationship, they are thus indicative that a RC between the buyer and the supplier is in place.

Applying these China-US data, we obtain empirical results firmly in line with the predictions of our theory. We first report a robust positive correlation between the use of RCs and the duration of individual Chinese export transactions. At the industry level, we consistently find that the share of LTCs and RC prevalence are positively correlated across product categories, thus suggesting that more enduring buyer-supplier relationships are more likely to involve elements of relational contracting.

Finally, we use the structure of our theoretical framework to develop an instrumental variable estimation approach that uncovers the *causal* effect of LTCs on the prevalence of RCs. In particular, our model shows that the distribution of suppliers' unit costs is a valid instrument for the length of match durations. The underlying intuition is that more cost dispersion in an industry makes search more attractive, essentially because it raises the benefits of having a good supplier with low unit costs. Yet, more cost dispersion does not directly affect the contractual nature of firm-supplier relationships in our model, but there is only an indirect (negative) effect via the reduced match durations in the time span of observation. Put differently, our estimation strategy exploits the prediction of our theory that the supplier's unit costs are negatively correlated with *duration* but uncorrelated with the *contractual nature* of a collaboration. Applying this estimation approach with our Chinese-US data, we find that cost dispersion is indeed a good predictor for supplier turnover. Our quantitative results in the second stage then suggest that, going from an industry with only one-shot to an industry with only repeated interactions, leads to a sizeable increase of RC arrangements by 38 to 67 percentage points.

Related literature: Our paper is related to different lines of research. A first strand of literature incorporates elements of search and matching into models of international trade, either referring to the search of sales agents and distributors (Rauch 1999; Antràs and Costinot 2011) or to input suppliers as in our model (Rauch and Watson 2003; Carballo, Ottaviano and Volpe-Martincus 2013; Bernard, Moxnes and Ulltveit-Moe 2014). While some of those models also deal with relationship durations, they are mostly silent on contractual choices and do not analyze relational agreements as our paper does.

A second set of models incorporates contractual choices into models of export dynamics, in particular Araujo, Mion and Ornelas (2014) and Aeberhardt, Buono and Fadinger (2014).[4] In these frameworks, a domestic exporter matches with a foreign sales agent who might behave opportunistically. Once the firm has found a reliable partner, she increases exports at the intensive margin. This is then consistent with empirical evidence showing that firms indeed engage in a trial-and-error search for suppliers. Rauch and Watson (2003) and Besedes (2008) argue that firms initially place small test orders when they first deal with a new supplier. If they are dissatisfied, they terminate the collaboration while

[4]Contractual choices between buyers and suppliers are also taken up in a recent paper by Bernard and Dhingra (2015), but in a somewhat different context.

they increase order volumes when they are satisfied. Like the other contributions, our model features such a pattern since it predicts increasing investments once a long-term relational agreement is established. We then differ in two main respects. First, while those papers focus on the exporter, we are mainly interested in the decisions of the importing firm. Second, we can explicitly distinguish long-term collaborations and relational agreements which is not possible in those models.

Third, our model adds to a literature which has studied theoretically and empirically how relational agreements can overcome hold-up problems in the context of firm organization and input sourcing. This includes, for example, the studies by Macchiavello and Morjaria (2014), Board (2011) and Corts and Singh (2004). The models by Kukharskyy (2015), Kamal and Tang (2015) and Kukharskyy and Pflüger (2011) are particularly closely related, as they are also based on the Antràs (2003) framework. However, none of these papers studies the interaction between RCs and supplier search and re-matching, but they focus on the ownership choice of integration versus outsourcing (from which we abstract) in a repeated game setup with a fixed partner.

Finally, our model is also more broadly related to a game theoretic literature which analyzes settings where players not only choose strategies but also with whom they play. Examples include Gosh and Ray (1996), Fujiwara-Greve and Okuno-Fujiwara (2009), Rob and Yang (2010) and Jackson and Watts (2010), as well as applied contributions which investigate partnership problems or matching and contracting in the labour market (e.g., Board and Meyer-ter-Vehn, 2015). Our model focuses on a different context, buyer-supplier relationships in input sourcing, and specifically investigates relational contracts in an infinitely repeated environment with costly partner search.

The rest of this paper is organized as follows. Section 2 present the basic model structure. In Section 3, we analyze the use of relational contracts as an equilibrium strategy when supplier re-matching is ruled out. We introduce this option in Section 4. Section 5 presents our empirical analysis. Section 6 concludes.

2 The Model

Our model is related to Antràs (2003) and Antràs and Helpman (2004) and extends this seminal framework to a dynamic setting with repeated interactions and supplier search and re-matching. We first characterize market demand and firm technology. There is a single firm producing a final good for which it faces the following iso-elastic demand

$$y = A\, p^{-\frac{1}{1-\alpha}}, \quad 0 < \alpha < 1,\, A > 0, \tag{1}$$

where y is quantity, p is price, $A > 0$ measures the demand level, and $\alpha \in (0, 1)$ is a parameter that governs the demand elasticity $1/(1-\alpha)$.

The production technology for the final output requires two inputs, h and m, and has the following Cobb-Douglas specification:

$$y = \left(\frac{h}{\eta}\right)^{\eta} \left(\frac{m}{1-\eta}\right)^{1-\eta}, \quad 0 < \eta < 1 \tag{2}$$

The firm's headquarter (called H) provides the input h herself, and the parameter η captures the headquarter-intensity of final good's production. The component m is an intermediate input which is sourced from an independent supplier (called M).[5] Both inputs are fully relationship-specific and have no value besides for the production of output y. The revenue from selling y can be written as:

$$R = p \cdot y = A^{1-\alpha} \left[\left(\frac{h}{\eta} \right)^{\eta} \left(\frac{m}{1-\eta} \right)^{1-\eta} \right]^{\alpha} \tag{3}$$

We assume that H and M have constant marginal costs of input production, denoted respectively by c_h and c_m, and face zero fixed costs.

In the following we characterize our setup for the repeated game which consists of an initialization phase and a stage game which is repeated ad infinitum.

Initialization phase

In the initialization phase key parameters of the game are determined. It has two steps:

1. H enters the market and learns its unit costs c_h, the headquarter intensity $\eta \in (0,1)$, the demand parameters α and A, and its time-discount factor $\delta \in (0,1)$. All these parameters are public information and remain the same for all periods of the repeated stage game.

2. We assume that there is a continuum of potential suppliers, which are all identical except for the marginal costs c_m that they would incur in the relationship with H. In particular, all potential suppliers have the same outside option ω_M and a common discount factor which is identical to the headquarter's δ.[6] The costs c_m are distributed across potential suppliers according to some distribution $g(c_m)$ with corresponding cumulative density function $G(c_m)$. The headquarter H gets initially matched with some supplier M_0 with unit costs c_m^0 randomly drawn from $g(c_m)$. This initial matching is costless.

Stage game

After the initialization phase the stage game starts. It has the following consecutive steps:

1. **Proposal stage** (cheap talk): H can make M a non-binding and non-contractible proposal specifying investment levels (h,m) and an ex-post bonus payment B to M. We call this proposal, which is essentially just cheap talk, a *relational contract*.

2. **Participation decision stage**: The supplier M decides upon his participation in the relationship with H according to his outside option ω_M.

3. **Investment stage**: The headquarter H and the supplier M simultaneously choose their non-contractible input investments (h,m).

4. **Information stage**: H and M learn the investment level of their production partner.

[5] In this paper we do not analyze the integration versus outsourcing decision central to the Antràs (2003) paper, but assume that M is an independent supplier with full ownership rights over his assets.

[6] This homogeneity assumption avoids various complications which arise in repeated games with heterogeneous time-preference rates as discussed in Lehrer and Pauzner (1999).

5. **Bargaining stage**: If a relational contract was proposed, H can decide to pay the bonus B to M. Otherwise the surplus is split according to an asymmetric Nash bargaining, where $\beta \in (0,1)$ is H's and $(1-\beta)$ is, respectively, M's bargaining power.
6. **Profit realization stage**: The final output is produced and sold. The surplus is divided as specified in stage 5.

At first we ignore supplier search and potential re-matching and assume that H contracts with the initial partner M_0 forever. Below we introduce the option to switch suppliers, which then extends the stage game by one further step.

Notice that our setting is a game of public monitoring. Thus, following the methods of Abreu (1988) we make use of the possibility to identify a simple strategy profile that implements the production of the first-best (joint profit-maximizing) output level as a subgame-perfect Nash equilibrium (SPNE) of the repeated game.[7] As such it will be sufficient to apply the one-step deviation principle to one arbitrary (but representative) stage game in order to confirm a strategy as an equilibrium. We make use of this property in the following implementation.

Static Nash Equilibrium

Before analyzing the repeated game, we briefly consider a static setting where the stage game is only played once. In such a case, we end up with identical hold-up and underinvestment problems as described by Antràs (2003) or Antràs and Helpman (2004).

By backward induction, first consider the bargaining stage 5. Since $\{(h,m),B\}$ is not legally enforceable, it is always optimal for the headquarter not to pay the bonus B. Hence, the two parties will engage in Nash bargaining. Anticipating this, at the investment stage 3, the headquarter and the supplier choose h and m, respectively, in order to maximize their individual payoffs

$$\max_h \ \beta R(h) - hc_h, \qquad \max_m \ (1-\beta)R(m) - mc_m.$$

The resulting investment levels are denoted by $\tilde{h}$, and $\tilde{m}$, and can be found from calculating mutual best responses. They are identical to the results from Antràs (2003):

$$\tilde{h} = \frac{\alpha\beta\eta}{c_h}\,\tilde{R}, \qquad \tilde{m} = \frac{\alpha(1-\beta)(1-\eta)}{c_m}\,\tilde{R}, \tag{4}$$

where

$$\tilde{R} = \mathcal{A}\left[\left(\frac{\beta}{c_h}\right)^{\eta}\left(\frac{1-\beta}{c_m}\right)^{1-\eta}\right]^{\frac{\alpha}{1-\alpha}} \qquad \text{with} \qquad \mathcal{A} \equiv A\,\alpha^{\frac{\alpha}{1-\alpha}}.$$

Thus, in this static world, the equilibrium payoffs are, respectively, given by

$$\pi_H^N = \beta\tilde{R} - \tilde{h}c_h, \qquad \pi_M^N = (1-\beta)\tilde{R} - \tilde{m}c_m, \tag{5}$$

and the participation of the supplier in stage 2 can simply be ensured by a low enough outside option, namely $\omega_M \leq \pi_M^N$, which we henceforth assume to hold.

[7] In the following, we use the terms joint profit-maximizing and first-best investment levels synonymously.

3 Relational contracts without supplier re-matching

We now turn to the repeated game and study how first-best input investments can be implemented by an informal agreement despite the absence of any legally binding contracts. In this section we still consider that the firm collaborates with a fixed partner, namely the initial supplier M_0. The strategy profile of our repeated game has two states, the cooperative and the non-cooperative one. The game starts in the cooperative state, where H promises a bonus payment B to M, and stipulates first-best (joint profit-maximizing) input quantities h^* and m^*. We denote joint first-best profits by

$$\pi^{JFB} = R(h^*, m^*) - h^* c_h - m^* c_m,$$

and the corresponding input quantities and revenue are given by

$$h^* = \frac{\alpha\eta}{c_h} R^* > \tilde{h}, \qquad m^* = \frac{\alpha(1-\eta)}{c_m} R^* > \tilde{m}, \qquad R^* = \mathcal{A}\left(c_h^{\eta} c_m^{1-\eta}\right)^{\frac{-\alpha}{1-\alpha}} > \tilde{R}. \tag{6}$$

In the cooperative state, H and M actually produce those input levels h^* and m^* in stage 3. In stage 5, H pays M the agreed upon bonus B (that we derive below) and keeps the residual revenue for herself. The per-period payoffs for the two parties under this *relational contract* (RC) are thus given by

$$\pi_H^{RC} = R^* - h^* c_h - B, \qquad \pi_M^{RC} = B - m^* c_m. \tag{7}$$

When a deviation from this informal agreement $\{(h^*, m^*), B\}$ occurs, the game switches to the non-cooperative state from the next period onwards and stays in that state forever. In the non-cooperative state, H and M make input investments as in the static game, see (4), and earn payoffs π_i^N in every period as given in (5).[8] In the deviation period, the respective deviator chooses his or her payoff-maximizing input investment given first-best investments of the other party, while anticipating that the surplus is divided via Nash bargaining since inputs can be observed before revenue is realized. We denote the deviation investments as h^D and m^D, respectively, which can be computed as follows

$$\max_{h^D} \ \beta R(h^D, m^*) - h^D c_h, \qquad \max_{m^D} \ (1-\beta) R(h^*, m^D) - m^D c_m.$$

This leads to

$$h^D = \frac{\alpha\eta\beta}{c_h} R(h^D, m^*), \qquad m^D = \frac{\alpha(1-\eta)(1-\beta)}{c_m} R(h^*, m^D), \tag{8}$$

where $R(h^D, m^*) = \mathcal{A}\beta^{\frac{\alpha\eta}{1-\alpha\eta}} \left(c_h^{\eta} c_m^{1-\eta}\right)^{\frac{-\alpha}{1-\alpha}}$ and $R(h^*, m^D) = \mathcal{A}(1-\beta)^{\frac{\alpha(1-\eta)}{1-\alpha+\alpha\eta}} \left(c_h^{\eta} c_m^{1-\eta}\right)^{\frac{-\alpha}{1-\alpha}}$.

[8] Nash reversion is a commonly assumed punishment strategy, and in our context it has the advantage of bringing the model back to the baseline framework by Antràs (2003). Below we also consider different penal codes, such as the "carrot-and-stick" punishment by Abreu (1988). Moreover, in the Appendix we also consider an extension of our model where the firm can only imperfectly monitor the quality of the provided input.

The deviation payoffs are thus given by

$$\pi_H^D = \beta R(h^D, m^*) - h^D c_h, \qquad \pi_M^D = (1-\beta) R(h^*, m^D) - m^D c_m. \tag{9}$$

In the following we show under which conditions the relational contract can be implemented as a SPNE of the repeated game, given the Nash reversion trigger strategy. First, after the proposal in stage 1, the supplier decides in stage 2 on participation according to:

$$B - m^* c_m \geq \omega_M \tag{PC-M}$$

This participation constraint (PC) requires that the bonus must at least compensate the full production costs at the first-best investment level m^*, plus his outside option ω_M.

Second, in stage 3, adhering to the agreement and choosing first-best investment levels must be better for both parties than deviating once and reverting to the non-cooperative state in the future. We face the following incentive compatibility (IC) constraints,

$$\frac{1}{1-\delta}\pi_M^{RC} \geq \pi_M^D + \frac{\delta}{1-\delta}\pi_M^N \tag{IC-M}$$

$$\frac{1}{1-\delta}\pi_H^{RC} \geq \pi_H^D + \frac{\delta}{1-\delta}\pi_H^N, \tag{IC-H}$$

where π_i^{RC}, π_i^D and π_i^N for $i = H, M$ are given by (5), (7), and (9). Rearranging (IC-H) and (IC-M) for B, we can identify the range of bonus payments for which relational contracting is incentive compatible:

$$B \leq R^* - h^* c_h - \left[\delta \pi_H^N + (1-\delta)\pi_H^D\right] \equiv B_H(\delta) \tag{IC-H'}$$

$$B \geq m^* c_m + \left[\delta \pi_M^N + (1-\delta)\pi_M^D\right] \equiv B_M(\delta) \tag{IC-M'}$$

where it follows from (IC-M') and (PC-M) that the supplier's participation constraint is always satisfied, since $\pi_M^D > \pi_M^N > \omega_M$ always holds.

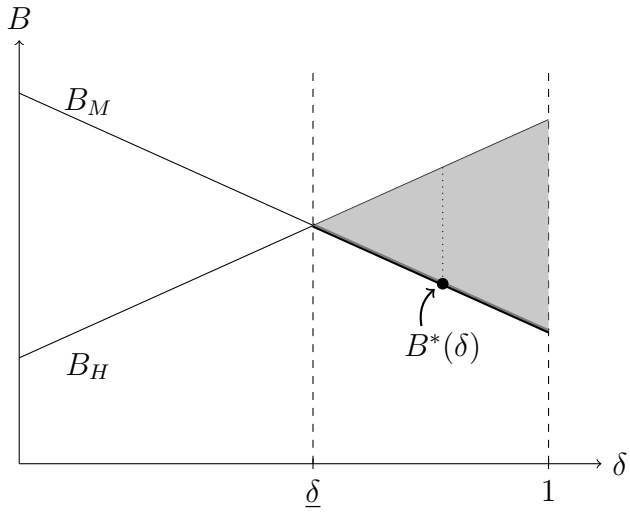

Figure 1: Feasible and optimal bonus payments

In Figure 1 we depict the δ-specific bonus B_i for which the respective player is indifferent between sticking to, and deviating from the RC. As can be seen, the firm is willing to transfer a higher maximum bonus $B_H(\delta)$ at higher δ, since the RC is relatively more attractive when agents are more patient. By the same argument, the minimum required bonus $B_M(\delta)$ to keep the supplier within the RC is lower at higher levels of δ. Clearly, the RC can only be incentive compatible if $B_H(\delta) \geq B_M(\delta)$, i.e., if the maximum bonus that H is willing to pay exceeds the minimum bonus required by M. Since $B_H(\delta)$ is linearly increasing and $B_M(\delta)$ linearly decreasing in δ, the following is true: If we can find a $\underline{\delta}$ for which $B_H(\underline{\delta}) = B_M(\underline{\delta})$ holds, then for all $\delta \geq \underline{\delta}$ the RC can be made incentive compatible with an appropriate bonus. As shown in Proposition 1 below, such a critical discount factor $\underline{\delta} \in (0, 1)$ always exists. The grey area in Figure 1 depicts the set of feasible bonus payments for which the RC is an equilibrium of the repeated game.

Notice that the firm has no incentive to transfer more than necessary to make the RC incentive compatible for M. Hence, she will offer a B such that (IC-M') binds with equality, and we denote the resulting optimal bonus payment by $B^*(\delta)$. Summing up,

Proposition 1. *Under the relational contract (RC) first best investment levels* (h^*, m^*) *can be supported as a SPNE of the repeated game for all*

$$\delta \geq \underline{\delta}(\alpha, \beta, \eta), \tag{10}$$

where $\underline{\delta}(\alpha,\beta,\eta) = \frac{1-\alpha-(1-\alpha\eta)\beta^{\frac{1}{1-\alpha\eta}}-(1-\alpha(1-\alpha\eta))(1-\beta)^{\frac{1}{1-\alpha(1-\eta)}}}{(1-\alpha\eta)\left[\beta^{\frac{1-\alpha(1-\eta)}{1-\alpha}}(1-\beta)^{\frac{\alpha(1-\eta)}{1-\alpha}}-\beta^{\frac{1}{1-\alpha\eta}}\right]+(1-\alpha(1-\eta))\left[\beta^{\frac{\alpha\eta}{1-\alpha}}(1-\beta)^{\frac{1-\alpha\eta}{1-\alpha}}-(1-\beta)^{\frac{1}{1-\alpha(1-\eta)}}\right]}$

and bonus payment $B^*(\delta) = m^* c_m + \left[\delta\pi_M^N + (1-\delta)\pi_M^D\right]$.

Proof. The result can be derived from (IC-H') and (IC-M') by solving $B_H(\underline{\delta}) = B_M(\underline{\delta})$ for $\underline{\delta}$. This yields $\underline{\delta} = \frac{\pi^{JFB}-\pi_H^D-\pi_M^D}{\pi_H^N+\pi_M^N-\pi_H^D-\pi_M^D} = \frac{\sum_i \pi_i^D - \pi^{JFB}}{\sum_i \pi_i^D - \sum_i \pi_i^N}$. Using $B^*(\delta)$ in (7), and comparing it to (5) and (9), it follows immediately that $\pi_i^D > \pi_i^{RC} > \pi_i^N > 0$ for $i = H, M$. This ordering of payoffs implies $\sum_i \pi_i^D > \pi^{JFB} > \sum_i \pi_i^N$, and we thus have $\underline{\delta} \in (0, 1)$. □

Several points are worth noting about this result. First, the agents in our model encounter a situation that may be thought of as a repeated prisoner's dilemma. This follows from the fact that the ordering of payoffs $\pi_i^D > \pi_i^{RC} > \pi_i^N$ holds both for the headquarter ($i = H$) and for the supplier ($i = M$). Similar as in collusion games from the industrial organization literature, the cooperative outcome (the RC) is thus only an equilibrium if agents are patient enough. Second, if the RC arises, it makes the supplier strictly better off compared to static Nash play, even though the firm only pays him the minimum required bonus. More specifically, we have $\pi_M^{RC} = B^* - m^* c_m = \left[\delta\pi_M^N + (1-\delta)\pi_M^D\right] > \pi_M^N$. The intuition is the strategic value of the deviation option which is capitalized in the optimal bonus payment $B^*(\delta)$. Third, we can also introduce an explicit notion of relational rent sharing, by assuming that H offers M a bonus payment $B_M(\delta) < B^*(\delta) < B_H(\delta)$. The rest of Proposition 1, in particular the critical discount factor $\underline{\delta}$, still apply in that case. Yet, even without this rent sharing, our model provides a rationale why informal RCs are to the mutual benefit of both parties in this context.

Fourth, and most interestingly for the remainder of this paper, notice from (10) that the demand parameter A, as well as the unit costs c_h and c_m do not affect the critical discount factor $\underline{\delta}$, which only depends on the demand elasticity parameter α, the headquarter intensity η, and the bargaining power β. In the current model setup without supplier re-matching, this means that the unit cost level c_m^0 of the perpetual supplier M^0 is irrelevant for the question whether a RC with him is feasible or not. If M^0 happens to be more efficient and has lower c_m^0, this would proportionally raise all payoffs π_i^D, π_i^{RC}, and π_i^N for $i = H, M$, as well as the bonus payment $B^*(\delta)$. However, it would not affect $\underline{\delta}$ as the linear curves in Figure 1 still intersect at the same level of δ. Intuitively, changes in the supplier's costs do not affect payoffs relatively stronger for one contractual arrangement than for another. Note that this property does not crucially hinge on the assumed functional forms in (1) and (2) with iso-elastic demand and Cobb-Douglas production.[9]

Finally, it can be shown that the essence of Proposition 1 also remains unchanged with different penal codes than the Nash reversion trigger strategy that we take as our benchmark. In particular, in Appendix A we also consider an optimal penal code along the lines of Abreu (1988) which involves "carrot-and-stick" punishment where, after a deviation, the other agent punishes the respective deviator with zero investments for T periods, and then returns to the cooperative state. For this case we obtain a different critical discount factor $\hat{\delta}$ that is decreasing in the punishment length T, but that is still independent of c_m. Moreover, as a further robustness check, we there also consider an extension of our model where the firm cannot infer the quality of the input m and is subject to fluctuations in the demand level A, similar as in the well-known collusion model by Green and Porter (1984). Also in that case we can implicitly determine a critical discount factor that now also depends on demand uncertainty, but still not on the unit costs c_m.

Our key result – the cost-orthogonality of the critical discount factor – therefore seems to hold under fairly general condition and does not hinge on the peculiarities of the model parametrization or the punishment strategies. We will make use of this cost-orthogonality of the critical discount factor below when developing our estimation strategy.

4 Supplier re-matching

We now assume that the headquarter has the opportunity to switch suppliers after every round of output realization. In particular, the following stage is added to the game, with all other stages remaining the same as before.

7. **Re-matching stage**: H can pay a publicly known fixed cost $F > 0$ and re-match to a new supplier. Let c_m^t be the unit cost of her current supplier, and c_m^{t+1} the unit cost of the new supplier that she has encountered. This cost c_m^{t+1} is randomly drawn from the distribution function $g(c_m)$ and is perfectly observable to the headquarter.

For the ease of exposition, we make two simplifying assumptions in the main text. First, we consider a particular parametrization of $g(c_m)$ with only two possible realizations

[9] Specifically, this independence of $\underline{\delta}$ holds for any production and demand function such that the resulting payoffs π_i^D, π_i^{RC}, and π_i^N are homogeneous of some common degree b in the parameters c_m, c_h, and A.

$\{c_m^l, c_m^h\}$, where $c_m^h > c_m^l > 0$. Hence, there are only low-cost and high-cost suppliers. The share of low-cost suppliers in the universe of potential suppliers, and thus the probability to find such a partner in every round of re-matching is denoted by P.[10] Second, we assume that the headquarter *must* re-match when she decides to engage in costly search, even when the supplier she encounters is less efficient than her current partner.

In the Appendix, we relax both simplifications. In particular, there we consider i) a general distribution $g(c_m)$ from which c_m is drawn, and ii) we separate the processes of supplier *search* and *re-matching*. Specifically, in that version of the model the firm pays a search cost F, encounters a candidate supplier, and re-matches only if that candidate has strictly lower unit costs than her current supplier (otherwise she sticks to the old partner). It turns out that these more complicated versions yield insights which are similar to our simpler benchmark model.

We now analyze the re-matching decision and its interrelation with the contractual choice on RCs. Recall from the previous Section, where a change of suppliers was ruled out, that the critical discount factor $\underline{\delta}$ from (10) is independent of c_m. Therefore, since δ is observable and identical for H and all (potential) M, the headquarter knows already at the beginning of the game if she will at any point be able to achieve a RC or not. We can therefore distinguish two cases in the subsequent analysis.

4.1 The case with impatient agents ($\delta < \underline{\delta}$)

With $\delta < \underline{\delta}$, RCs can never be implemented regardless of the supplier type. Due to their impatience, agents would always deviate from the prescribed first-best behaviour, and thus Nash bargaining and the associated underinvestment emerge for sure.

Turning to the re-matching decision, at the end of every period the headquarter trades off the fixed cost F against the expected payoff difference when changing partners. Recalling that there are only two supplier types, this re-matching decision has two dimensions:

1. if the current supplier is a high-cost type ($c_m^t = c_m^h$), does the firm re-match?
2. if the current supplier is a low-cost type ($c_m^t = c_m^l$), does she stick with that partner?

It is convenient to answer the second question first. Here it is important to notice that for the case of impatient agents the answer is unambiguously positive, that is, further re-matching makes no sense if the firm has found (or is initially matched with) a low-cost type. The reasons are twofold: i) a payoff improvement is only possible if the current supplier is a high-cost type, so that the firm has the chance P to find a low-cost supplier, but not if the current supplier is already a low-cost type, and ii) since both partners only make Nash investments due to their impatience, there exist no profitable deviations from this equilibrium strategy of mutual best responses for neither partner. It is thus never optimal for the firm to re-match further, even if F is very low and/or P is very high, but Nash play with the low-cost supplier will sustain over the entire game. Intuitively,

[10] We thus assume a perfectly elastic supply of good component manufacturers. In the Chinese context considered below in the empirical part, this assumption appears quite plausible because the number of Chinese component manufacturers has been vastly increasing over time.

re-matching is not a credible threat, and it cannot be exploited strategically in order to squeeze rents from the supplier. Both agents are aware that the same hold-up problem arises in every stage game round, so it is always better to avoid paying the fixed cost F.

Turning to the first question, suppose the initial supplier is a high-cost type ($c_m^0 = c_m^h$). We then need to analyze if the firms prefers to change suppliers at any point in time. Let $\pi_H^{N,l}$ and $\pi_H^{N,h}$ denote the headquarter's Nash payoff per period when matched with a low-cost, respectively, a high-cost supplier, with $\pi_H^{N,l} > \pi_H^{N,h}$ as directly evident from (5). The following condition must be satisfied in order for re-matching to occur:

$$E[\pi_H^{RM,N} \mid c_m^0 = c_m^h] > \frac{1}{1-\delta}\pi_H^{N,h}, \tag{11}$$

i.e., the expected payoff when changing partners must be higher than the continued Nash payoff with the initial high-cost supplier. Suppose it is profitable for the headquarter to re-match in period 0. She thus pays F and with probability $(1-P)$ gets matched to another high-cost supplier. If that happens, she earns $\pi_H^{N,h}$ from this relationship in period 1, and will then re-match again at the end of that round because re-matching must be profitable after round 1 when it was profitable after round 0, and so forth. With probability P, however, she finds a low-cost supplier at the end of the round. If that happens, she stops changing partners and stays with this low-cost supplier in a long-term collaboration of repeated Nash-bargainings as shown before.

This decision problem can be written formally in the following way,

$$\begin{aligned} V_0 &= \pi_H^{N,h} - F + \delta V_1 \\ V_i &= (1-P)\left(\pi_H^{N,h} - F + \delta V_{i+1}\right) + P\frac{1}{1-\delta}\pi_H^{N,l}, \quad i = 1, 2, ... \end{aligned}$$

Observing that the decision problem is the same in every round ($V_i = V_{i+1}$ for $i = 1, 2, ...$), the program can be simplified to

$$V_0 = \pi_H^{N,h} - F + \delta V_1, \qquad V_1 = \frac{1}{1-\delta(1-P)}\left((1-P)(\pi_H^{N,h} - F) + \frac{P}{1-\delta}\pi_H^{N,l}\right),$$

and solving for V_0 we get the following expected profit when starting to change suppliers:

$$E[\pi_H^{RM,N} \mid c_m^0 = c_m^h] = \frac{1}{1-\delta(1-P)}\left[\pi_H^{N,h} - F + \frac{\delta P}{1-\delta}\pi_H^{N,l}\right]. \tag{12}$$

Combining expressions (11) and (12), we then obtain the following re-matching condition,

$$F < \frac{\delta P}{1-\delta}\left[\pi_H^{N,l} - \pi_H^{N,h}\right] \equiv \overline{F}^1, \tag{13}$$

which shows that the cost F must be low enough in order for re-matching to occur.[11] Notice that the critical search cost level $\overline{F}^1$ depends positively on δ, i.e., re-matching is

[11] Making use of a structural assumption on F, we can also express this critical search cost level in terms of model fundamentals. In particular, assume that $F \equiv fR^*(c_m = c_m^h)$, $f > 0$, which is without loss of generality because $R^*(c_m = c_m^h) > 0$ and constant for given $g(c_m)$. We can then rewrite (13) as $f < \frac{\delta P}{1-\delta}(1-\alpha\eta)\beta^{\frac{1-\alpha(1-\eta)}{1-\alpha}}(1-\beta)^{\frac{\alpha(1-\eta)}{1-\alpha}}\left[\left(c_m^h/c_m^l\right)^{\frac{\alpha(1-\eta)}{1-\alpha}} - 1\right]$, which is an equivalent expression for $\overline{F}^1$.

more attractive the more patient the agents are. The intuition is that future profits then matter more, and hence it becomes more important to find an efficient low-cost supplier. Moreover, $\overline{F}^1$ depends positively on P and on the term $[\pi_H^{N,l} - \pi_H^{N,h}]$. This shows that re-matching is more likely the higher is the chance to find a low-cost supplier, and the larger is the headquarter's per-period payoff difference with a good and a bad partner.

Summing up, in the impatient agents case with $\delta < \underline{\delta}$, we can conclude that no supplier turnover will ever arise if search costs are too high ($F \geq \overline{F}^1$) or if the initial supplier is already a low-cost type. On the other hand, if the initial supplier is a high-cost type and if $F < \overline{F}^1$, the firm will change suppliers in every period until a low-cost supplier is found (which happens with strictly positive probability in finite time) and then stops re-matching forever. Notice that there is non-cooperative hold-up and underinvestment behaviour all the time, both with and without search and before and after the firm has found her ultimate match, because agents are too impatient to forego the deviation temptation.

4.2 Re-matching behaviour with patient agents ($\delta \geq \underline{\delta}$)

Among patient agents RCs are now feasible in principle, and would emerge if supplier re-matching was ruled out. With respect to the re-matching decision, as before, it has two dimensions: i) start re-matching if $c_m^t = c_m^h$, and ii) stop re-matching if $c_m^t = c_m^l$.

Again we start with the second aspect. This analysis now becomes more involved, because a "cheat-and-run" behaviour might emerge as a profitable off-equilibrium strategy. In particular, suppose $c_m^0 = c_m^l$, and the firm offers a RC to her initial low-cost supplier. However, rather than actually continuing the relationship with M_0, the firm may now deviate given the first-best supplier investment m^* and earn $\pi_H^{D,l}$ in that period, re-match to a new supplier in order to avoid the punishment (Nash reversion) of the old partner, and then deviate-and-rematch in every subsequent round. Clearly, if F is low and P is high, this can be more attractive than the RC with the initial low-cost supplier. The reason is, essentially, that a profitable off-equilibrium deviation strategy with payoff π_H^D now exists, in contrast to the impatient agents case where everybody plays Nash.[12] To rule out "cheat-and-run", we have to assume that the search costs are above a threshold, otherwise a RC would never form. In Appendix B we derive this lower bound $\widetilde{F}$ and show that it is increasing in P. The condition $F > \widetilde{F}$ then ensures that "cheat-and-run" will not emerge. As a consequence, given that $c_m^t = c_m^l$, a long-term RC relationship is established once a low-cost supplier is found and is sustained forever without further re-matching.

Turning to the first question, re-matching among patient agents may thus only occur if the initial supplier is a high-cost type, $c_m^0 = c_m^h$, and as before it will actually occur if search costs F are low enough. In particular, the expected payoff when starting to switch partners must be higher than the continuation payoff with the initial high-cost supplier,

$$E[\pi_H^{RM,RC} \mid c_m^0 = c_m^h] > \frac{1}{1-\delta}\pi_H^{RC,h}. \tag{14}$$

[12] It is important to realize that "cheat-and-run" is an off-equilibrium strategy, because in this full information game M will realize the firm's deviation and deviate himself, thus leading to Nash investments in equilibrium.

If that condition is violated, the headquarter forms a RC with her initial high-cost supplier which is sustainable since $\delta > \underline{\delta}$. If condition (14) is satisfied, however, the firm starts re-matching. Once she encounters a low-cost supplier, re-matching stops and she forms a long-term RC collaboration with that partner as shown before. But in every round where she switches partners, the firm encounters a high-cost supplier with probability $(1-P)$. It may thus take several periods before supplier turnover stops once and for all. Note that during this transition period with one-shot interactions both parties only make Nash investments, since they anticipate the supplier replacement in each round so that RCs are not credible.[13] In other words, the RC only forms once the firm is satisfied with her match, stops searching and decides to launch a long-term collaboration (LTC).

Analogous to the case of impatient agents, the decision problem for patient agents can be formalized by the following program:

$$V_0 = \pi_H^{N,h} - F + \delta V_1, \qquad V_1 = \frac{1}{1-\delta(1-P)}\left((1-P)(\pi_H^{N,h} - F) + \frac{P}{1-\delta}\pi_H^{RC,l}\right),$$

which by solving for V_0 implies these expected profit when engaging in re-matching:

$$E[\pi_H^{RM,RC} \mid c_m^0 = c_m^h] = \frac{1}{1-\delta(1-P)}\left[\pi_H^{N,h} - F + \frac{\delta P}{1-\delta}\pi_H^{RC,l}\right] \tag{15}$$

Plugging (15) into (14), we then obtain the critical search cost level $\overline{F}^2$ for this case,

$$F < -\left(\pi_H^{RC,h} - \pi_H^{N,h}\right) + \frac{\delta P}{1-\delta}\left[\pi_H^{RC,l} - \pi_H^{RC,h}\right] \equiv \overline{F}^2, \tag{16}$$

which also depends positively on δ, P, and on the payoff difference between a good and a bad partner that is now given by $[\pi_H^{RC,l} - \pi_H^{RC,h}]$. Moreover, $\overline{F}^2$ depends negatively on the term $[\pi_H^{RC,h} - \pi_H^{N,h}]$, that is, re-matching becomes less likely if launching a RC with the current high-cost supplier generates a larger payoff gain for the headquarter. Finally, for consistency, the upper bound $\overline{F}^2$ must be larger than the lower bound $\widetilde{F}$ derived before. As also shown in Appendix B, this ranking $\widetilde{F} < \overline{F}^2$ can be guaranteed by making appropriate restrictions on $g(c_m)$, namely that c_m^h and c_m^l are not too similar.

4.3 Summary of theoretical results and estimation strategy

The following Proposition summarizes the key theoretical findings derived so far:

Proposition 2. *a) Suppose the headquarter is initially matched with a low-cost supplier ($c_m^0 = c_m^l$). Patient agents (with $\delta > \underline{\delta}$) will collaborate with that supplier forever in a relational contract (RC) agreement with $\{h^*, m^*\}$ and $B^*(\delta)$, assuming that re-matching costs are not too low ($F > \widetilde{F}$). Impatient agents (with $\delta < \underline{\delta}$) will form a long-term collaboration of repeated Nash bargainings with that supplier.*

[13] This is different when the firm is not forced to re-match when paying F but has the option to keep the old supplier. In that model version, considered in the Appendix, it can be optimal for very patient agents to offer a RC to a high-cost supplier in the transitory period, despite the ongoing search for a better partner.

b) With $c_m^0 = c_m^h$, impatient agents with $\delta < \underline{\delta}$ start re-matching if $F < \overline{F}^1$ and continue until they find a low-cost supplier. The RC can never be implemented.

c) With $c_m^0 = c_m^h$, patient agents with $\delta > \underline{\delta}$ start re-matching if $F < \overline{F}^2$ and continue until they find a low-cost supplier. The RC forms with the ultimate low-cost supplier, but not with any high-cost supplier.

For the single headquarter firm considered in our model, supplier turnover will thus either occur not at all, or otherwise take place in every stage game round until a low-cost supplier is found. Figure 2 illustrates these results. The left panel depicts the case where the current supplier is still a high-cost type, while the right panel focuses on the constellation with a low-cost supplier. The patience level δ is on the vertical, and the search costs level F is on the horizontal axis in both panels. The critical discount factor $\underline{\delta}$ is shown as the horizontal line. Re-matching and thus one-shot interactions occur in the darkly shaded area if F falls short of the respective critical level.

In the right panel, we never observe any re-matching, since the firm already has a low-cost supplier. Below the solid line, agents are impatient and engage in a LTC with repeated Nash bargainings. Above the solid line, agents are patient and the RC forms.

Turning to the left panel, for the impatient agents (below the horizontal line) we generally observe Nash play. This is true, i) when search costs are high and the firm sticks with her initial high-cost supplier forever, and ii) also when search costs are low and the firm is in the process of finding a better partner. This can take multiple rounds, and after the impatient firm has succeeded in finding a low-cost supplier, the constellation switches to the right panel of Figure 2.

Now consider the patient agents case (above the horizontal line) in the left panel. When search costs are high, the firm goes for the RC with her initial high-cost supplier. When search costs are low, however, the firm seeks for a better partner. During that transition, when it has not yet found a low-cost type but still deals with high-cost types, we observe Nash play despite the fact that agents are in principle sufficiently patient. The RC only forms once a low-cost supplier has been found, after which the constellation switches to the right panel. Put differently, our theory implies that a LTC is not a sufficient condition for a RC. In fact, for impatient agents, collaborations can be long-term without being a RC.[14] Moreover, and even more importantly, there is a positive *causal* relationship between the two. In particular, patient agents only form a RC once they decide not to have any further turnover, that is, the launching of a LTC with their current supplier *causes* the start of a RC. The aim of our empirical analysis is to quantify this causal link.

[14] In our baseline model, a LTC is a necessary condition for a RC. This is different in the appendix version where search and re-matching are separated. In that case, RCs can even form on a short-term basis during the search for a better partner (if agents are very patient), and a LTC in neither sufficient nor necessary for a RC.

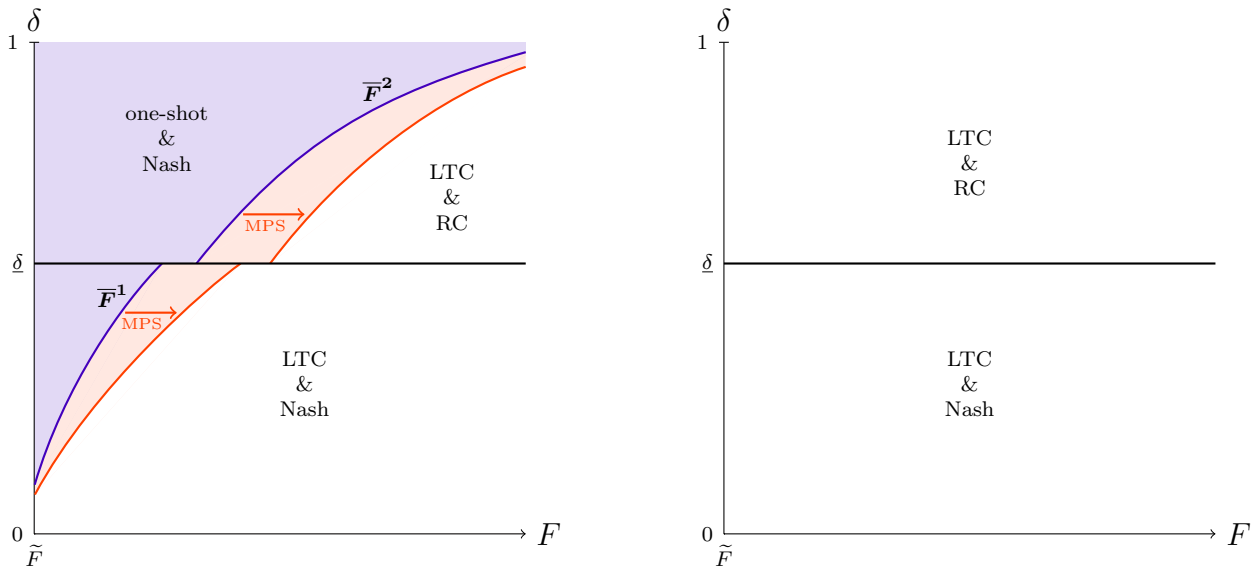

(a) $c_m^0 = c_m^h$ (b) $c_m^0 = c_m^l$

Figure 2: Contractual and re-matching decision

This exercise is difficult for two reasons. First, as argued in the introduction, current data typically does not allow to observe the detailed match-specific contractual arrangements in individual buyer-supplier relationships. To address this problem, we rely on proxies for match durations and the prevalence of RCs as discussed in the next section.

The second problem is conceptual. Suppose we run the following simple ordinary least squares (OLS) estimation at the level of individual buyer-supplier matches i:

$$\text{Relational contract}_i = b_0 + b_1 \cdot \text{match duration}_i + \text{controls}_i + \nu_i.$$

By the previous arguments, we expect $b_1 > 0$. However, the coefficient we obtain from this empirical exercise does *not* identify the causal effect of LTCs on RCs. The reason for the inherent endogeneity problem is apparent from the left panel of Figure 2: More patient agents are not only more likely to form RCs, they are also likely to search and re-match more, since they care more about future profits and, thus, have a stronger interest to collaborate with a low-cost supplier for any given search cost level. This can be seen by noting that the darkly shaded re-matching area becomes larger at higher levels of δ. Hence, there is also a *negative* effect working in the opposite direction, from RCs to LTCs. If we could control for the agents' patience level, we could sort out this issue of reverse causality and identify the true coefficient b_1. However, the parameter δ is typically unobservable to the econometrician, and we thus face an omitted variable problem that jeopardizes the causal interpretation of the coefficient b_1 in the OLS estimation above.

To address this identification issue, we exploit the structure of our theoretical model in order to develop an instrumental variable (IV) estimation approach that allows us to isolate the causal effect we are aiming for. For this empirical approach, we need an instrument for match durations that is uncorrelated with the critical discount factor, and thus, with the prevalence of relational contracts. Recalling Proposition 1, such an

instrument can be constructed from the suppliers' unit costs c_m which are orthogonal to $\underline{\delta}$. More specifically, to establish our identification strategy, we consider a mean-preserving spread (MPS) in the distribution of supplier costs $g(c_m)$ across all potential suppliers in this industry, which we model as an increase in the difference between c_m^l and c_m^h with a constant probability P. This leads to the following Proposition:

Proposition 3. *A mean-preserving spread in the distribution $g(c_m)$ (a larger difference between c_m^l and c_m^h at constant P) increases the critical search cost levels $\overline{F}^1$ and $\overline{F}^2$, and thereby expands the parameter range where the headquarter engages in re-matching.*

Proof. The result follows directly from equations (13) and (16), because this MPS increases the terms $(\pi_H^{RC,l} - \pi_H^{RC,h})$ and $(\pi_H^{N,l} - \pi_H^{N,h})$ as π_H^{RC} and π_H^N are both monotonically decreasing in c_m. In the left panel of Figure 2, the MPS is illustrated by the shift of $\overline{F}^1$ and $\overline{F}^2$ to the right, which enlarges the area where re-matching occurs.

From Proposition 3, we can therefore conclude that a more dispersed cost distribution in the industry increases the headquarter's re-matching propensity, and therefore negatively impacts on the observed match durations in a certain time span for given values of δ, search costs F, and the other parameters. The intuition is that the relative efficiency of low-cost suppliers versus high-cost suppliers has increased, which makes finding a low-cost partner more attractive. Yet, this higher cost dispersion does *not* affect the critical discount factor $\underline{\delta}$ and, thus, it does not directly affect the prevalence of RCs. It only affects this contractual choice indirectly, because it leads to more supplier turnover and thereby can induce some (patient) agents not to offer RCs immediately, but only when they have found their right match. Put differently, more cost dispersion will not *directly* affect contractual choices, but only indirectly via its negative effect on match durations.

Notice that the cost dispersion instrument is *industry*-specific. We therefore conduct our main estimations below at the level of narrowly defined product categories (six-digit industries), indexed by j.[15] It can be represented by the following simultaneous system of equations that we estimate with two-stage least squares (2SLS) techniques:

$$\textbf{First stage:} \quad \text{Match duration}_j = a_0 + a_1 \text{Cost Dispersion}_j + \text{controls}_j + \epsilon_j$$

$$\textbf{Second stage:} \quad \text{Relational contracts}_j = b_0 + b_1 \widehat{\text{Match duration}}_j + \text{controls}_j + e_j$$

The model predicts that the first-stage coefficient a_1 is negative, i.e., more cost dispersion leads to more search and shorter match durations as shown in Proposition 3. Moreover, cost dispersion is the theoretically ideal instrumental variable given the structure of our model, because the exclusion restriction is satisfied.[16] In the second stage, we include

[15] The identifying assumption is that all parameters except the search costs F are industry-specific and common to all individual matches i within industry j and abstract from all cross-industry effects. This pins down the critical search cost levels $\overline{F}_j^1$ and $\overline{F}_j^2$, and implies that all firms and suppliers in j encounter the same constellation as shown in Figure 2 above. We then allow for headquarter-specific variation $F_{i(j)}$ in order to capture the within-industry variation in match durations and contractual designs that we observe in the data.

[16] The exclusion restriction would be violated if firms were *directly* more reluctant to engage in relational contracting with a supplier of low quality, even if search and re-matching are ruled out. As argued in Section 3, however, the cost-orthogonality of the critical discount factor holds under fairly general conditions.

standard proxies for the parameters α, β and η as controls since they directly affect the critical discount factor and thereby, potentially, contractual choices.

With this IV estimation approach, we identify our main coefficient b_1 only from the variation in match durations explained by cost dispersion, which is graphically captured by the *shifts* in the critical search cost levels shown above in Figure 2. It is, by construction, unaffected by the unobservable patience level δ, and it therefore yields the causal effect of LTCs on RCs that our model predicts.

5 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: the prevalence of relational contracts (RC), and whether a buyer-supplier match is a long-term collaboration (LTC). We first present some preliminary correlations, both at the micro-level of individual transactions and at the level of narrowly defined industries. Afterwards, we move to the instrumental variable estimation using data for industry-wide cost dispersion.

5.1 Data and variables

Export transactions in the US of fresh Chinese suppliers: The empirical evaluation of the model exploits transaction-level custom data from the Chinese General Administration of Customs for the period 2000 to 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 she decides to re-match. To reflect this, we only consider transactions of *fresh* Chinese exporters which 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 US, which is the main destination for Chinese exports besides Hong-Kong and Macao. This restriction allows us to abstract from the possible variation in the destination markets' contracting environments. Specifically, our sample is composed of Chinese firms i which export a product j to the US in year t, where $t = \{2001, 2002, 2003\}$, but which have not exported *anything* to *any* destination outside China during the previous years since 2000.[17] We furthermore exclude all product categories j with fewer than 10 active firms and those where control variables are not available. The final sample then includes 16,150 fresh Chinese exporters, starting a total of 63,580 new export transactions in the US, and spans 1,004 HS6-digit product categories.

[17] As a robustness check, we also consider an alternative approach and only require that firms did not export anything to anywhere in year $(t-1)$. Results turn out to be very similar to those reported below.

Match durations: Following them over time, we observe which firms that started an exporting activity in the US market in year t still export the *same* product j to the US after three years (in $t+3$), as opposed to those which have terminated that exporting activity in the meantime. Of the 63,580 transactions in our sample, it turns out that 27,572 (43.3%) are terminated after less than one year, while 20,402 (32.1%) of them endure for more than 3 years.[18] Aggregating to the HS6 product level, we then compute the share of transactions in industry j that got started in $t = \{2001, 2002, 2003\}$ and that are still active in year $(t+3)$. This gives us our measure on the average match duration in industry j, which we use as our main explanatory variable: the share of LTCs.

Relational contracts: Measuring the dependent variable is more challenging, because the various explicit and implicit match-specific contractual arrangements between the American buyer and the Chinese supplier are, of course, not observable to the researcher. Yet, the Chinese custom data provides some information about the type of arrangement between the partners that we can exploit to construct an empirical proxy. In particular, as a measure for the prevalence of RCs, we use the transactions realized using processing trade arrangements – a data feature that Feenstra and Hanson (2005) have previously utilized in a somewhat different, but related context.[19]

Why is processing trade a reasonable proxy for relational contracts? The rationale is the following: First, processing arrangements require to deal with a specific foreign client. An approval from provincial-level commerce departments must be sought by the Chinese firm before it can engage in processing trade and must include details on the foreign partner. The documents for application notably require the draft agreement about the deal signed by the Chinese supplier with its foreign partner, thus requiring a certain amount of familiarity between the two parties. Second, and more importantly, processing arrangements have some specific features directly related to our theoretical modelling of RCs. In particular, under fixed-fee agreements, Chinese companies export the finished goods and receive only a fixed processing fee. This fee is proportional to the number of processed products, which in turn limits the scope for bargaining and is close in spirit to the bonus payment featured in our model. Moreover, another related feature of these arrangements is that it allow the US importer to send equipments and machineries to the Chinese suppliers duty-free. The machineries have to be used in the production line specifically dedicated to the production of goods for the respective foreign (US) partner, and thus, these transactions indicate a strong interlinked relationship between the buyer and the supplier. Such processing trade arrangements with supplied equipment therefore capture the essence of RCs fairly well in our view, and they are also similar in spirit to the specific measures of RCs that Marchiavello and Morjaria (2014) use in their case study.

[18] Ideally, we would like to observe if the Chinese exporter still deals with the same US importer, but the data do not allow us to do so. Nevertheless, our duration measure provides a very similar pattern 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 (43.3%).

[19] More detailed about processing trade arrangement and their specific features can be found in the recent papers by Fernandes and Tang (2015) and Manova and Yu (2014).

Summing up, in our empirical analysis we consider two specific proxies for the prevalence of RCs in an industry. The first and somewhat wider measure is the overall share of all transactions realized using processing trade arrangements. Then, as a second and more narrow measure, we only consider processing arrangements with fixed-fee agreements and/or with imported equipments supplied by the foreign party, and then build the share of these particular transactions among all transactions within the respective industry.

5.2 Descriptive evidence

Table 1 reports the duration of the transactions under processing and ordinary trade arrangements and reveals quite striking differences. With processing trade, almost 59 % of the transactions last three years or more, while only 17 % do not go beyond the first year. By contrast, the majority of ordinary arrangements (47 %) actually seem to be one-shot deals and only 28.4 % of them have a duration of more than three years.

Digging deeper into the different types, we find that 20% of all processing trade transactions are either a fixed-fee arrangement or one with supplied equipment by the US partner; around 13% of them actually have both features. Table 2 then provides similar statistics for the durations of these two sub-types, with fixed fees (columns 1 and 2) and with supplied equipments (columns 3 and 4). As can be seen, only 14 to 18 % of those transactions last for less than a year, while 55 to 60 % seem to be resilient long-term collaborations. The share of LTCs within these two subgroups is, thus, rather similar to the one among all processing trade transactions.

Table 1: Duration of processing and non-processing transactions

Duration of	With Processing		Without Processing	
transactions	# transactions	Percentage	# transactions	Percentage
< 1 year	1,325	17.0	26,247	47.0
> 1 year	1,047	13.5	8,451	15.1
> 2 years	840	10.8	5,268	9.4
> 3 years	4,561	58.7	15,841	28.4
Total	7,773	100.0	55,807	100.0

Table 2: Duration of processing transactions with fixed-fees and supplied equipments

	Processing		Processing	
Duration of	with fixed-fees		with supplied equipments	
transactions	# transactions	Percentage	# transactions	Percentage
< 1 year	445	18.3	160	14.1
> 1 year	357	14.7	160	14.1
> 2 years	279	11.5	132	11.6
> 3 years	1,353	55.6	684	60.2
Total	2,434	100.0	1,136	100.0

An important observation is that quite a few of the Chinese exporters in our data are owned by multinational corporations. It is important to take these ownership structures into account, as Chinese foreign affiliates may differ systematically in their economic behaviour from independent contractors.[20] Tables 3 and 4 are analogous to Tables 1 and 2, but exclude all transactions from Chinese foreign affiliates which are dependent subsidiaries or part of a joint venture.

Table 3: Duration of processing and non-processing transactions - excluding multinational affiliates

Duration of	With Processing		Without Processing	
transactions	# transactions	Percentage	# transactions	Percentage
< 1 year	430	18.6	23,344	49.7
> 1 year	334	14.5	7,262	15.4
> 2 years	290	12.6	4,344	9.2
> 3 years	1,254	54.3	12,063	25.7
Total	2,308	100.0	47,013	100.0

Table 4: Duration of processing transactions with fixed-fees and supplied equipments - excluding multinational affiliates

	Processing		Processing	
Duration of	with fixed-fees		with supplied equipments	
transactions	# transactions	Percentage	# transactions	Percentage
< 1 year	274	16.9	143	13.6
> 1 year	235	14.5	145	13.8
> 2 years	203	12.5	126	12.0
> 3 years	906	56.0	638	60.6
Total	1,618	100.0	1,052	100.0

Two interesting facts emerge: First, while the number of transactions using ordinary trade decreases by only 15.8% (from 55,807 to 47,013 transactions), the overall number of processing trade transactions drops dramatically by 70.3% (from 7,773 to 2,308), although this drop is less pronounced for the two sub-types of processing trade with fixed-fee and/or supplied equipment. In other words, the use of processing trade seems to be particularly common within the boundaries of the firm, but it is at the same time not limited to intra-firm trade relationships. Second, the data suggest that foreign affiliates' transactions are a bit more resilient. In fact, dropping the foreign affiliates from the sample decreases the fraction of long-term collaborations and increases the share of one-shot transactions by 2-4 percentage points both for processing and ordinary trade. However, it has basically

[20]This ownership dimension does not feature in our theoretical model, where we assume that the supplier maintains all property rights over his assets. Notice, however, that 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 (also see Kukharskyy and Pflüger, 2011). To reflect this, we conduct our estimations on the overall sample of all Chinese exporters and in our robustness checks, on a reduced sample where all Chinese foreign affiliates are excluded.

no impact on the duration of processing transactions with fixed fees or with supplied equipments. More importantly, it does not alter the sharp differences in duration between ordinary and processing trade transactions. Even when we exclusively focus on independent Chinese exporters, we still find a clear pattern that the latter type of arrangements have substantially longer durations. This is particularly clearly visible in Table 4: For the type of processing trade with fixed-fee and/or supplied equipment we find that the share of long-term collaborations is substantially higher than on average.

Summing up, the Chinese customs data shows that processing trade transactions have higher durations than ordinary ones, thus suggesting a positive correlation between match durations and our empirical proxy for RCs. This is especially true when a fixed-fee agreement and/or the supply of equipment by the US partner is involved, which we believe are particularly good indicators that the two partners engage in relational contracting.

5.3 Ordinary least squares (OLS) regressions

To address this correlation in more detail, we next present some OLS results at the individual transaction-level. Specifically, we use all new transactions by the fresh Chinese exporting firm i in industry j. We then regress a dummy variable, which takes the value one if the transaction was carried out under a processing trade arrangement, on a dummy variable indicating if the transaction is still observed in $t+3$ (value 1) or if it has ended in the meantime (value 0), while controlling for the ownership status of the Chinese firm. In order to obtain transaction-level estimates that are directly comparable to the industry-level estimations presented below, we weight each transaction such that all products j receive the same overall weight in the regression.

The results are reported in Table 5. Column 1 indicates a strong positive correlation between the use of processing agreements and the dummy for transaction continuation. In fact, when the collaboration is long-term (LTC), it has a 8.9 percentage points higher probability of being a processing trade arrangement. In column 2, we introduce a set of industry-level control variables, such as the capital-labor ratio, the human capital intensity, the R&D-sales ratio and the relationship-specificity by Nunn (2013) which are standard measures for the headquarter-intensity η from the model. In order to capture α, we also include products' demand elasticity taken from Broda and Weinstein (2006).[21] The introduction of these additional control variable leaves the coefficient unchanged, and this also remains true if we were to introduce product-fixed effects. Finally, in Columns 3 and 4, we use the narrower proxy for the use of RCs. Specifically, we replace the dependent variable by another dummy indicator taking the value 1 if the transaction is a processing trade arrangement with fixed-fees and/or with supplied equipments by the foreign partner.[22] The same picture emerges as for the wider proxy: If the relationship is a long-term collaboration, it is more likely a RC.

[21] As data are disaggregated at the HS 10 digit level, we use the median value for each HS6 product.

[22] Notice how the R^2 largely decreases when using our narrower RC proxy (see columns 3 and 4). Also observe that the coefficient associated with foreign ownership is much smaller when using our second proxy for RCs. In fact, foreign ownership is a good predicator of the use of processing arrangements, but mainly when the processing transactions do not make use of fixed-fees and supplied equipment.

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

Dependent variable:	Processing transactions		Processing with Fixed-fees and/or supplied equipments	
LTC dummy	0.089***	0.088***	0.031***	0.031***
	(0.006)	(0.005)	(0.003)	(0.003)
Foreign ownership	0.286***	0.282***	0.018***	0.017***
	(0.009)	(0.008)	(0.004)	(0.004)
ln Capital/Worker		-0.031***		-0.012***
		(0.006)		(0.004)
ln Skilled/Worker		-0.093***		-0.047***
		(0.020)		(0.011)
ln R&D Sales		-0.001		-0.010***
		(0.004)		(0.002)
Nunn measure		0.229***		0.129***
		(0.039)		(0.020)
ln demand elasticity		0.019**		0.009**
		(0.007)		(0.004)
Observations	63,580	63,580	63,580	63,580
R^2	0.162	0.176	0.007	0.018

Notes: Dependent variable: Dummy variable which takes the value one in the case part or all the transaction is realized under processing agreement, and zero otherwise. In column 3-4, the binary variable takes a value one for processing agreements with fixed-fees and/or supplied equipments, and zero otherwise. Each transaction is weighted such that each HS6 digit product is given the same weight. Robust standard-errors in brackets using clustered standard errors at the HS6 level.

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

Dependent variable:	Share of processing transactions		Share with fixed-fees and/or supplied equipments	
LTC Share	0.088***	0.083***	0.050***	0.051***
	(0.024)	(0.024)	(0.015)	(0.014)
Foreign ownership	0.517***	0.492***	0.089***	0.085***
	(0.033)	(0.030)	(0.017)	(0.017)
ln Capital/Worker		-0.020***		-0.009***
		(0.006)		(0.004)
ln Skilled/Worker		-0.081***		-0.044***
		(0.018)		(0.011)
ln R&D Sales		-0.003		-0.011***
		(0.003)		(0.002)
Nunn measure		0.212***		0.132***
		(0.038)		(0.019)
ln demand elasticity		0.014**		0.009**
		(0.006)		(0.004)
Observations	1,004	1,004	1,004	1,004
R^2	0.304	0.363	0.047	0.131

Notes: Dependent variable: In column 1-2, the dependent variable is the share of exports transactions using processing agreement on the overall number of transactions within a HS6 product. In column 3-4, the dependent variable is the share of exports transactions using processing agreement with fixed-fees and/or supplied equipments. Robust standard-errors in brackets using clustered standard errors at the HS6 product level.

In Table 6, we provide analogous estimations after aggregating the data to the HS6 product level. In columns 1-2, we regress the share of processing trade arrangements in industry j on the share of long-term transactions (lasting 3 years or more). In columns 3-4, we use the share of processing transactions with fixed-fee and/or supplied equipment. Moreover, in columns 2 and 4 we introduce the additional industry-level controls.

Also with this approach we obtain a robust positive correlation between the share of long-term transactions and the use of processing trade agreements in an industry j, in the range between 5 and 9 percentage points. This supports the positive impact of long-term collaborations on the use of relational contracts as featured in our model.

5.4 Instrumental variable (IV) estimation

As explained in Section 4, the correlations from Tables 5 and 6 may not capture the *causal* effect of LTCs on RCs, as they may include confounding effects running in the opposite direction. Moreover, other sources of endogeneity may now also arise given the specific proxies for the use of RCs that we have chosen for our empirical analysis, i.e., different reasons not captured in our model why processing trade might affect match durations.

To address those issues, we now consider the IV approach discussed above. In order to construct an empirical proxy for the industry-wide cost dispersion (our instrument), 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 HS6 industry during our period of observation. As usual, marginal costs are difficult to observe directly. Our underlying assumption, thus, 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 the results of our model, where re-matching is more likely the higher are the marginal costs of the current supplier.[23]

The results are presented in Table 7. For ease of exposition, for the first-stage, we only report the coefficient and the standard errors associated with our instrument.[24] Following Angrist and Pischke (2009, pp. 217-18), we furthermore report their conditional F-statistic which should range above the cutoff value of 10 according to Stock and Yogo (2005). Table 7 shows that cost dispersion, indeed, seems to be a strong instrument for match durations. This is the case for the baseline specifications in columns 1 and 3, and also when the other control variables are included (columns 2 and 4).

[23]Manova and Zhang (2012) argue that differences in unit-values could also reflects quality variations. In this case, our empirical proxy would be capturing heterogeneity in suppliers' quality. However, under iso-elastic demand and monopolistic competition, a marginal cost or a quality differences are isomorphic in the sense that they both 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 increase of LTCs. Hence, the interpretation of our results would be robust to the way we interpret differences in unit-values.

[24]We assume that the variables are continuous rather than bounded variables (between zero and one) and that their respective structural equations are linear. This could be relaxed by implementing a fractional logit.

Table 7: Relational-contracts and long-term collaborations: Instrumental variable approach

Dependent variable: (2nd stage)	Share of processing transactions		Share with fixed-fees and/or supplied equipments	
LTC Share	0.476**	0.681***	0.437***	0.403***
	(0.191)	(0.177)	(0.128)	(0.105)
Foreign ownership	0.438***	0.375***	0.011	0.016
	(0.048)	(0.045)	(0.032)	(0.029)
ln Capital/Worker		-0.020**		-0.009*
		(0.008)		(0.005)
ln Human cap/Worker		-0.085***		-0.047***
		(0.021)		(0.013)
ln R&D Sales		-0.007		-0.013***
		(0.004)		(0.003)
Nunn measure		0.190***		0.119***
		(0.050)		(0.029)
ln demand elasticity		0.013*		0.008*
		(0.007)		(0.004)
Dependent variable: (1st stage)	Share of long-term transactions			
Unit value SD	-0.045***	-0.060***	-0.045***	-0.060***
	(0.010)	(0.011)	(0.010)	(0.011)
Observations	1,004	1,004	1,004	1,004
F-Stats	22.28	31.94	22.28	31.94

Notes: Dependent variable: In column 1-2, the dependent variable is the share of exports transactions using processing agreement on the overall number of transactions within a HS6 product. In column 2-4, the dependent variable is the share of exports transactions using processing agreement with fixed-fees and/or supplied equipments. Robust standard-errors in brackets using clustered standard errors at the HS6 product level.

Overall, the IV estimation approach yields empirical results firmly in line with the predictions of our theory. In particular, in the first-stage we find that higher cost dispersion negatively affects match durations, which is consistent with Proposition 3. This negative effect is robust across all specifications and statistically highly significant. In the second-stage, we then confirm the positive impact of LTCs on RCs. Quantitatively, the results suggest that, going from an industry with only one-shot to an industry with only repeated interactions, is associated with a sizeable increase of RC arrangements by 40 to 68 percentage points, depending on the set of controls and which proxy for RCs is used. These numbers are substantially higher than the corresponding simple OLS correlations (which range between 5 and 9 percentage points), and this downward bias is consistent with the intuition of our theory: the model predicts that more patient agents are, per se, more prone towards RCs and also tend to search more. Thus, they are likely to have shorter match durations in the time period we observe in the data. The IV approach shuts down this reverse causality, and thus, yields a stronger positive effect of LTCs on RCs. Stated differently, the coefficients in Table 7 are only identified by the variation in match durations induced by cost dispersion. Hence, we can think of them as capturing the causal effect of launching a LTC, by deciding to stop searching, on the introduction of a RC between the US importer and the Chinese exporter.

5.5 Robustness checks and discussion

Tables 8 and 9 provide a battery of robustness checks. As before, the left half uses the wide proxy for RCs as the dependent variable, namely the overall share of processing trade arrangements among all transactions, while the right half considers the narrower proxy, the processing trade arrangements with fixed-fees and/or supplied equipments.

First, in Columns 1 and 5 of Tables 8, 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 as we still exclude all products with less than 10 observations. The main coefficient for the impact of LTCs on RCs slightly decreases, but it remains significant at the 1% level, and the F-statistic decreases somewhat to 14.29. Overall, this indicates that our results remain robust also in a smaller sample only consisting of independent Chinese-owned component manufacturers.

Our second robustness check consists of dropping all the HS6 products with zero processing transactions. In fact, our theoretical model predicts that some industries populated by impatient agent (with $\delta < \underline{\delta}$), never experience any relational contracting. And indeed, about 23 percent of all HS6 products do not have any processing trade transactions. Columns 2 and 6 show that the main coefficient for the LTC share increases when those industries are excluded, but the qualitative results remain the same.

Third, since our model focuses on exported intermediate inputs, we exclude in Columns 3 and 7 all the products which are classified as "final goods" according to the UN classification of Broad Economic Categories (BEC). Notice that even those goods could be intermediate inputs in the sense that they are not directly sold to American consumers, but still undergo some branding or minor modification by the US headquarter. By excluding those cases, we apply a stricter interpretation of what a manufacturing component actually is. Still, also within this sample we find that the coefficients associated with LTC share remain in the same ballpark and stay significant at the 1 percent level.

Fourth, since many firms export multiple products, we exclude all the transactions which are not part of the firm's main sector of activity which we define by the 2-digit HS product for which we observe the largest number of exported HS6 products. In Columns 4 and 8 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.

Fifth, in column 1 and 5 of Tables 9, we use the unit values of export transactions to Japan instead of the US to construct the industry-specific instruments which should, from a theoretical perspective, be independent of the export destination. We choose Japan as it is the second main destination of Chinese exports after the US (besides Hong-Kong and Macao). In addition, it has a similar level of contracting environment, which makes the two countries comparable in this sense. It turns out that the results remain very similar, and the associated instrument is also quite strong with a large F-statistic of 27.19.

Sixth, we use the export transactions to Japan to build our dependent variables (RCs, wide and narrow definition) and our main explanatory variable (the LTC share), while using the transactions in the US market to build the instrument. The results are again very similar and the instrument remains strong. Our findings thus do not seem to be specific to US-China trade, but also hold for Chinese trade with other developed countries.

Table 8: Relational-contracts and long-term collaborations: Industry-level

	Drop foreign	Drop zero processing	Drop final products	Main sector	Drop foreign	Drop zero processing	Drop final products	Main sector
Dependent variable: (2nd stage)	Share of processing transactions				Share with fixed-fees and/or supplied equipments			
LTC Share	0.492***	0.825***	0.629***	0.763***	0.380**	0.504***	0.177**	0.433***
	(0.187)	(0.236)	(0.236)	(0.266)	(0.155)	(0.177)	(0.089)	(0.150)
Foreign ownership		0.311***	0.450***	0.454***		-0.067	0.062***	0.035
		(0.060)	(0.046)	(0.041)		(0.066)	(0.021)	(0.027)
ln Capital/Worker	-0.015**	-0.001	-0.017*	-0.013	-0.013**	0.008	-0.006*	-0.007
	(0.007)	(0.013)	(0.009)	(0.012)	(0.006)	(0.011)	(0.004)	(0.007)
ln Skilled/Worker	-0.088***	-0.089***	-0.083***	-0.138***	-0.078***	-0.059***	-0.031***	-0.063***
	(0.024)	(0.025)	(0.024)	(0.029)	(0.021)	(0.022)	(0.011)	(0.017)
ln R&D Sales	-0.005	-0.004	-0.009	-0.016*	-0.005	-0.011***	-0.003	-0.019***
	(0.004)	(0.006)	(0.006)	(0.008)	(0.003)	(0.004)	(0.003)	(0.005)
Nunn measure	0.136***	0.153**	0.117**	0.251***	0.124***	0.144**	0.026	0.132***
	(0.039)	(0.070)	(0.052)	(0.062)	(0.030)	(0.063)	(0.019)	(0.034)
ln elasticity	0.019***	0.015	0.013*	0.017*	0.012**	0.015*	0.001	0.010*
	(0.007)	(0.011)	(0.007)	(0.010)	(0.005)	(0.008)	(0.003)	(0.006)
Dependent variable: (1st stage)	Share of long-term transactions							
Unit value SD	-0.047***	-0.059***	-0.048***	-0.055***	-0.047***	-0.059***	-0.048***	-0.055***
	(0.012)	(0.014)	(0.012)	(0.013)	(0.012)	(0.014)	(0.012)	(0.013)
Observations	809	773	665	765	809	509	665	765
F-Stats	14.29	17.41	16.33	16.89	14.29	17.41	16.33	16.89

Notes: Dependent variable: In column 1-4, the dependent variable is the share of exports transactions using processing agreement on the overall number of transactions within a HS6 product. In column 5-8, the dependent variable is the share of exports transactions using processing agreement with fixed-fees and/or supplied equipments. Robust standard-errors in brackets using clustered standard errors at the HS6 product level.

Seventh, we go back to our original specification from Table 7, but exclude all the transactions with processing arrangements when constructing our instrumental variable. Hence, our instrument becomes the standard-deviation of the log of the unit-value of transactions using ordinary trade only. We do so, since one might worry that our baseline instrument captures differences in pricing rules associated with processing arrangements. Results from Column 3 and 7 suggest, however, that this is not the case: all results remain unchanged or, if anything, become even stronger.

Eighth, we exclude all short-term transactions (less than 3 years) when constructing the instrument. Here again, one may be worried that short-term transactions may be priced differently, and so may affect our cost dispersion measure. Yet again, columns 4 and 8 yield very similar results.

Finally, we have conducted several further robustness checks for which we omit the detailed results. For example, we have changed the definition of long-term collaborations, and use only transactions that started in 2001 or 2002 which lasted for more than four years. Yet, it turns out that our results remain robust: we find a positive effect of LTCs on RCs, with cost dispersion being a valid instrument for match durations.

Table 9: Relational-contracts and long-term collaborations: Industry-level

Data	US	Japan	US	US	US	Japan	US	US
Instrument	Japan	US	No Proc	No LT	Japan	US	No Proc	No LT
Dependent variable: (2nd stage)	Share of processing transactions				Share with fixed-fees and/or supplied equipments			
LTC Share	0.813***	0.693***	0.487***	0.809***	0.497***	0.811***	0.338***	0.512***
	(0.208)	(0.175)	(0.151)	(0.238)	(0.132)	(0.170)	(0.091)	(0.145)
Foreign ownership	0.349***	0.421***	0.413***	0.350***	-0.003	0.190***	0.028	-0.006
	(0.052)	(0.043)	(0.040)	(0.056)	(0.035)	(0.041)	(0.026)	(0.038)
ln Capital/Worker	-0.019**	-0.055***	-0.020***	-0.019**	-0.009	-0.052***	-0.009**	-0.009
	(0.009)	(0.013)	(0.007)	(0.009)	(0.006)	(0.012)	(0.005)	(0.006)
ln Skilled/Worker	-0.086***	-0.120***	-0.084***	-0.086***	-0.047***	-0.033	-0.046***	-0.047***
	(0.023)	(0.028)	(0.020)	(0.023)	(0.015)	(0.030)	(0.012)	(0.015)
ln R&D Sales	-0.007	-0.008	-0.005	-0.007	-0.014***	-0.033***	-0.013***	-0.014***
	(0.005)	(0.006)	(0.004)	(0.005)	(0.003)	(0.006)	(0.002)	(0.003)
Nunn measure	0.186***	0.257***	0.197***	0.186***	0.116***	0.347***	0.122***	0.115***
	(0.056)	(0.071)	(0.043)	(0.056)	(0.034)	(0.069)	(0.026)	(0.034)
ln elasticity	0.012	0.036***	0.013*	0.012	0.008*	0.030***	0.008**	0.008*
	(0.008)	(0.010)	(0.007)	(0.008)	(0.005)	(0.009)	(0.004)	(0.005)
Dependent variable: (1st stage)	Share of long-term transactions							
Unit value SD	-0.051***	-0.078***	-0.062***	-0.050***	-0.051***	-0.078***	-0.062***	-0.050***
	(0.010)	(0.012)	(0.010)	(0.011)	(0.010)	(0.012)	(0.010)	(0.011)
Observations	1,004	860	1,004	1,004	1,004	860	1,004	1,004
F-Stats	27.19	45.76	35.16	19.94	27.19	45.76	35.16	19.94

Notes: Dependent variable: In column 1-4, the dependent variable is the share of exports transactions using processing agreement on the overall number of transactions within a HS6 product. In column 5-8, the dependent variable is the share of exports transactions using processing agreement with fixed-fees and/or supplied equipments. Robust standard-errors in brackets using clustered standard errors at the HS6 product level.

6 Conclusion

In this paper, we have developed a dynamic property rights model of global sourcing where a domestic headquarter seeks to obtain an intermediate input from a foreign input supplier. Both, the specific foreign partner and the mode of play with him are endogenously determined in our model, and we have shown that the option to re-match crucially affects the timing when an efficient relational contract can be implemented. In particular, our model shows that there is a positive *causal* relationship between match durations (LTCs) and the stability of relational contracts (RCs). This key prediction of our model is empirically supported in the context of US input sourcing from China. The data show that more enduring export transactions from China are more likely to involve elements of relational contracting, which we have captured by particular types of processing trade arrangements. We have then developed a theory-driven instrumental variable estimation approach, using industry-wide cost dispersion as an instrument, which supports our findings.

Our model is simple and can be extended in various directions. For example, we have focussed on the decision of one single headquarter firm and assumed that potential suppliers are abundantly available. While this assumption seems plausible on average in the context of US input sourcing from China, there are also cases with large and powerful suppliers (e.g. Foxconn) which have good outside options and bargaining power, and which actively seek to collaborate with the best possible headquarter firms. Such scenarios could

be captured in a model with heterogeneity on both sides, and with assortative matching of headquarters and suppliers. However, our key result – the positive causal effect of match durations on relational contracting – arises already in our simpler environment.

On the empirical side, one has to deal with the fact that match-specific contractual arrangements between a buyer and a supplier are often tacit and implicit, and thus, by their very nature unobservable. Existing work on relational contracts has tackled this difficulty in detailed case studies about particular industries in developing countries. In this paper, we have taken a different avenue and considered all Chinese manufacturing exports in a particular time frame, and worked with proxies for relational contracts that capture – in our view – the essence of such trust-based cooperative arrangements fairly well. This approach based on representative panel data by the Chinese Customs authority might open the door for many further studies, especially if it can be matched with the Customs data of further countries.

References

[1] Abreu, D. On the theory of infinitely repeated games with discounting. *Econometrica 56*, 2 (1988), 383–396.

[2] 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.

[3] Angrist, J. D., and Pischke, J.-S. *Mostly harmless econometrics: An empiricist's companion.* Princeton University Press, 2009.

[4] Antràs, P. Firms, contracts, and trade structure. *The Quarterly Journal of Economics 118*, 4 (2003).

[5] Antràs, P. *Global Production: Firms, Contracts and Trade Structure.* 2015.

[6] Antràs, P., and Costinot, A. Intermediated trade. *The Quarterly Journal of Economics 126* (2011), 1319–1374.

[7] Antràs, P., and Helpman, E. Global sourcing. *Journal of Political Economy 112*, 3 (2004), 552–580.

[8] Araujo, L., Mion, G., and Ornelas, E. Institutions and export dynamics. Working paper, 2014.

[9] 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.

[10] Bernard, A. B., and Dhingra, S. Contracting and the division of the gains from trade. Working paper, 2015.

[11] 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.

[12] Besedeš, T. A search cost perspective on formation and duration of trade. *Review of International Economics 16*, 5 (2008), 835–849.

[13] Board, S. Relational contracts and the value of loyalty. *American Economic Review 101*, 7 (2011), 3349–67.

[14] Board, S., and Meyer-ter Vehn, M. Relational contracts in competitive labour markets. *The Review of Economic Studies 82*, 2 (2015), 490–534.

[15] Broda, C., and Weinstein, D. E. Globalization and the gains from variety. *The Quarterly Journal of Economics 121*, 2 (2006), 541–585.

[16] Carballo, J., Ottaviano, G., and Volpe-Martincus, C. The buyer margins of firms' exports. Discussion Paper 1234, CEP, 2013.

[17] 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.

[18] Dwyer, R. F., Schurr, P. H., and Oh, S. Developing buyer-seller relationships. *The Journal of Marketing 51*, 2 (1987), 11–27.

[19] Eaton, J., Eslava, M., Jinkins, D., Krizan, C. J., and Tybout, J. A search and learning model of export dynamics. Working paper, 2014.

[20] Feenstra, R. C., and Hanson, G. H. Ownership and control in outsourcing to china: Estimating the property-rights theory of the firm. *The Quarterly Journal of Economics 120*, 2 (2005), 729–761.

[21] Fernandes, A. P., and Tang, H. Scale, scope, and trade dynamics of export processing plants. *Economics Letters 133* (2015).

[22] Fujiwara-Greve, T., and Okuno-Fujiwara, M. Voluntarily separable repeated prisoner's dilemma. *The Review of Economic Studies 76*, 3 (2009), 993–1021.

[23] Ghosh, P., and Ray, D. Cooperation in community interaction without information flows. *The Review of Economic Studies 63*, 3 (1996), 491–519.

[24] Green, E. J., and Porter, R. H. Noncooperative collusion under imperfect price information. *Econometrica 52*, 1 (1984), 87–100.

[25] 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.

[26] Hart, O. D., and Moore, J. Property rights and the nature of the firm. *Journal of Political Economy 98*, 6 (1990), 1119–58.

[27] Jackson, M. O., and Watts, A. Social games: Matching and the play of finitely repeated games. *Games and Economic Behavior 70*, 1 (2010), 170–191.

[28] Johnson, R. C., and Noguera, G. Accounting for intermediates: Production sharing and trade in value added. *Journal of International Economics 86*, 2 (2012), 224–236.

[29] 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.

[30] Kamal, F., and Tang, H. Relational trade: Evidence from the united states. Mimeo, 2015.

[31] Kukharskyy, B. Quality of institutions and global sourcing. Mimeo, 2015.

[32] Kukharskyy, B., and Pflüger, M. Relational contracts and the economic well-being of nations. Discussion paper, BGPE, 2011.

[33] Lehrer, E., and Pauzner, A. Repeated games with differential time preferences. *Econometrica 67*, 2 (1999), 393–412.

[34] Macchiavello, R., and Morjaria, A. Competition and relational contracts: Evidence from rwanda's coffee mills. Working paper, 2014.

[35] Manova, K., and Yu, Z. Firms and credit constraints along the global value chain: Processing trade in china. Working paper, 2014.

[36] Manova, K., and Zhang, Z. Export prices across firms and destinations. *The Quarterly Journal of Economics 127* (2012), 379–436.

[37] Melitz, M. J., and Redding, S. J. Heterogeneous firms and trade. *Handbook of International Economics 4* (2014), 1–54.

[38] Monarch, R. It's not you, it's me: Breakups in u.s.-china trade relationships. Working paper, 2015.

[39] Nunn, N. Relationship-specificity, incomplete contracts, and the pattern of trade. *The Quarterly Journal of Economics 122*, 2 (2007), 569–600.

[40] Rauch, J. E. Networks versus markets in international trade. *Journal of International Economics 48*, 1 (1999), 7–35.

[41] Rauch, J. E., and Watson, J. Starting small in an unfamiliar environment. *International Journal of Industrial Organization 21*, 7 (2003), 1021–1042.

[42] Rob, R., and Yang, H. Long-term relationships as safeguards. *Economic Theory 43*, 2 (2010), 143–166.

[43] Stock, J. H., and Yogo, M. Testing for weak instruments in linear iv regression. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg* (2005).

APPENDIX

A "Carrot-and-stick" punishment

This Appendix documents that the main result from Proposition 1, namely the independence of the critical discount factor of A, c_m, and c_h, also holds for the case of optimal penal codes. Consider the following carrot-and-stick strategy profile along the lines of Abreu (1988): The game starts in the cooperative state. After a deviation of player i in period t, players min-max each other for $T-1$ periods, where $T \geq 2$, beginning in period $t+1$. In period $t+T$ the game switches back to the cooperative state if both players carried out the punishment, otherwise the punishment is repeated.

The following expressions characterize the incentive compatibility constraints of the strategy profile, where (IC-on) and (IC-off) denote the constraints of player $i = \{H, M\}$ on the equilibrium path and, respectively, in the first post-deviation period. Note that the min-max-payoffs are attained at $h = m = 0$, which results in punishment payoffs equal to zero for both players.

$$\frac{\pi_i^{RC}}{1-\delta} \geq \pi_i^D + 0 \cdot \sum_{j=1}^{T-1} \delta^j + \frac{\delta^T}{1-\delta}\pi_i^{RC} \qquad (\text{IC}_i\text{-on})$$

$$0 \cdot \sum_{j=0}^{T-2} \delta^j + \frac{\delta^{T-1}}{1-\delta}\pi_i^{RC} \geq 0 \cdot \sum_{j=0}^{T-1} \delta^j + \frac{\delta^T}{1-\delta}\pi_i^{RC} \qquad (\text{IC}_i\text{-off})$$

First, consider (IC-off) and note that any deviation of player i yields at most zero when restarting the punishment. Simplifying gives $\pi_i^{RC} \geq \delta\pi_i^{RC}$, which is trivially satisfied for any discount factor. Now, consider (IC-on). Simplifying gives:

$$(1-\delta^T)\pi_i^{RC} \geq (1-\delta)\pi_i^D$$

We define the critical discount factor under the carrot-and-stick strategy as $\hat{\delta}$. Analogous to the discussion of Nash reversion in the main text, at $\delta = \hat{\delta}$, both (IC_H-on) and (IC_M-on) hold with equality. Plugging in π_i^{RC}, solving the resulting expressions for the bonus payment $B^*(\hat{\delta})$, and equalizing gives the following implicit expression for $\hat{\delta}$:

$$\hat{\delta}^T\pi^{JFB} - \hat{\delta}(\pi_H^D + \pi_M^D) + (\pi_H^D + \pi_M^D - \pi^{JFB}) = 0 \qquad (17)$$

From (17) it is evident that $\hat{\delta}$ is decreasing in T. That is, the longer the punishment phase, the more stable becomes the cooperative RC. Even more importantly, since π^{JFB} and π_i^D are homogeneous of a common degree b, it also follows that the critical discount factor $\hat{\delta}$ is independent of c_m, as well as of A and c_h. The cost-orthogonality of the critical discount factor therefore also holds with this alternative penal code.

Imperfect monitoring and demand uncertainty. We now study the robustness of the cost-orthogonality result from Proposition 1 when the headquarter can only imperfectly monitor the quality of the supplier's input and demand is uncertain in a similar

way as in the seminal model by Green and Porter (1984). Specifically, suppose demand realizations are stochastic and i.i.d. over periods, and in each period demand is either in a low state ($A = 0$) with probability θ, or in a high state ($A = 1$) with probability $1 - \theta$. The firm and the supplier do not know the state of demand when they make their input investments, nor can they infer it at any later point. Second, suppose the supplier M_0 has the option to supply any quantity of the input m either with a high quality ($I = 1$) or with a low quality ($I = 0$). With $I = 1$, unit costs are c_m^0 as before but the low-quality input can be supplied at zero costs for the supplier. Input quality cannot be inferred by the headquarter at any time, hence, the firm cannot disentangle if zero revenue in the last step of a stage game is due to a low state of demand or to low quality of the input. To study this variation of our model, we adjust the *stage game* as follows:

1'. **Proposal stage** (cheap talk): H can make M a non-binding and non-contractible proposal specifying investment levels (h,m) and an ex-post bonus payment B to M.

2'. **Participation decision stage**: The supplier M decides upon his participation in the relationship with H according to his outside option ω_M.

3'. **Investment stage**: The headquarter H and the supplier M simultaneously choose their non-contractible input investments (h,m). Additionally, M decides on the quality $I \in \{0, 1\}$ of his input. With $I = 0$, the input m is fully incompatible for the production of the final output and M incurs zero costs of input provision. With $I = 1$, the supplier incurs unit costs c_m as before and input m is usable for production.

4'. **Information stage**: H and M learn the investment level (quantity) of their production partner, but H cannot observe the quality I of the input.

5'. **Bargaining stage**: If a relational contract was proposed, H can decide to pay the bonus B to M. The bonus payment is made immediately (liquidity constraints are ruled out). Otherwise, the surplus is split according to an asymmetric Nash bargaining conditional on the revenue realized in 6', where $\beta \in (0, 1)$ is H's and $(1 - \beta)$ is, respectively, M's bargaining power.

6'. **Profit realization stage**:

- If $A = 0$ and/or $I = 0$ no final output can be produced and revenue is zero.
- If $A = 1$ and $I = 1$, the final output is produced and sold. The surplus is divided as specified in 5'.

We consider the following strategy profile. The game starts in the *cooperative state* in which behaviour is essentially identical to the one described in the main text. Additionally, M sets $I = 1$ in the cooperative state. Now, whenever i) an observable deviation from the RC occurs, or ii) when no final output can be sold, the game enters a *punishment phase* which lasts for T periods. In the punishment phase, H and M make zero investments and thus follow the "carrot-and-stick" penal code studied above. After the punishment phase ends, the players revert to the cooperative state.

We define by V^{+i} and V^{-i} player i's present discounted value of payoffs in a period in the cooperative state and in a period where the punishment phase has just started:

$$V^{+i} = (1-\theta)(\pi_i^{RC} + \delta V^{+i}) + \theta\delta V^{-i}, \qquad V^{-i} = \delta^T V^{+i}$$

The solution to this system of equations is given by:

$$V^{+i} = \frac{(1-\theta)\pi_i^{RC}}{1-(1-\theta)\delta - \theta\delta^{T+1}}, \qquad V^{-i} = \frac{\delta^T(1-\theta)\pi_i^{RC}}{1-(1-\theta)\delta - \theta\delta^{T+1}}$$

In order to rationalize the relational contract, we need to set up the incentive constraints for H and M. By optimally deviating, H can attain V^{dH} and M can achieve V^{dM}, respectively. Notice that the supplier would never choose the verifiable deviation (supplying a wrong quantity $m \neq m^*$), but if he wants to deviate, he would always prefer to supply the correct quantity m^* in low quality ($I = 0$) as he will then receive the full bonus in the deviation period *before* revenue is realized.

$$V^{dH} = (1-\theta)\pi_H^D + \delta V^{-H}, \qquad V^{dM} = B + \delta V^{-M}$$

The IC-constraints, which generally read $V^{+i} \geq V^{di}$, $i = H, M$, can then be written as:

$$(1-\delta^{T+1})V^{+H} \geq (1-\theta)\pi_H^D \qquad (\text{IC}_H)$$
$$(1-\delta^{T+1})V^{+M} \geq B, \qquad (\text{IC}_M)$$

and plugging in V^{+H} and V^{+M} from above, we can rewrite the constraints as follows, where $\psi \equiv \frac{(1-\delta^{T+1})(1-\theta)}{1-(1-\theta)\delta-\theta\delta^{T+1}}$:

$$\psi\pi_H^{RC} - (1-\theta)\pi_H^D \geq 0$$
$$\psi\pi_M^{RC} - B \geq 0 \qquad (18)$$

Plugging π_M^{RC} into (18) the latter can be rearranged as $B \geq \frac{\psi}{1-\psi}m^*c_m$, where a necessary condition for (IC_M) to hold is $\theta < \frac{\delta-\delta^{T+1}}{1+\delta-2\delta^{T+1}}$. In words, the cooperative RC cannot be sustained if the probability of low demand is too large, similar as in the collusion model by Green and Porter (1984).

Let us now define the critical discount factor above which relational contracting is incentive compatible as $\tilde{\delta}$, and as before observe that at this point both IC-constraints bind with equality. Merging the constraints from (18), we can compute $\tilde{\delta}$ implicitly from the following equation, which inter alia depends also on θ and T:

$$\psi\left(R^* - h^*c_h - \frac{\psi}{1-\psi}m^*c_m\right) - (1-\theta)\pi_H^D = 0 \qquad (19)$$

For patience levels $\delta > \tilde{\delta}$ the cooperative RC can be sustained in this model version, but similar as in Green and Porter (1984), we will observe periods of punishment and zero investments followed by cooperative periods with optimal (high-quality) investments. It is also possible to compute an *optimal* punishment length T^* which maximizes V^{+H} subject

to (19). However, the more important observation for our purpose is that c_m can be factored out of the LHS of (19). From here, it follows immediately that $\tilde{\delta}$ is independent of the unit cost level c_m^0 which reinforces our cost-orthogonality result from Proposition 1.

B Preventing "cheat-and-run"

To rule out "cheat-and-run", the following condition must be satisfied to ensure that the firm stops deviating once it is matched to a low-cost supplier:

$$\frac{1}{1-\delta}\pi_H^{RC,l} \geq E[\pi_H^{cheat-and-run} \mid c_m^0 = c_m^l] \tag{20}$$

Since low-cost suppliers are only found with probability P in every round, the "cheat-and-run" payoff can be computed by the following program:

$$V_0 = \pi_H^{D,l} - F + \delta P V_0 + \delta(1-P)V_1, \qquad V_1 = \pi_H^{D,h} - F + \delta P V_0 + \delta(1-P)V_1,$$

and solving for V_0 gives

$$E[\pi_H^{cheat-and-run} \mid c_m^0 = c_m^l] = \frac{1}{1-\delta}\left[(1-\delta(1-P))\pi_H^{D,l} + \delta(1-P)\pi_H^{D,h} - F\right].$$

With this, we can rewrite condition (20) as follows:

$$F \geq (1-\delta(1-P))\pi_H^{D,l} + \delta(1-P)\pi_H^{D,h} - \pi_H^{RC,l}. \tag{21}$$

Since the right hand side of (21) is increasing in P, "cheat-and-run" can be precluded for all P by requiring

$$F > \pi_H^{D,l} - \pi_H^{RC,l} \equiv \widetilde{F}. \tag{22}$$

For consistency it is necessary to show that the lower bound $\widetilde{F}$ does not per se rule out search and re-matching for $c_m^0 = c_m^h$ in the case of patient agents. For this it is necessary to show that $\widetilde{F} < \overline{F}^2(\delta)$ can hold for every relevant patience level $\delta \in (\underline{\delta}, 1)$. In the following we derive a sufficient condition for $\widetilde{F} < \overline{F}^2(\underline{\delta})$, which then automatically implies $\widetilde{F} < \overline{F}^2(\delta)$ since $\overline{F}^2(\delta)$ is increasing in δ. Rearranging $\widetilde{F} < \overline{F}^2(\underline{\delta})$ gives:

$$P > \frac{(1-\underline{\delta})(\pi_H^{D,l} - \pi_H^{RC,l}(\underline{\delta}) + \pi_H^{RC,h}(\underline{\delta}) - \pi_H^{N,h})}{\underline{\delta}(\pi_H^{RC,l}(\underline{\delta}) - \pi_H^{RC,h}(\underline{\delta}))} \equiv \widetilde{P} \tag{23}$$

$\pi_H^{RC}(\underline{\delta})$ denotes RC-profits at the critical discount factor. We need $\widetilde{P} \in (0,1)$ for $G(c_m)$ to be well-defined. Obviously $\widetilde{P} > 0$. Assuming $\widetilde{P} < 1$ and rearranging we get:

$$(1-\underline{\delta})\left[\pi_H^{D,l} - \pi_H^{N,h}\right] < \pi_H^{RC,l} - \pi_H^{RC,h} \quad \Leftrightarrow \quad (1-\underline{\delta})\pi_H^{D,l} - \pi_H^{RC,l} < (1-\underline{\delta})\pi_H^{N,h} - \pi_H^{RC,h} \tag{24}$$

Notice that both RHS and LHS of this last inequality are negative if the sufficient condition holds. Factorizing c_m^l from the LHS, and c_m^h from the RHS, respectively, we can rewrite

this sufficient condition in the following form:

$$\frac{c_m^h}{c_m^l} > C(\alpha, \beta, \eta), \tag{25}$$

where C is a positive term that depends on α, β, and η. Summing up, if the cost dispersion $\frac{c_m^h}{c_m^l}$ is sufficiently large, condition (25) is satisfied and we have $\widetilde{F} < \overline{F}^2(\delta)$.

C Continuous supplier cost distribution

In the following, we generalize our analysis to the case of a continuous distribution of supplier costs. Specifically, we assume that the cumulative distribution function $G(c_m)$ is continuously differentiable, and denote by $g(c_m)$ the corresponding probability density function with support $[\underline{c_m}, \overline{c_m}]$.

Let $\widetilde{c_m}$ denote the cost realization of the firm's current supplier, and $\widetilde{\pi_H^i}$ is the firm's payoff at this cost realization given the contract structure $i = \{N, RC\}$. In order to derive whether the firm starts re-matching or not in period 0, where $c_m^0 = \widetilde{c_m}$, the same logic as in the main text applies. To start re-matching, it must be the case that the expected gains must be larger than the continuation payoffs with the current match:

$$E[\pi_H^{RM,i} \mid c_m^0 = \widetilde{c_m}] > \frac{1}{1-\delta}\widetilde{\pi_H^i}, \tag{26}$$

where, as before, the two cases of impatient and patient agents must be distinguished that ultimately lead to the contract structure i once the firm has stopped searching.

For the calculation of $E[\pi_H^{RM,i} \mid c_m^0 = \widetilde{c_m}]$, the following decision rule applies: If the firm re-matches after round 0, she collaborate with a new supplier with unit cost level c_m^t in round t, where c_m^t is randomly drawn from $g(c_m)$. If that unit cost is below a certain threshold level, namely $c_m^t \leq c_m^*$, she stops re-matching and sticks with that supplier henceforth. Otherwise, if $c_m^t > c_m^*$, she continues re-matching until a supplier with unit costs below c_m^* is found. Evidently, the firm wants to choose this cost cutoff c_m^* optimally in order to maximize the expected payoff when starting to re-match. In a first step, we set up $E[\pi_H^{RM,i} \mid c_m^0 = \widetilde{c_m}]$ by the following program which we then optimize for c_m^*:

$$\begin{aligned} V_0 &= \widetilde{\pi_H^N} - F + \delta V_1 \\ V_t &= Pr(c_m \leq c_m^*)\frac{E[\pi_H^i \mid c_m \leq c_m^*]}{1-\delta} + Pr(c_m > c_m^*)\left(E[\pi_H^N \mid c_m > c_m^*] - F + \delta V_t\right), \quad t = 1, 2, 3, ... \end{aligned}$$

The term V_t can be rewritten as:

$$V_1 = \frac{1}{1-\delta(1-G(c_m^*))}\left[(1-G(c_m^*))(E[\pi_H^N \mid c_m > c_m^*] - F) + \frac{G(c_m^*)E[\pi_H^i \mid c_m \leq c_m^*]}{1-\delta}\right],$$

and simplifying yields

$$V_1 = \frac{1}{1-\delta(1-G(c_m^*))}\left[-(1-G(c_m^*))F + \int_{c_m^*}^{\overline{c_m}} \pi_H^N g(c_m)\, dc_m + \frac{\int_{\underline{c_m}}^{c_m^*} \pi_H^i g(c_m)\, dc_m}{1-\delta}\right]$$

Consequently, in order to maximize $E[\pi_H^{RM,i} \mid c_m^0 = \widetilde{c_m}]$ the firm will set

$$c_m^* \in \arg\max V_1.$$

Analogously as in Section 4, from (26) we can derive the fixed cost thresholds $\overline{F}^k$, where $k = 1$ for impatient and $k = 2$ for patient agents, respectively. As before, the thresholds are defined in a way such that for all $F < \overline{F}^k$ re-matching is incentive compatible, while otherwise it is not. We can provide the following result:

Proposition 4. *a) Impatient agents ($\delta < \underline{\delta}$): For every $F \in (0, \overline{F}^1)$ there exists an optimal supplier efficiency threshold $c_m^{*1}(F)$ that maximizes $E[\pi_H^{RM,N} \mid c_m^0 = \widetilde{c_m}]$. If the firm draws a supplier with $c_m^t < c_m^{*1}(F)$ in period t, she stops re-matching. With $c_m^t > c_m^{*1}(F)$ she continues re-matching in the following period. The cutoff $c_m^{*1}(F)$ is strictly increasing and concave in F.*

b) Patient agents ($\delta \geq \underline{\delta}$): For every $F \in (\widetilde{F}, \overline{F}^2)$ there exists an optimal supplier efficiency threshold $c_m^{*2}(F)$ that maximizes $E[\pi_H^{RM,RC} \mid c_m^0 = \widetilde{c_m}]$. If the firm draws a supplier with $c_m^t < c_m^{*2}(F)$ in t, she stops re-matching. With $c_m^t > c_m^{*2}(F)$ she continues re-matching in the following period. $c_m^{*2}(F)$ is strictly increasing and concave in F.

Proof. Since V_1 is continuous in c_m^* on the closed and bounded interval $[\underline{c_m}, \overline{c_m}]$ by the Extreme Value Theorem a global maximum of V_1 exists. Consequently, we know that for every F there exists a $c_m^*(F)$ that maximizes V_1 (and hence also $E[\pi_H^{RM,i} \mid c_m^0 = \widetilde{c_m}]$).

Next we show that $c_m^*(F)$ is increasing and concave in F (i.e. $\frac{\partial c_m^*}{\partial F} > 0$, and $\frac{\partial^2 c_m^*}{\partial F^2} < 0$). Partially differentiating V_1 for c_m^* we get the following first-order condition:

$$\underbrace{F - \pi_H^N(c_m^*) + \frac{\pi_H^i(c_m^*)}{1-\delta} - \frac{\delta}{1-\delta(1-G(c_m^*))}\left[-(1-G(c_m^*))F + \int_{c_m^*}^{\overline{c_m}} \pi_H^N g(c_m)\, dc_m + \frac{\int_{\underline{c_m}}^{c_m^*} \pi_H^i g(c_m)\, dc_m}{1-\delta}\right]}_{\equiv H(c_m^*, F)} = 0 \tag{27}$$

By the Implicit Function Theorem we can determine $\frac{\partial c_m^*}{\partial F}$ as:

$$\frac{\partial c_m^*}{\partial F} = -\frac{\frac{\partial H}{\partial F}}{\frac{\partial H}{\partial c_m^*}}$$

Since we know that a maximum at c_m^* exists, $\frac{\partial H}{\partial c_m^*} < 0$. It is easy to see from (27) that $\frac{\partial H}{\partial F} > 0$ must hold. Hence $\frac{\partial c_m^*}{\partial F} > 0$. $\frac{\partial H}{\partial c_m^*}$, $\frac{\partial H}{\partial F}$ and therefore $\frac{\partial c_m^*}{\partial F}$ can also be calculated explicitly from (27). By partially differentiating the resulting expression w.r.t. F it is straightforward to show that $\frac{\partial^2 c_m^*}{\partial F^2} < 0$ holds. □

Intuitively, this result shows that the firm will search less and suffice with a worse current supplier the higher the search costs F are. For completeness, it can also be shown that for all $F \in (0, \overline{F}^1)$, the cutoff $c_m^{*1}(F)$ never lies on the boundaries of the cost interval $[\underline{c_m}, \overline{c_m}]$. To see this, suppose first that $c_m^{*1} = \underline{c_m}$ and note that this implies that re-matching will continue at any level of supplier costs. But this cannot be optimal, since paying $F > 0$ and re-matching only once and earning $E[\pi_H^N]$ in every period from tomorrow onwards is strictly better than paying F and receiving $E[\pi_H^N]$ in every period.

Second, suppose that $c_m^{*1} = \overline{c_m}$ and note that this implies that re-matching does not occur at all. But by definition of $\overline{F}^1$ this cannot be the case, since for all $F < \overline{F}^1$ re-matching must occur at least for some c_m^t. From this we conclude that for all $F \in (0, \overline{F}^1)$, the cutoff $c_m^{*1}(F)$ must be an interior point on $[\underline{c_m}, \overline{c_m}]$. By the same reasoning, for all $F \in (\widetilde{F}, \overline{F}^2)$, the cutoff for patient agents $c_m^{*2}(F)$ never lies on the boundaries of the cost interval. Note that in order to rule out cheat-and-run behaviour (which we extensively discussed for the case of binary cost distribution in Section 4.2) we need to set some lower bound on re-matching fixed costs, $\widetilde{F} > 0$. This guarantees that repeated deviation and re-matching is ruled out on the equilibrium path. By applying the logic of Appendix A, for any $\overline{F}^2 > 0$ it is possible to find a cost distribution $G(c_m)$ that results in $\overline{F}^2 > \widetilde{F}$.

Finally, if search costs happen to be equal to the respective threshold value $\overline{F}^k > 0$, the firm is just indifferent between keeping her current supplier $\widetilde{c_m}$ and re-matching to a new partner. This can be easily shown as follows:

Corollary to Proposition 4.
a) Given that $\overline{F}^1 > 0$, at $F = \overline{F}^1$ it holds that $c_m^{*1}(F) = \widetilde{c_m}$.
b) Given that $\overline{F}^2 > \widetilde{F}$, at $F = \overline{F}^2$ it holds that $c_m^{*2}(F) = \widetilde{c_m}$.

Proof. In a first step we derive $\overline{F}^k$ from (26) by using the same steps as in Section 4. This yields:

$$\overline{F}^k = (1-\delta(1-G(c_m^*)))\left[\widetilde{\pi_H^N} - \frac{\widetilde{\pi_H^i}}{1-\delta}\right] + \delta\left[\int_{c_m^*}^{\overline{c_m}} \pi_H^N g(c_m)\, dc_m + \frac{\int_{\underline{c_m}}^{c_m^*} \pi_H^i g(c_m)\, dc_m}{1-\delta}\right] \quad (28)$$

It now is key to observe that the FOC which determines c_m^*, (27), can be rewritten as

$$F = (1-\delta(1-G(c_m^*)))\left[\pi_H^N(c_m^*) - \frac{\pi_H^i(c_m^*)}{1-\delta}\right] + \delta\left[\int_{c_m^*}^{\overline{c_m}} \pi_H^N g(c_m)\, dc_m + \frac{\int_{\underline{c_m}}^{c_m^*} \pi_H^i g(c_m)\, dc_m}{1-\delta}\right]. \quad (29)$$

Comparing (28) and (29) we can see that at $F = \overline{F}^k$ it must hold that

$$\widetilde{\pi_H^N} - \frac{\widetilde{\pi_H^i}}{1-\delta} = \pi_H^N(c_m^*) - \frac{\pi_H^i(c_m^*)}{1-\delta}. \quad (30)$$

Observing that c_m can be factored out from both sides of (30) the latter equation can be rewritten as

$$\widetilde{c_m}^{\frac{-\alpha(1-\eta)}{1-\alpha}}\phi = (c_m^*)^{\frac{-\alpha(1-\eta)}{1-\alpha}}\phi, \quad (31)$$

where ϕ is a function of all the remaining payoff function parameters and δ. Obviously (31) can only be fulfilled if $c_m^* = \widetilde{c_m}$. $\square$

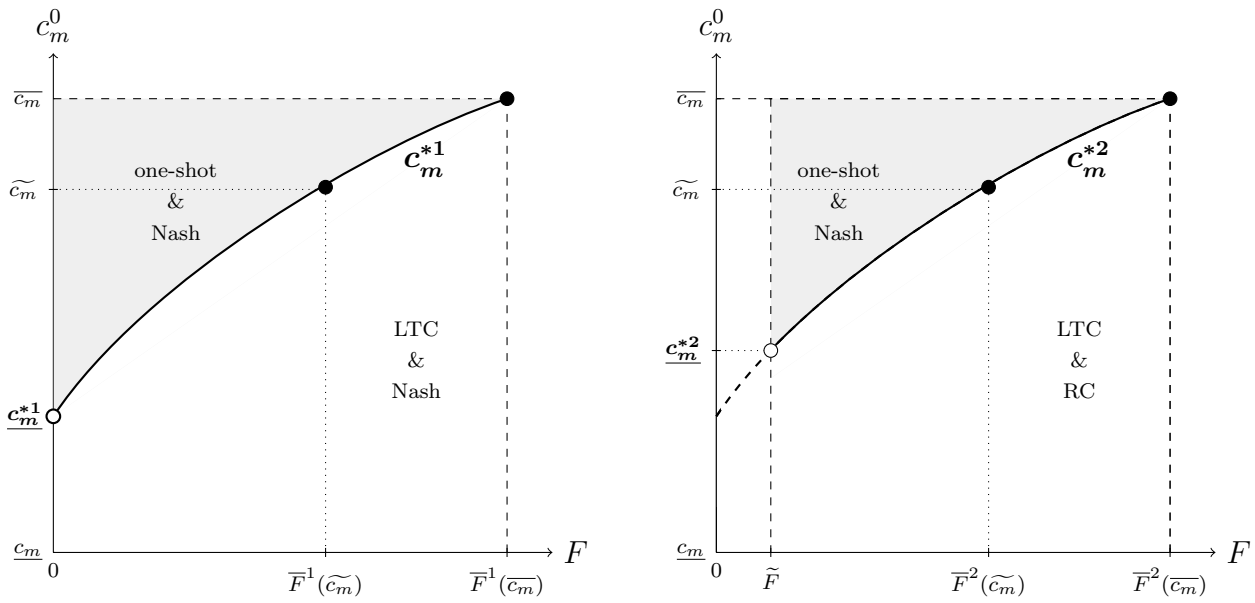

(a) Impatient agents

(b) Patient agents

Figure 3: Optimal re-matching behaviour with continuous supplier cost distribution

Figure 3 summarizes the results for the case of a continuous distribution of supplier costs. By Proposition 4, c_m^{*1} and c_m^{*2} are increasing and concave in F and furthermore we know that $c_m^{*k}\left(\overline{F}^k(\widetilde{c_m})\right) = \widetilde{c_m}$. The plots illustrate the general pattern that a higher current cost realization $\widetilde{c_m}$ implies a larger range of re-matching costs F for which the firm is willing to re-match its supplier. On the other hand, for the case of impatient (patient) agents, the closer the current cost realization gets to $\underline{c_m^{*1}}$ ($\underline{c_m^{*2}}$) the smaller the re-matching cost range. For all $\widetilde{c_m} < \underline{c_m^{*1}}$ in the impatient agents case and for all $\widetilde{c_m} < \underline{c_m^{*2}}$ in the patient agents case, no re-matching will be observed.

Proposition 5. *A mean-preserving spread $\overline{G}(c_m)$ on the (continuous) distribution function $G(c_m)$ increases the critical search cost levels $\overline{F}^1$ and $\overline{F}^2$, and thereby expands the parameter range where the headquarter engages in re-matching.*

Proof. With the Corollary to Proposition 4 we can write $\overline{F}^k$ as:

$$\overline{F}^k = (1-\delta(1-G(\widetilde{c_m})))\left[\widetilde{\pi_H^N} - \frac{\widetilde{\pi_H^i}}{1-\delta}\right] + \delta\left[\int_{\widetilde{c_m}}^{\overline{c_m}} \pi_H^N g(c_m)\,dc_m + \frac{\int_{\underline{c_m}}^{\widetilde{c_m}} \pi_H^i g(c_m)\,dc_m}{1-\delta}\right] \quad (32)$$

Now consider the distribution $\overline{G}(c_m)$ to be a MPS of $G(c_m)$ and let $\overline{g}(c_m)$ be the associated probability density function. Denote by $\hat{F}^k$, $k = 1, 2$, the re-matching thresholds resulting under the distribution $\overline{G}(c_m)$. In order to proof the Proposition we need to show

$$\hat{F}^k - \overline{F}^k > 0 \quad , \qquad k = 1, 2.$$

This expression can be restated as:

$$(\overline{G}(\widetilde{c_m})-G(\widetilde{c_m}))\left[\widetilde{\pi_H^N}-\frac{\widetilde{\pi_H^i}}{1-\delta}\right]+\int_{\widetilde{c_m}}^{\overline{c_m}}\pi_H^N\overline{g}(c_m)\,dc_m-\int_{\widetilde{c_m}}^{\overline{c_m}}\pi_H^N g(c_m)\,dc_m+\frac{\int_{\underline{c_m}}^{\widetilde{c_m}}\pi_H^i\overline{g}(c_m)\,dc_m-\int_{\underline{c_m}}^{\widetilde{c_m}}\pi_H^i g(c_m)\,dc_m}{1-\delta}>0,\quad i=N,RC$$

Integration by parts and simplifying gives:

$$\int_{\widetilde{c_m}}^{\overline{c_m}}\frac{\partial\pi_H^N}{\partial c_m}G(c_m)\,dc_m-\int_{\widetilde{c_m}}^{\overline{c_m}}\frac{\partial\pi_H^N}{\partial c_m}\overline{G}(c_m)\,dc_m+\frac{\int_{\underline{c_m}}^{\widetilde{c_m}}\frac{\partial\pi_H^i}{\partial c_m}G(c_m)\,dc_m-\int_{\underline{c_m}}^{\widetilde{c_m}}\frac{\partial\pi_H^i}{\partial c_m}\overline{G}(c_m)\,dc_m}{1-\delta}>0,\quad i=N,RC \tag{33}$$

Since $\int_{\widetilde{c_m}}^{\overline{c_m}}\frac{\partial\pi_H^N}{\partial c_m}G(c_m)\,dc_m-\int_{\widetilde{c_m}}^{\overline{c_m}}\frac{\partial\pi_H^N}{\partial c_m}\overline{G}(c_m)\,dc_m<0$, a sufficient condition for (33) to hold is

$$\int_{\underline{c_m}}^{\widetilde{c_m}}\frac{\partial\pi_H^i}{\partial c_m}G(c_m)\,dc_m-\int_{\underline{c_m}}^{\widetilde{c_m}}\frac{\partial\pi_H^i}{\partial c_m}\overline{G}(c_m)\,dc_m\geq\int_{\widetilde{c_m}}^{\overline{c_m}}\frac{\partial\pi_H^N}{\partial c_m}\overline{G}(c_m)\,dc_m-\int_{\widetilde{c_m}}^{\overline{c_m}}\frac{\partial\pi_H^N}{\partial c_m}G(c_m)\,dc_m,\quad i=N,RC.$$

The latter can be rewritten as:

$$\int_{\underline{c_m}}^{\widetilde{c_m}}\left(\frac{\partial\pi_H^i}{\partial c_m}-\frac{\partial\pi_H^N}{\partial c_m}\right)G(c_m)\,dc_m\geq\int_{\underline{c_m}}^{\widetilde{c_m}}\left(\frac{\partial\pi_H^i}{\partial c_m}-\frac{\partial\pi_H^N}{\partial c_m}\right)\overline{G}(c_m)\,dc_m,\quad i=N,RC \tag{34}$$

Since $\frac{\partial\pi_H^i}{\partial c_m}-\frac{\partial\pi_H^N}{\partial c_m}\leq 0$, $i=N,RC$, together with the fact that $\overline{G}(c_m)$ is a MPS of $G(c_m)$, expression (34) is always true. □

D Separating search and re-matching

Finally, in this Appendix we propose an alternative specification for the firm's decision of supplier re-matching. Contrary to the baseline model, the firm can now separately decide whether or not she wants to re-match with a supplier that she encounters during the search process. In particular, we modify the re-matching stage as follows:

7. **Re-matching stage**: H can decide to search for a new supplier. When deciding to search, she incurs a publicly known fixed cost $F>0$. Let c_m^t be the unit cost level of her current supplier, and c_m^{t+1} the unit cost of the new supplier that she has encountered during her search. The cost c_m^{t+1} is randomly drawn from the distribution function $G(c_m)$ which is i.i.d. over periods. If the cost draw is such that $c_m^{t+1}<c_m^t$, the headquarter re-matches and continues the game with the new supplier. If $c_m^{t+1}\geq c_m^t$ she keeps her previous supplier.

Our assumption is thus that the headquarter can observe the efficiency level c_m^{t+1} of the candidate supplier, and will only re-match if he is more efficient than her current partner. Analogously to the baseline model the search decision has two dimensions:

1. if the initial supplier is a high-cost type ($c_m^0=c_m^h$), does the firm start searching?
2. if the initial supplier is a low-cost type ($c_m^0=c_m^l$), or if the firm has found a low-cost type during her search, does she stop searching and stick with that partner?

For the *case of impatient agents* ($\delta<\underline{\delta}$) the analysis remains the same as in the baseline model. That is, if the firm finds (or is initially matched with) a low-cost supplier, she will stick to that partner since further costly search makes no sense. When the firm

is still matched with a high-cost supplier and if search costs are low enough, the firm starts searching. Search stops once she finds a low-cost supplier, but this can require several periods. The only difference to the baseline model is that, during this search period, the firm now has the option to stick with her initial high-cost supplier, rather than switching to another high-cost supplier in every round. However, in both scenarios we have identical (Nash) investments and (Nash) payoffs in the respective stage game round, hence the formal analysis from the baseline model remains unchanged. The search condition $F < \overline{F}^1$ from (13) therefore still applies.

For the *case of patient agents* ($\delta > \underline{\delta}$), the analysis become more intricate. First, consider the second aspect. As before, "cheat-and-run" can be ruled out by ensuring that the search process stops whenever the firm is matched to a low-cost supplier:

$$\frac{1}{1-\delta}\pi_H^{RC,l} \geq E[\pi_H^{cheat-and-run} \mid c_m^0 = c_m^l] \tag{35}$$

Compared to (20), the expected "cheat-and-run"-payoffs have to be slightly modified for the case where search and re-matching are two separate decisions. They can be formalized by the following program:

$$V_0 = \pi_H^{D,l} - F + \delta P V_0 + \delta(1-P)V_1, \qquad V_1 = z - F + \delta P V_0 + \delta(1-P)V_1,$$

where $z = \max\{\pi_H^{N,l}, \pi_H^{D,h}\}$. Note that in the baseline case we had $z = \pi_H^{D,h}$. Now, if the firm encounters a high-cost supplier during her search, depending on the difference in unit costs c_m^h and c_m^l, it may be better to stick to the current low-cost partner and play Nash with him, rather than to re-match and then cheat on the high-cost supplier. As a consequence:

$$E[\pi_H^{cheat-and-run} \mid c_m^0 = c_m^l] = \frac{1}{1-\delta}\left[(1-\delta(1-P))\pi_H^{D,l} + \delta(1-P)z - F\right].$$

Using the equivalent steps from the baseline model, from (35) we can derive a lower threshold on search costs, $\widetilde{F}'$, where $F > \widetilde{F}'$ rules out "cheat-and-run" behaviour.

Turning to the first question, search may thus only occur if the initial supplier is a high-cost type, $c_m^0 = c_m^h$, and as before it will actually occur if search costs F are low enough. In particular, and equivalently to (14), the expected payoff when engaging in search must be higher than the continuation payoff with the initial high-cost supplier, i.e.,

$$E[\pi_H^{search,RC} \mid c_m^0 = c_m^h] > \frac{1}{1-\delta}\pi_H^{RC,h}. \tag{36}$$

If that condition is violated and the firm decides not to search, she forms a RC with her initial high-cost supplier which is sustainable since $\delta > \underline{\delta}$. If condition (36) is satisfied, the firm starts searching. Once she has found a low-cost supplier, search stops and she forms a long-term RC collaboration with that partner as shown before. But in every search round, the firm encounters a high-cost supplier with probability $(1-P)$, so it may take several periods before the once-and-for-all supplier turnover actually takes places.

Separating the firm's decisions of supplier search and re-matching introduces the possibility to further distinguish the patient agents into two groups that behave differently

during the periods of supplier search. In the following we show that for a subset of very patient agents it can be incentive compatible to engage in a RC with the initial high-cost supplier during the search periods, despite the ongoing search for a better partner. Specifically, the RC is better for the firm than Nash play with M_0 if

$$\begin{aligned}\pi_H^{RC,h} &+ \left[(1-P)\pi_H^{RC,h} + P\pi_H^{RC,l}\right] \cdot \left(\delta + \delta^2(1-P) + \delta^3(1-P)^2 + \dots\right) \\ &> \pi_H^{D,h} + \left[(1-P)\pi_H^{N,h} + P\pi_H^{RC,l}\right] \cdot \left(\delta + \delta^2(1-P) + \delta^3(1-P)^2 + \dots\right),\end{aligned}$$

and for M_0 the RC is incentive compatible if

$$\begin{aligned}\pi_M^{RC,h} &+ \pi_M^{RC,h}\left(\delta + \delta^2(1-P) + \delta^3(1-P)^2 + \dots\right) \\ &> \pi_M^{D,h} + (1-P)\pi_M^{N,h}\left(\delta + \delta^2(1-P) + \delta^3(1-P)^2 + \dots\right).\end{aligned}$$

These incentive compatibility constraints thus boil down to

$$\frac{1}{1-\delta(1-P)}\pi_i^{RC,h} > \pi_i^{D,h} + \frac{\delta(1-P)}{1-\delta(1-P)}\pi_i^{N,h} \qquad \text{for} \quad i = H, M,$$

which are similar to (IC-H) and (IC-M) from above, but capture the replacement probability P in every period. A temporary RC with the high-cost initial supplier until replacement is thus optimal if $\delta \geq \frac{1}{1-P}\frac{\pi^{JFB}-\pi_H^D-\pi_M^D}{\pi_H^N+\pi_M^N-\pi_H^D-\pi_M^D} = \underline{\delta}/(1-P) > \underline{\delta}$, where $\underline{\delta}$ is the previously derived critical discount factor given in (10). Put differently, very patient agents with a discount factor above $\underline{\delta}/(1-P)$ would always form a RC, both with the initial high-cost and with the final low-cost supplier. By contrast, mildly patient agents with $\underline{\delta} < \delta < \underline{\delta}/(1-P)$ only form the RC once they have found their final low-cost supplier, but not in the temporary search phase with the high-cost supplier.

Having distinguished the very patient and the mildly patient agents, we can now complete the model extension and derive the critical search cost levels for the very patient agents (the ones for mildly patient agents are the same as in the baseline model and described by $\overline{F}^2$). For the very patient agents who always engage in RCs ($\delta > \underline{\delta}/(1-P)$), the search decision can be formalized as

$$V_0 = \pi_H^{RC,h} - F + \delta V_1, \qquad V_1 = \frac{1}{1-\delta(1-P)}\left((1-P)(\pi_H^{RC,h} - F) + \frac{P}{1-\delta}\pi_H^{RC,l}\right),$$

which yields these expected profits

$$E[\pi_H^{search,RCstrong} \mid c_m^0 = c_m^h] = \frac{1}{1-\delta(1-P)}\left[\pi_H^{RC,h} - F + \frac{\delta P}{1-\delta}\pi_H^{RC,l}\right] \tag{37}$$

Plugging (37) into (36) and rearranging we obtain the critical search cost level $\overline{F}^3$ for the very patient agents case:

$$F < \frac{\delta P}{1-\delta}\left[\pi_H^{RC,l} - \pi_H^{RC,h}\right] \equiv \overline{F}^3. \tag{38}$$

Comparing (13), (16) and (38), it can be verified that $\overline{F}^3 > \overline{F}^2$ always holds. Moreover, as $\overline{F}^1$ and $\overline{F}^2$, also $\overline{F}^3$ is increasing in δ. Finally, we can make a similar consistency

argument as in Appendix A to guarantee $\overline{F}^2 > \widetilde{F}'$ which in turn implies $\overline{F}^3 > \widetilde{F}'$ since $\overline{F}^3 > \overline{F}^2$.

We summarize our results under separated search and re-matching in the following Propositions 6 and 7, which are analogous to Propositions 2 and 3 from the main text and now incorporate the distinction of mildly patient and very patient agents.

Proposition 6. *a) Suppose the headquarter is initially matched with a low-cost supplier ($c_m^0 = c_m^l$). Patient agents (with $\delta > \underline{\delta}$) will collaborate with that supplier forever in a relational contract (RC) agreement with $\{h^*, m^*\}$ and $B^*(\delta)$, assuming that re-matching costs are not too low ($F > \widetilde{F}'$). Impatient agents (with $\delta < \underline{\delta}$) will form a long-term collaboration of repeated Nash bargainings with that supplier.*

b) With $c_m^0 = c_m^h$, impatient agents with $\delta < \underline{\delta}$ search if $\widetilde{F} < F < \overline{F}^1$ and continue searching until they find a low-cost supplier. The RC can never be implemented.

c) With $c_m^0 = c_m^h$, mildly patient agents with $\underline{\delta} < \delta < \underline{\delta}/(1-P)$ search if $\widetilde{F} < F < \overline{F}^2$ and continue searching until they find a low-cost supplier. The RC forms with the final low-cost supplier, but not with the initial high-cost supplier.

d) With $c_m^0 = c_m^h$, very patient agents with $\delta > \underline{\delta}/(1-P)$ search if $\widetilde{F} < F < \overline{F}^3$ and continue searching until they find a low-cost supplier. The RC forms both with the final low-cost supplier, and with the initial high-cost supplier during the search period.

Moreover, we state the following result referring to the impacts of a mean-preserving spread (MPS) in the distribution of supplier costs:

Proposition 7. *A mean-preserving spread in the distribution of supplier costs (a larger difference between c_m^l and c_m^h at constant P) increases the critical search cost levels $\overline{F}^1$, $\overline{F}^2$ and $\overline{F}^3$, and thereby expands the parameter range where the headquarter engages in search.*

We illustrate the changes that result from separating the decisions of search and re-matching in Figure 4, which is comparable to Figure 2. The dotted horizontal line at $\underline{\delta}/(1-P)$ indicates the critical discount factor above which agents are very patient and an RC is formed also during the search for a better partner. We label this a *short-term collaboration* (STC), since it is neither a one-shot nor a truly long-term interaction. In the range between the dotted and the solid horizontal line (at $\underline{\delta}$) agents are mildly patient, and play Nash during the STC-phase and only turn to the RC once the LTC is launched. This is the causal effect $LTC \rightarrow RC$ studied in the main text. Finally, below the solid line agents are impatient and always play Nash.

A MPS on the distribution of supplier costs has the same effect on $\overline{F}^3$ as it has on $\overline{F}^2$ and $\overline{F}^1$ and shifts the $\overline{F}^3$-function outwards. In total, also when separating the search decision from the re-matching decision, a MPS on $G(c_m)$ unambiguously increases the firm's propensity to search. For the mildly patient agents, this indirectly affects contractual structures because they only offer a RC once they have found their ultimate low-cost supplier. Yet, as in the baseline model, the MPS does not directly affect contractual structures since $\underline{\delta}$ and $\underline{\delta}/(1-P)$ are both independent of c_m.

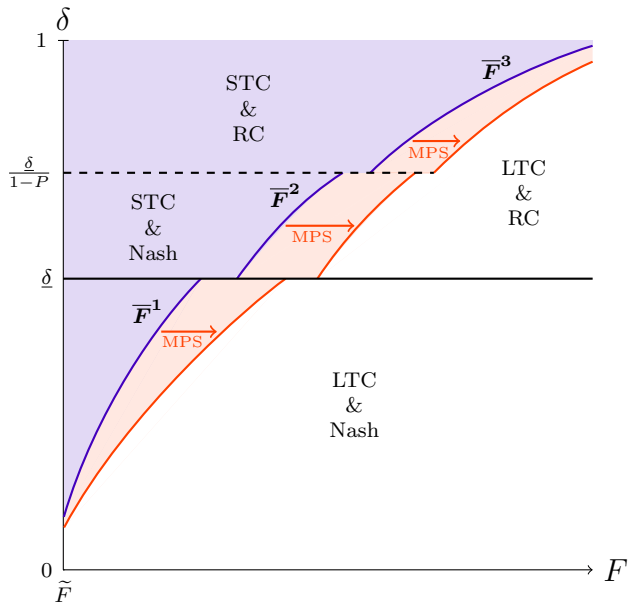

Figure 4: Contractual and search/re-matching decision ($c_m^0 = c_m^h$)