

Essays on Repeated Strategic Interactions and the Organization of Multinational Firms

D I S S E R T A T I O N

zur Erlangung des akademischen Grades
doctor rerum politicarum (Dr. rer. pol.)
im Fach Volkswirtschaftslehre

eingereicht an der Wirtschaftswissenschaftlichen Fakultät
der Heinrich-Heine-Universität Düsseldorf



von:	Christian Fischer, M.Sc. geboren am 5. November 1987 in Nürnberg
Erstgutachter:	Prof. Dr. Jens Südekum
Zweitgutachter:	Prof. Dr. Hans-Theo Normann
Abgabedatum:	30. September 2016

Acknowledgments

I am grateful for all the support that I received during my work on this dissertation. Most of all, I would like to thank my two supervisors Jens Südekum and Hans-Theo Normann for their terrific guidance and motivational support. In joint research projects I could learn a lot from them, most importantly how to structure and communicate my work. Especially, I want to thank Jens for encouraging me to pursue my interests in the academic field of international trade and the organization of multinational firms. Also, I want to express my gratitude to Hans for introducing me to the world of economic laboratory experiments. It was a great pleasure to work with both of them.

This thesis was written during my time as a doctoral researcher at the *Düsseldorf Institute for Competition Economics (DICE)* at Heinrich-Heine-University Düsseldorf. I perceived my time at DICE as a privilege: Not only did I benefit from excellent conditions to be academically productive but also from the strong support that I received to communicate my results to an academic audience through conferences and workshops. This makes the institute undoubtedly an exceptional place to pursue a doctorate in economics.

I also want to thank my co-author Fabrice Defever for fruitful conversations and collaborations, and for a productive research visit at the University of Nottingham that he made possible. Moreover, this thesis benefited greatly from discussions with numerous researchers at DICE and beyond. I am particularly grateful to Florian Knauth, Miriam Thöne, and my office mates Volker Benndorf and Shiva Shekhar.

Finally, I want to thank my parents Egon and Doris for their continued love and unconditional support, and my brother Andreas who was always very helpful in mathematical discussions.

Contents

1	Introduction	1
2	Relational contracts and supplier turnover in the global economy	5
2.1	Introduction	6
2.1.1	Theory	6
2.1.2	Empirical application	8
2.1.3	Related literature	9
2.2	The model	10
2.3	Relational contracts without supplier re-matching	13
2.4	Supplier re-matching	18
2.4.1	Searching an efficient partner	18
2.4.2	Searching for a reliable partner	25
2.5	Empirical application	28
2.5.1	Data and variables	28
2.5.2	Relational contracts and long-term collaborations	30
2.5.3	The impact of cost dispersion	35
2.6	Conclusion	37
2.7	References	45
2.A	Appendix	48
2.B	Supplementary appendix	50
3	Supplier search and re-matching in global sourcing – theory and evidence from China	67
3.1	Introduction	68
3.2	The model	71
3.2.1	Environment	71
3.2.2	Separable search and re-matching	73
3.2.3	Non-separable search and re-matching	76
3.2.4	Summary of theoretical results and main prediction	79
3.3	Empirical analysis	80
3.3.1	Data and variables	80
3.3.2	Descriptive evidence	82

3.3.3	Robustness checks and discussion	85
3.4	Conclusion	86
3.5	References	89
4	Payment contracts and trade finance in export relationships	93
4.1	Introduction	94
4.2	The model	97
4.3	The nature of payment contracts	99
4.3.1	Cash in advance	99
4.3.2	Open account	102
4.3.3	Comparing the payment contract equilibria	104
4.4	Dynamic selection of payment contracts	105
4.4.1	Main result	105
4.4.2	Discussion	107
4.5	Trade credit insurance	109
4.5.1	Insured payment contracts in isolation	110
4.5.2	When to employ trade credit insurance	111
4.5.3	Comparative statics	115
4.6	Conclusion	116
4.7	References	117
4.A	Appendix	119
5	Collusion and bargaining in asymmetric cournot duopoly – an experiment	121
5.1	Introduction	122
5.2	Model and benchmarks	124
5.3	Experimental design	127
5.4	Procedures	129
5.5	Results	129
5.5.1	Overview and behavior across supergames	129
5.5.2	The effect of communication	131
5.5.3	The effect of earned roles	134
5.5.4	Bargaining outcomes	135
5.5.5	How do firms collude in the Talk treatments?	136
5.5.6	Text-mining analysis	139
5.6	Conclusion	140
5.7	References	142
5.A	Appendix	145
6	Conclusion	149

List of Tables

2.1	Duration of processing and non-processing transactions	31
2.2	Duration of processing transactions with fixed-fees and supplied equip- ments	31
2.3	Duration of processing and non-processing transactions – excluding multinational affiliates	32
2.4	Duration of processing transactions with fixed-fees and supplied equip- ments – excluding multinational affiliates	32
2.5	Relational-contracts and long-term collaborations: Transaction-level .	39
2.6	Relational-contracts and long-term collaborations: Product-level . . .	40
2.7	Relational-contracts and long-term collaborations – robustness checks	41
2.8	The impact of cost dispersion	42
2.9	The impact of cost dispersion – robustness checks (I)	43
2.10	The impact of cost dispersion – robustness checks (II)	44
3.1	Firms, products and destinations	81
3.2	Duration of export transactions by ownership	83
3.3	Duration of export transactions by export destination	83
3.4	Relational-contracts and long-term collaborations: Transaction-level .	85
3.5	Relational-contracts and long-term collaborations: Industry-level . . .	87
3.6	Relational-contracts and long-term collaborations: Transaction-level .	88
5.1	Benchmarks	127
5.2	2×2 factorial treatment design and the corresponding number of par- ticipants.	128
5.3	Summary of treatment averages for outputs, profits, and welfare . . .	130
5.4	Ordinary least-squares regressions for output, payoff, and welfare . . .	131
5.5	Probit regression with dummy <i>deviation</i> as the dependent variable . .	138
5.6	Text-mining analysis	141

List of Figures

2.1	Incentive-compatible bonus payments	16
2.2	Contractual and re-matching decision	24
2.3	Contractual and search/re-matching decision ($c_m^0 = c_m^h$)	56
3.1	Illustration of search threshold and MPS	79
4.1	Payment contract determines the stage game timing	98
4.2	Payment contract choice over time	108
4.3	Payment contract choice over time with trade credit insurance	113
4.4	Ceteris paribus analysis, part 1	115
4.5	Ceteris paribus analysis, part 2	116
5.1	Feasible payoffs and bargaining solutions, exact plot for experimental parameters.	126
5.2	Aggregate quantities, Q , across supergames	130
5.3	Kernel density estimates for outputs by treatment and firm type . . .	132
5.4	Kernel density estimates for profits by treatment and firm type . . .	133
5.5	Session means in the π_1 - π_2 space and benchmarks	135
5.6	Distribution of agreement types, conditional on reaching an agree- ment in the first place.	136
5.7	Deviation frequencies from agreements	137
5.8	Schematic summary	148

Chapter 1

Introduction

“The theory of repeated games [...] is concerned with the evolution of fundamental patterns of interaction between people [...]. Its aim is to account for phenomena such as cooperation, altruism, revenge, threats (self-destructive or otherwise), etc. – phenomena which may at first seem irrational – in terms of the usual ‘selfish’ utility-maximizing paradigm of game theory and neoclassical economics.”

— Robert Aumann, 1981, p. 11,
Survey of Repeated Games

Several decades have passed since Robert Aumann and others have laid the foundations for a theory of repeated games. The analysis of long-term relationships is relevant for virtually every area of economics and researchers have ever since pushed the boundaries of this very theory. By today, it provides a multitude of powerful tools making it possible to evaluate economic questions ranging from industrial organization, over international trade and labor economics to monetary policy and behavioral theory.

This dissertation contributes to an applied literature that studies cooperation in repeated games by offering novel models and predictions on firm relationships in a national and transnational context. It reinforces the predictive power of the method by showing that many of the theoretical findings in this thesis can be confirmed by empirical evaluations. Even though the equilibria of repeated games can rationalize many of at first sight seemingly irrational economic phenomena, there may remain justified skepticism that it is able to explain all of the mechanisms relevant in long-termed economic interactions. The experimental essay of this thesis reports findings that stand in direct conflict with the equilibrium predictions of the theory.

Overall, this dissertation is concerned with two distinct applications. On the one side, it analyzes how firms organize their value chains in a globalized world. On the other side, it studies horizontal agreements between businesses in homogeneous product markets. Chapters 2, 3, and 4 investigate how firms can handle their multinational sourcing and export activities when they are challenged to cope with search and contractual frictions. While in these essays continued, long-term firm relationships induce overly positive and desirable economic outcomes, the ensuing Chapter 5 points at how firms may abuse their coordinative capabilities through collusive conduct. It studies experimentally if and how firms harmonize their production when interactions are repeated.

Chapter 2 on *“Relational contracts and supplier turnover in the global economy”* is joint work with Fabrice Defever and Jens Südekum. It starts out with the observation that headquarter firms 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 chapter, 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. Using firm-product-level data of fresh Chinese exporters to the US, we obtain empirical evidence in line with the predictions of our theory.

The third chapter on *“Supplier search and re-matching in global sourcing – theory and evidence from China”* emerged as a spin-off product from Chapter 2 and is also joint work with Fabrice Defever and Jens Südekum. It complements the analysis of Chapter 2 in some important dimensions. Using several simplifications we focus exclusively on the firm’s search and matching problem and generalize the former analysis in several respects. First, rather than just assuming two supplier types, we work with a general density function that describes how the types of potential suppliers are distributed in the economy. Second, and even more importantly, we elicit the role of uncertainty when the firm searches for a new partner, and encounters a candidate supplier in her random search. Finally, we address the empirical relevance of our theoretical predictions using Chinese customs data.

In Chapter 4 entitled *“Payment contracts and trade finance in export relationships”* I examine the use of payment contracts in export partnerships when the importer’s payment moral is uncertain and there is a time gap between the production and sale of traded goods. I find that by deciding upon the payment format the exporter is challenged with a trade-off between maximizing the profits from the current transaction and making use of screening opportunities that differ between format types. Such screening can help to attain information about the reliability of importers which is crucial to develop efficient and intensive export relationships. Trade credit insurance can catalyze export growth and transaction profitability but its positive impact is limited to the initial phase of export relationships.

The fifth chapter on *“Collusion and bargaining in asymmetric cournot duopoly – an experiment”* is joint work with Hans-Theo Normann. In asymmetric Cournot markets without side payments, maximizing joint profits is implausible, firms disagree on the collusive price, and production is necessarily inefficient. We investigate experimentally how firms collude (implicitly and explicitly), if at all, in such situations. When explicit communication is available, joint profits increase above the static Nash level, but nearly all the gains from talking go to the inefficient firm. Bargaining solutions do not satisfactorily predict collusive outcomes, except for the equal split solution. When the role of the efficient firm is earned in a contest, the efficient firm earns higher profits and firms often collude by producing equal amounts.

Confirming previous results, our data also show that, without communication, firms fail to collude and essentially play the static Nash equilibrium.

References

AUMANN, R. (1981): “Survey of Repeated Games,” in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*. Bibliographisches Institut, Mannheim.

Chapter 2

Relational contracts and supplier turnover in the global economy

Co-authored with Jens Südekum and Fabrice Defever

2.1 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* (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 (see, e.g., Egan and Mody 1992) 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 an 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. 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 or if the current partner does not behave cooperatively. We endogenously determine *with whom* and *how* the buyer interacts, and show that the possibility to switch partners crucially affects the contractual nature of buyer-supplier relationships.

2.1.1 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 her current match, she decides whether to engage in a costly search for a better supplier. 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 input investments are contractible. Yet, if agents are patient enough on average, an efficient long-term RC can be implemented in equilibrium that overcomes hold-up and underinvestment problems.

At first we study the firm's search for an *efficient* partner, who can provide the input of a given quality at lower unit costs. In the main text we make several assumptions to keep this analysis as simple as possible. For instance, we consider a grim trigger strategy in that repeated game setting, where a deviation is punished with infinite Nash reversion. Furthermore, we assume that the suppliers' unit costs are perfectly observable to the headquarter, that there are only two types of suppliers (with high and low unit costs, respectively), and that the producer *must* switch partners when she decides to pay the fixed re-matching cost. In a supplementary appendix, we generalize our model along all those lines. Specifically, we there consider an *optimal* punishment strategy as in Abreu (1988), imperfect monitoring as in the seminal collusion model by Green and Porter (1984), a continuous distribution from which the supplier's productivity is drawn (in a related fashion to Melitz 2003), and similarly as in Burdett (1978), we separate the processes of *search* and *re-matching* by assuming that the firm may also stick to her current partner if she encounters a candidate that is less efficient than her current match. In a second step of our analysis, we then allow for heterogeneity in the time discount rate across potential suppliers and study the firm's search for a *reliable* partner. This analysis shows that for sufficiently low search costs the firm seeks to find a patient supplier, who tends to be more cooperative as he cares more about future profits.

Our model delivers a key insight whose essence remains robust in all variations and extensions: long-term collaborations (LTCs) and relational contracts (RCs) sharply need to be distinguished, but the two concepts are correlated. To see the first part, notice that if agents are too 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. Hence, a LTC is not the same as a RC. On the second part, our model clarifies that an efficient RC is only feasible among patient agents, and it would surely emerge in equilibrium if supplier re-matching was ruled out. But once this realistic option is introduced, the firm only starts a RC once it is satisfied with the quality of her current supplier and decides to stop searching for a better partner. In other words, there is a positive correlation since the endogenous decision to launch a LTC can start the emergence of a RC.

This result adds a novel perspective to the literature on the international organization of value chains, and the theory of the firm more generally. The upshot of that literature is that sufficiently patient agents may overcome inefficiencies that bad legal institutions and weak contract enforcement may create. This insight also applies in the context of input sourcing, but our theory goes beyond this point. It explicates that the relevant decision *with whom* to interact is crucially important

for the interrelated question *how* the contractual nature of this interaction will look like.

2.1.2 Empirical application

While we mostly focus on the model, we also aim for an empirical application to address the practical relevance of this theoretical prediction. Putting our model to a test is challenging, however, because several key aspects are unobservable to the researcher. Our theory is about particular matches of domestic buyers and foreign suppliers, but current data only rarely allows observing such matches.¹ 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. Recent work by Macchiavello and Morjaria (2014) has addressed those problems by turning to a detailed case study for one specific industry (coffee beans in Rwanda), where firms are surveyed about their buyer-supplier arrangements.

In this paper, we adopt a different approach. We exploit Chinese customs data that encompass all exports to the US and other destination markets between 2000 and 2006 at the firm-product level. This context is thus fairly broad and well suited for our purpose, because the paradigm of incomplete contracts appears to be quite plausible when it comes to input sourcing from China (Antràs 2015). The cost we have to bear, however, is that we need to rely on coarser empirical proxies whether a particular buyer-supplier match involves a relational agreement, namely by specific categories of processing trade arrangements (Feenstra and Hanson 2005) that resemble the features of RCs fairly closely.

In line with our theory, we robustly find a positive correlation between LTCs and RCs in the data. More specifically, the longer a fresh Chinese exporter sells a particular product to the US, the more likely this is done under those processing arrangements with strong relational elements. Furthermore, we also test another prediction of our model, namely that the degree of heterogeneity across potential suppliers negatively affects LTCs and RCs, because it increases the firm's incentive to search and therefore postpones the formation of an efficient long-term collaboration. We empirically confirm this prediction as we find that a more dispersed distribution of suppliers' unit values is negatively associated with observed match durations, and with the share of RCs at the product level.

¹A notable exception is the paper by Eaton *et al.* (2014) who are able to construct pairs of Colombian exporters and their US importers. 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, 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).

2.1.3 Related literature

A first strand of related literature incorporates elements of search and matching into models of international trade, either referring to 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 relational agreements.

A second strand incorporates contractual choices into models of export dynamics, see Araujo, Mion and Ornelas (2016) and Aeberhardt, Buono and Fadinger (2014). In those frameworks, a domestic exporter matches with a foreign sales agent who might behave opportunistically, and once the firm has found a reliable partner, she increases exports at the intensive margin. This is consistent with empirical evidence by Rauch and Watson (2003) and Besedes (2008), who find that firms initially place small test orders when they first deal with a new supplier, and 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. Examples includes Macchiavello and Morjaria (2014), Board (2011), or Corts and Singh (2004). The models by Kukharskyy (2016), Kukharskyy and Pflüger (2016), and Kamal and Tang (2015) are particularly closely related, as they are also based on Antràs (2003). However, none of them studies the interaction between RCs and supplier 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, this paper is related to a small literature on repeated games with heterogeneous time discount factors (Harrington 1989; Lehrer and Pauzner 1999; Haag and Lagunoff 2007). Our model applies some techniques of this literature to the context of global sourcing, which seems especially relevant since economic environments are usually more volatile in less developed countries, so that suppliers from those economies may be less patient than large domestic headquarter corporations. Moreover, from a conceptual point of view, our model seems to be the first to consider this heterogeneity in a setting where agents not only choose their strategy *how* to interact, but also *with whom* they play.

The rest of this paper is organized as follows. Section 2.2 present the basic model. In Section 2.3, we analyze RCs when supplier re-matching is still ruled

out. We introduce this option in Section 2.4. Section 2.5 presents our empirical application. Section 2.6 concludes.

2.2 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:

$$y = A p^{-\frac{1}{1-\alpha}}, \quad 0 < \alpha < 1, A > 0, \quad (2.1)$$

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:

$$y = \left(\frac{h}{\eta}\right)^\eta \left(\frac{m}{1-\eta}\right)^{1-\eta}, \quad 0 < \eta < 1 \quad (2.2)$$

The firm's headquarter (called H) is located in the domestic country and provides the input h herself. 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 foreign supplier (called M).² 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. Moreover, $\tau > 1$ units of the input m need to be shipped for one unit to arrive in the domestic country. Finally, following Antràs (2003) and Antràs and Helpman (2004), we assume that 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:

$$R = p \cdot y = A^{1-\alpha} \left[\left(\frac{h}{\eta}\right)^\eta \left(\frac{m}{1-\eta}\right)^{1-\eta} \right]^\alpha \quad (2.3)$$

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.

²In this paper we do not analyze the integration versus outsourcing decision central to the Antràs (2003)-model, but we assume that M is an independent supplier with full ownership rights over his assets. Moreover, we assume that there is just one single intermediate input, and thus one supplier to contract with. See Schwarz and Suedekum (2014) for a (static) model of global sourcing with multiple inputs and multilateral bargaining.

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 $\delta_H \in (0, 1)$. 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 face the same trade cost τ , have the same outside option ω_M , and a common discount factor $\delta_M \in (0, 1)$.³ 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. **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.

³Notice that δ_M need not coincide with the headquarter's δ_H . See the discussion below. Moreover, in Section 2.4.2 we also consider the case where suppliers are heterogeneous in their discount factors δ_{Mi} .

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 Abreu (1988) we can 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.⁴ 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

$$\max_h \beta R(h, m) - c_h h, \quad \max_m (1 - \beta) R(h, m) - \tau c_m m.$$

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):

$$\tilde{h} = \frac{\alpha \beta \eta}{c_h} \tilde{R}, \quad \tilde{m} = \frac{\alpha(1 - \beta)(1 - \eta)}{\tau c_m} \tilde{R}, \quad (2.4)$$

where

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

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

$$\pi_H^N = \beta \tilde{R} - c_h \tilde{h}, \quad \pi_M^N = (1 - \beta) \tilde{R} - \tau c_m \tilde{m}, \quad (2.5)$$

⁴In the following, we use the terms joint profit-maximizing and first-best investment levels synonymously.

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.

Before proceeding, it is worth stressing that the assumption of full relationship-specificity may be restrictive, as intermediate inputs in reality often have some value on an outside market. To capture this, let us assume that the supplier could also sell m at price $p_m > \tau c_m$ elsewhere. This, in turn, improves his bargaining power vis-a-vis H . In particular, similar as in Baker et al. (2002), H has to compensate M for this alternative use of the input, and bargain over the surplus value of the relationship. In stage 3, the two agents then face the following problems which can be solved in a similar fashion as before:

$$\begin{aligned} \max_h \quad & -p_m m + \beta R(h, m) - c_h h, \\ \max_m \quad & p_m m + (1 - \beta) (R(h, m) - p_m m) - \tau c_m m = (1 - \beta) R(h, m) - (\tau c_m - \beta p_m) m, \end{aligned}$$

where we assume that $\tau c_m > \beta p_m$. It is easy to see that the resulting investments \tilde{m} and \tilde{h} and the payoff π_M^N are increasing, while π_H^N is decreasing in p_m .⁵ Intuitively, an improvement in the outside market conditions incentivizes the supplier to produce, as he is fully compensated for the opportunity costs. The headquarter, in turn, also contributes more to the relationship, but ends up with a lower Nash payoff due to the required transfer.

To stay consistent with the baseline model from the literature, we set $p_m = 0$ in the remainder of this paper so that (2.4) and (2.5) are the solutions to the static stage game.

2.3 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

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

and the corresponding input quantities and revenue are given by

$$h^* = \frac{\alpha \eta}{c_h} R^* > \tilde{h}, \quad m^* = \frac{\alpha(1 - \eta)}{\tau c_m} R^* > \tilde{m}, \quad R^* = \mathcal{A} (c_h^\eta (\tau c_m)^{1-\eta})^{\frac{-\alpha}{1-\alpha}} > \tilde{R}. \quad (2.6)$$

⁵An increase of p_m has, indeed, analogous effects as a reduction of the supplier's unit costs c_m .

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

$$\pi_H^{RC} = R^* - c_h h^* - B, \quad \pi_M^{RC} = B - \tau c_m m^*. \quad (2.7)$$

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 (2.4), and earn payoffs π_i^N in every period as given in (2.5).⁶ 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

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

This leads to

$$h^D = \frac{\alpha \eta \beta}{c_h} R(h^D, m^*), \quad m^D = \frac{\alpha(1 - \eta)(1 - \beta)}{\tau c_m} R(h^*, m^D), \quad (2.8)$$

with $R(h^D, m^*) = \mathcal{A} \beta^{\frac{\alpha \eta}{1 - \alpha \eta}} (c_h^\eta (\tau c_m)^{1 - \eta})^{\frac{-\alpha}{1 - \alpha}}$ and $R(h^*, m^D) = \mathcal{A} (1 - \beta)^{\frac{\alpha(1 - \eta)}{1 - \alpha + \alpha \eta}} (c_h^\eta (\tau c_m)^{1 - \eta})^{\frac{-\alpha}{1 - \alpha}}$. The deviation payoffs are thus given by

$$\pi_H^D = \beta R(h^D, m^*) - c_h h^D, \quad \pi_M^D = (1 - \beta) R(h^*, m^D) - \tau c_m m^D. \quad (2.9)$$

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:

$$B - \tau c_m m^* \geq \omega_M \quad (\text{PC-M})$$

This participation constraint (PC) requires that the bonus must at least compensate the supplier's full production costs (including trade costs) plus his outside option ω_M .

⁶Nash 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). In the supplementary appendix we consider different penal codes as robustness checks, see the discussion below.

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,

$$\frac{1}{1 - \delta_H} \pi_H^{RC} \geq \pi_H^D + \frac{\delta_H}{1 - \delta_H} \pi_H^N, \quad (\text{IC-H})$$

$$\frac{1}{1 - \delta_M} \pi_M^{RC} \geq \pi_M^D + \frac{\delta_M}{1 - \delta_M} \pi_M^N \quad (\text{IC-M})$$

where π_i^{RC} , π_i^D and π_i^N for $i = H, M$ are given by (2.5), (2.7), and (2.9). Rearranging (IC-H) and (IC-M) for B , we can identify the range of bonus payments for which relational contracting is incentive compatible:

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

$$B \geq \tau c_m m^* + [\delta_M \pi_M^N + (1 - \delta_M) \pi_M^D] \equiv B_M(\delta) \quad (\text{IC-M}')$$

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.

In Figure 2.1 we illustrate this bonus B_i for which the respective player is indifferent between sticking to, and deviating from the RC for the special case where the headquarter and the supplier have a common discount factor ($\delta_H = \delta_M = \delta$). 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 $B_M(\delta)$ 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. This critical discount factor $\underline{\delta} \in (0, 1)$ lies at the intersection of the two linear curves, and is characterized formally in Proposition 1. The grey area in Figure 2.1 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 $B^*(\delta)$ which is decreasing in δ . Thus, when agents are more patient, it is cheaper for the firm to sustain the cooperation with supplier M_0 and allows her to keep a larger relational rent.




Figure 2.1: Incentive-compatible bonus payments

This logic carries over to the realistic case where discount factors are heterogeneous, i.e., where the foreign suppliers are less patient than the domestic headquarter corporation ($\delta_H > \delta_M$). As shown in the following Proposition 1, the RC can still be implemented in equilibrium when agents are patient enough *on average*:

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

$$\tilde{\delta}(\delta_H, \delta_M, \alpha, \beta, \eta) \equiv \frac{\sum_i \delta_i (\pi_i^D - \pi_i^N)}{\sum_i (\pi_i^D - \pi_i^N)} \geq \frac{\sum_i \pi_i^D - \pi^{JFB}}{\sum_i (\pi_i^D - \pi_i^N)} \equiv \underline{\delta}(\alpha, \beta, \eta), \quad (2.10)$$

with bonus payment $B^*(\delta_M) = \tau c_m m^* + [\delta_M \pi_M^N + (1 - \delta_M) \pi_M^D]$.

Proof. See Appendix 2.A.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 ($i = M$). 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 $\pi_M^{RC} = B^* - \tau c_m m^* = [\delta_M \pi_M^N + (1 - \delta_M) \pi_M^D] > \pi_M^N$. The intuition is the strategic value of the deviation option which is capitalized in the optimal bonus payment $B^*(\delta_M)$. We could also introduce an explicit notion of relational rent sharing, by assuming that H offers M a bonus payment $B_M(\delta_M) < B^* < B_H(\delta_H)$. But our model provides a rationale why informal RCs are to the mutual benefit of both parties even without this rent sharing.

Second, in spirit of the literature on repeated games with heterogeneous time preferences (see Harrington 1989; Haag and Lagunoff 2007), Proposition 1 shows that sustaining the RC in equilibrium requires a sufficiently high *average* patience level, but not necessarily that *both* agents are patient. More specifically, in our model, the weighted average of δ_H and δ_M (that we denote by $\tilde{\delta}$) must exceed some critical level $\underline{\delta}$, where the weights are given by the relative deviation incentives $(\pi_i^D - \pi_i^N)$ of the firm and the supplier, respectively. If condition (2.10) holds, the (patient) firm can ensure the cooperation of the (impatient) supplier M_0 by an appropriate bonus payment $B^*(\delta_M)$, which is smaller the higher δ_M is.⁷ Also notice that (2.10) boils down to $\tilde{\delta} = \delta \geq \underline{\delta}$ if both agents have the same discount factor, in which case the RC can be implemented if δ is high enough.

Third, it is evident from (2.10) that the demand parameter A , as well as the unit costs c_h and c_m and the trade costs τ do not affect the critical discount factor $\underline{\delta}$, which solely depends on the demand elasticity parameter α , the headquarter intensity η , and the bargaining power β . Furthermore, the parameters A , c_h , c_m , and τ also cancel from the weighting terms $(\pi_i^D - \pi_i^N)$ on the LHS of (2.10), which depends only on δ_H , δ_M , α , β and η (see Appendix 2.A.1). 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 , or if he becomes more efficient and reduces c_m^0 over time (e.g., through learning effects), this would proportionally raise all payoffs π_i^D , π_i^{RC} , π_i^N for $i = H, M$, and the bonus $B^*(\delta_M)$, but it would *not* affect the threshold $\underline{\delta}$. In Figure 2.1, for example, the lines would still intersect at the same level of δ . Intuitively, cost changes do not affect payoffs relatively stronger for one contractual arrangement than for another.⁸

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 the supplementary appendix we first 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 that is decreasing in the punishment length T , but that is still independent of c_m and τ . Moreover, as a

⁷Lehrer and Pauzner (1999) show that the set of feasible repeated game payoffs can be larger than the convex hull of the underlying stage game payoffs when both players i are very patient in absolute terms, but still inhibit differing values of δ_i . Intuitively, this works through temporary side payments from the very patient to the slightly less patient agent. In this paper, we sidestep those issues by restricting our attention to stationary bonus payments that remain constant over time, and by ignoring mixed strategy equilibria. Agents in our model can, thus, at most coordinate on π^{JFB} via the RC but not on joint payoffs beyond.

⁸This property does not crucially hinge on the assumed functional forms in (2.1) and (2.2) with iso-elastic demand and Cobb-Douglas production, but 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 these parameters.

further robustness check, we have secondly developed 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 collusion model by Green and Porter (1984). Also in that case we can implicitly determine a critical discount factor that now depends on demand uncertainty, but still not on the production or the trade costs. The cost-orthogonality of the critical discount factor therefore seems to hold under fairly general conditions, and does not hinge on the peculiarities of the model parametrization or the punishment strategies.

2.4 Supplier re-matching

We now introduce the option for the headquarter to switch suppliers after every round of output realization. In our main analysis in Section 2.4.1, all potential suppliers differ only in their unit costs c_m but are identical otherwise. This captures the firm's search for an *efficient* partner who is able to provide the desired input with a given quality at low costs. In Section 2.4.2, we then relax the assumption that all suppliers have common time preferences and analyze the firm's search for a *reliable* partner who is more patient.

2.4.1 Searching an efficient partner

We add the following stage to the game, with all other stages remaining unchanged:

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 $\{c_m^l, c_m^h\}$, 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 .⁹ Second, we assume that the headquarter *must* re-match when she

⁹We 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. Furthermore, in an interesting recent model of assortative buyer-supplier matching under complete contracts, Cajal-Grossi (2016) argues that search costs may also depend on the heterogeneity of potential suppliers. Our model is simpler in this respect, as it focuses on a single firm, and for simplicity assumes that search costs F are fixed and independent of $g(c_m)$.

decides to engage in costly search, even when the supplier she encounters is less efficient than her current partner.

In the supplementary 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 (2.10) is independent of c_m . Since δ_M is identical for all (potential) M, this together with the given δ_H fixes the weighted average patience level $\tilde{\delta}$ which also does not depend on c_m . The headquarter therefore 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.

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

With $\tilde{\delta} < \underline{\delta}$, RCs can never be implemented regardless of the supplier's cost type. Agents would always deviate from the prescribed first-best behaviour, since they are too impatient on average, and thus Nash bargaining and the associated under-investment 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 average 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, 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 round, so it is 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 firm 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 (2.5). The following condition must be satisfied in order for re-matching to occur:

$$E[\pi_H^{RM,N} \mid c_m^0 = c_m^h] > \frac{1}{1 - \delta_H} \pi_H^{N,h}, \quad (2.11)$$

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,

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

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

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

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

$$E[\pi_H^{RM,N} \mid c_m^0 = c_m^h] = \frac{1}{1 - \delta_H(1 - P)} \left[\pi_H^{N,h} - F + \frac{\delta_H P}{1 - \delta_H} \pi_H^{N,l} \right]. \quad (2.12)$$

Combining expressions (2.11) and (2.12), we then obtain the following re-matching condition,

$$F < \frac{\delta_H P}{1 - \delta_H} \left[\pi_H^{N,l} - \pi_H^{N,h} \right] \equiv \bar{F}^1, \quad (2.13)$$

which states that the cost F must be low enough in order for re-matching to occur. Notice that the critical search cost level \bar{F}^1 depends positively on δ_H , i.e., re-matching is more attractive the more patient the headquarter is. The intuition is that future profits then matter more, and hence it becomes more important to find an efficient low-cost supplier. The patience level δ_M , by contrast, does not directly affect the firm's search decision, although we have to keep in mind that $\tilde{\delta} < \underline{\delta}$ is always assumed to hold. Moreover, \bar{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 $\tilde{\delta} < \underline{\delta}$, we can conclude that no supplier turnover will ever arise if search costs are too high ($F \geq \bar{F}^1$) 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 $F < \bar{F}^1$, 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 on average to sustain cooperation.

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

Among patient agents RCs are feasible in principle, and would emerge for sure 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, even in a setting where agents are patient. 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.¹⁰ 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 2.A.2 we derive

¹⁰It 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.

this lower bound \tilde{F} and show that it is increasing in P . The condition $F > \tilde{F}$ 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,

$$E[\pi_H^{RM,RC} \mid c_m^0 = c_m^h] > \frac{1}{1 - \delta_H} \pi_H^{RC,h}. \quad (2.14)$$

If that condition is violated, the headquarter forms a RC with her initial high-cost supplier which is sustainable since $\tilde{\delta} > \underline{\delta}$. If condition (2.14) 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 every round so that RCs are not credible.¹¹ 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:

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

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

$$E[\pi_H^{RM,RC} \mid c_m^0 = c_m^h] = \frac{1}{1 - \delta_H(1 - P)} \left[\pi_H^{N,h} - F + \frac{\delta_H P}{1 - \delta_H} \pi_H^{RC,l} \right] \quad (2.15)$$

Plugging (2.15) into (2.14), we then obtain the critical search cost level \bar{F}^2 for this case,

$$F < - \left(\pi_H^{RC,h} - \pi_H^{N,h} \right) + \frac{\delta_H P}{1 - \delta_H} \left[\pi_H^{RC,l} - \pi_H^{RC,h} \right] \equiv \bar{F}^2, \quad (2.16)$$

which also depends positively on δ_H , 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, \bar{F}^2 depends

¹¹This 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 supplementary 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.

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 \bar{F}^2 must be larger than the lower bound \tilde{F} derived before. As also shown in Appendix 2.A.2, this ranking $\tilde{F} < \bar{F}^2$ can be guaranteed by making appropriate restrictions on $g(c_m)$, namely that c_m^h and c_m^l are not too similar.

The impact of cost dispersion

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 ($c_m^0 = c_m^l$). When the agents are sufficiently patient on average (if $\tilde{\delta} \geq \underline{\delta}$), the firm will collaborate with that supplier forever in a relational contract (RC) agreement with $\{h^*, m^*\}$ and $B^*(\delta_M)$, assuming that re-matching costs are not too low ($F > \tilde{F}$). When the agents are not patient enough on average (if $\tilde{\delta} < \underline{\delta}$) the firm establishes a long-term collaboration of repeated Nash bargainings with that supplier.*

b) With $c_m^0 = c_m^h$ and $\tilde{\delta} < \underline{\delta}$, the firm starts re-matching if $F < \bar{F}^1$ and continues until she finds a low-cost supplier. The RC can never be implemented.

c) With $c_m^0 = c_m^h$ and $\tilde{\delta} \geq \underline{\delta}$, the firm starts re-matching if $F < \bar{F}^2$ and continues until she finds 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 2.2 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 weighted average patience level $\tilde{\delta}$ is fully determined by the given parameter values δ_H , δ_M , α , β and η , and is depicted on the vertical axis. 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 too impatient on average and thus engage in a long-term collaboration (LTC) with repeated Nash bargainings. Above the solid line, agents are patient enough on average and the RC forms immediately.

Turning to the left panel, we there depict the critical search cost levels \bar{F}^i (with $i = 1, 2$) derived in (2.13) and (2.16), below which impatient (patient) agents start re-matching when their initial supplier is a high-cost type. Recall that both are increasing and convex in δ_H , which for a given value of δ_M implies the concave

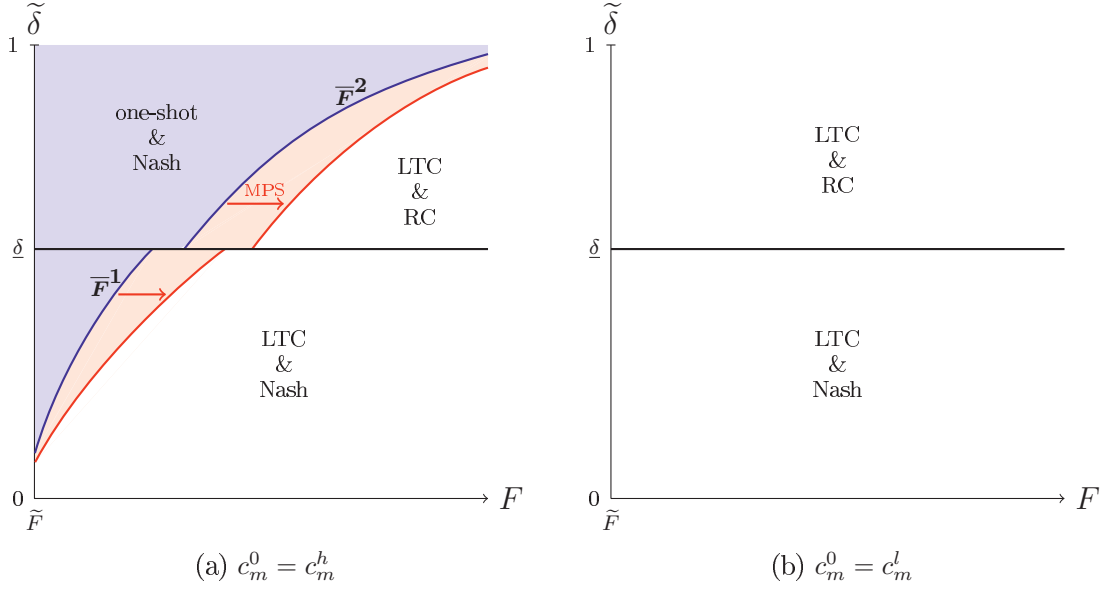


Figure 2.2: Contractual and re-matching decision

relationship between $\tilde{\delta}$ and \bar{F}^i as shown in Figure 2.2. For the case of 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 firm has succeeded in finding a low-cost supplier, the constellation switches to the right panel of Figure 2.2. Now consider the patient agents case (above the horizontal line) in the left panel. When search costs are high, the firm sticks with her initial high-cost supplier forever and immediately establishes a RC. When search costs are low, however, the firm seeks for a better partner and re-matches until a low-cost type is found. During the transition, when the firm 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, the RC is established once the firm decides to stop re-matching, that is, the launching of a LTC with the current supplier *causes* the start of the RC.¹²

¹²In our baseline model, a LTC is thus a necessary but not a sufficient condition for a RC, as all RCs are also LTCs but not vice versa. When search and re-matching are separated (see supplementary appendix) a LTC is not even necessary, as RCs can then form on a short-term basis during the search for a better partner if agents are very patient. Yet, also in that case, there is a positive causal relationship, as the start of a LTC causes the RC for mildly patient agents, while impatient (very patient) agents never (always) have a RC in operation.

In Proposition 3, we summarize how an increase in the suppliers' cost heterogeneity affects re-matching and contractual choices for the firm.

Proposition 3. *A mean-preserving spread (MPS) in the distribution $g(c_m)$ (a larger difference between c_m^l and c_m^h at constant P) increases \bar{F}^1 and \bar{F}^2 , and thus expands the parameter range where the headquarter engages in re-matching. For the patient agents case ($\tilde{\delta} > \underline{\delta}$), this implies that the formation of RCs is postponed when the firm starts re-matching after the parameter change.*

Proof. These results follow directly from (2.13) and (2.16), because the terms $(\pi_H^{RC,l} - \pi_H^{RC,h})$ and $(\pi_H^{N,l} - \pi_H^{N,h})$ are increasing when the MPS is performed, since π_H^{RC} and π_H^N are both monotonically decreasing in c_m .

In the left panel of Figure 2.2, we illustrate this Proposition by the shifts of \bar{F}^1 and \bar{F}^2 to the right. More precisely, when the firm's $(\tilde{\delta}, F)$ -combination falls inside the orange range in Figure 2.2, it has not searched before the parameter change, but after the MPS it will engage in re-matching so that supplier turnover is observed during a transitory period. The intuition is an increase in the relative benefit of collaborating with a low-cost supplier, which makes searching for such a type more attractive at unchanged values of the search costs F , average patience $\tilde{\delta}$, and the other parameters. At the same time, notice that the higher cost dispersion neither affects the critical discount factor $\underline{\delta}$ nor the weighted average patience $\tilde{\delta}$. Hence, the MPS does not *directly* affect whether the firm is able to establish a RC in the long-run. This contractual choice is only affected *indirectly* and *temporarily*, because more search and supplier turnover occurs. In the patient agents case, this induces the firm not to offer RCs immediately, but only when it has found the right match. Put differently, higher cost dispersion does not alter the firm's long-run decision whether to establish a RC, but it may postpone the formation.

Finally, notice that an increase of the demand level A or a decrease in trade costs τ also increase \bar{F}^1 and \bar{F}^2 , because they also raise the relative benefit of collaborating with a low-cost supplier. Those parameter changes in A or τ have therefore analogous effects to the MPS, i.e., they also lead to more search and fewer RCs during a transitory period.

2.4.2 Searching for a reliable partner

In the previous subsection, we have formalized the firm's search for a partner with low unit costs. Yet, as emphasized by Egan and Mody (1992), headquarter corporations in reality do not only care about the quality of the provided inputs and the *efficiency* of their suppliers, but also about their *reliability* when it comes to the agreed upon terms of delivery and other contractual details. We capture this notion of reliability here in a particular way, by formalizing the firm's search for a partner

with high δ_M . The idea is that more patient suppliers are, *ceteris paribus*, more willing to sustain the cooperation with the firm, by reliably adhering to what has been stipulated in the relational agreement.

For this analysis, we assume that potential suppliers differ in their time preference rate but are identical in all other respects, in particular with respect to their unit costs c_m which are now the same for all suppliers. The re-matching stage is adjusted as follows:

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

As before, we focus in the main text on a simple parameterization of $g(\delta_M)$ with only two realizations $\{\delta_M^{pat}, \delta_M^{imp}\}$, while deferring the general analysis to the supplementary appendix. The population of potential suppliers here consists of only two types, *patient* and *impatient* ones with δ_M^{pat} and δ_M^{imp} , respectively, where $\delta_M^{pat} > \delta_M^{imp}$ and where P_δ is the share of patient suppliers in this population. Moreover, we restrict the parameter values of δ_M^{pat} , δ_M^{imp} and δ_H such that $\tilde{\delta}(\delta_H, \delta_M^{pat}) > \underline{\delta}$ in (2.10), while we allow $\tilde{\delta}(\delta_H, \delta_M^{imp})$ to be above or below $\underline{\delta}$. That is, if supplier search was ruled out as in Section 2.3, the RC would for sure be implemented with the patient type, but we consider both cases for the impatient type.¹³ Finally, we again assume that the firm *must* re-match after every round where she has decided to search for a new supplier.

First, realize that the firm has no incentive to engage in re-matching when the initial supplier M_0 is already a patient type, but she would start the RC with him immediately. The reason is that the firm cannot improve her payoff, since there is no better (more patient) supplier to be found, so it is always better avoiding the search cost F . Re-matching can only make sense if the initial supplier is an impatient type. When $\tilde{\delta}(\delta_H, \delta_M^{imp}) < \underline{\delta}$, the RC cannot be launched with M_0 and the agents would always play Nash, hence the firm has an incentive to search in order to establish a beneficial relational agreement with someone else. When $\tilde{\delta}(\delta_H, \delta_M^{imp}) > \underline{\delta}$, the firm could already start the RC with the initial (impatient) partner. Still there is an incentive to search, because both patient and impatient suppliers provide the same quality input $m^*(c_m)$ and generate the same revenue $R^*(c_m)$ in the RC, but the firm can reduce the required bonus from $B^*(\delta_M^{imp})$ to $B^*(\delta_M^{pat})$ when finding a patient supplier, thus raising her own payoff.

¹³This rules out cases where the weighted average patience $\tilde{\delta}$ is so low that the RC could never be implemented, regardless of the initial supplier's type. That case strikes us as being less interesting. Recall that unit costs and trade costs are irrelevant for those case distinctions.

Formally, when the firm engages in re-matching, it is clear that agents play Nash in the respective period since the supplier anticipates his replacement. The RC only starts once a patient supplier is found, which happens with probability P_δ in every round of search. The firm, thus, solves the following problem:

$$V_0 = \pi_H^N - F + \delta_H V_1, \quad V_1 = \frac{1}{1 - \delta_H(1 - P_\delta)} \left((1 - P_\delta)(\pi_H^N - F) + \frac{P_\delta}{1 - \delta_H} \pi_H^{RC,pat} \right),$$

where $\pi_H^{RC,k}$ is the firm's payoff within the RC when collaborating with a supplier of patience level $k = \{imp, pat\}$, where $\pi_H^{RC,pat} > \pi_H^{RC,imp}$.¹⁴ Solving for V_0 implies these expected profits when starting to re-match in search for a more reliable partner:

$$E[\pi_H^{RM} \mid \delta_M^0 = \delta_M^{imp}] = \frac{1}{1 - \delta_H(1 - P_\delta)} \left[\pi_H^N - F + \frac{\delta_H P_\delta}{1 - \delta_H} \pi_H^{RC,pat} \right] \quad (2.17)$$

For the firm to start re-matching, this term in (2.17) must exceed the continuation payoff with the initial impatient supplier, which is $\pi_H^N/(1 - \delta_H)$ in the first case with $\tilde{\delta}(\delta_H, \delta_M^{imp}) < \underline{\delta}$, and $\pi_H^{RC,imp}/(1 - \delta_H)$ in the second case with $\tilde{\delta}(\delta_H, \delta_M^{imp}) > \underline{\delta}$. This, in turn, allows us to derive critical search cost level below which the firm would actually start re-matching when her initial supplier is an impatient type. Those read as

$$\bar{F}^3 = \frac{P_\delta \delta_H}{1 - \delta_H} (\pi_H^{RC,pat} - \pi_H^N) \quad \text{and} \quad \bar{F}^4 = -(\pi_H^{RC,imp} - \pi_H^N) + \frac{P_\delta \delta_H}{1 - \delta_H} (\pi_H^{RC,pat} - \pi_H^{RC,imp})$$

in the first and in the second case, respectively.¹⁵

Summing up, when the firm searches for a *reliable* partner, the basic logic from the previous analysis – the search for an *efficient* supplier – carries over: The firm only re-matches when the initial supplier is not already a “good” (i.e., patient or low-cost) type, and if supplier search is not too costly. Moreover, the RC does not start before the firm is ultimately satisfied with her match and decides to stop searching once and for all. In fact, also in this model version, the launching of a LTC causes the start of the RC, and in the second case considered above we find that search per se alters the contractual nature of buyer-supplier relationships. In particular, the firm would immediately start the RC with the initial impatient supplier if $F > \bar{F}^4$, but it will play Nash and replace him when search is not so costly (if $F < \bar{F}^4$) and only start the RC later.¹⁶

¹⁴Note that π_H^N would be the same for both supplier types, but is relevant here only for impatient ones.

¹⁵As before, we introduce a lower bound for search costs \tilde{F} in order to rule out “cheat-and-run”-behaviour. See the supplementary appendix.

¹⁶One subtle difference between search for an efficient and a reliable partner should be noted. For the latter we find that, whenever the firm decides to start re-matching, any future LTC will necessarily be a RC, while repeated Nash interactions can only occur if the firm never searches. This is different in Section 2.4.1, because impatient agents may start re-matching and eventually establish a LTC without ever being able to start a RC.

2.5 Empirical application

Our model implies a positive correlation between the prevalence of relational contracts (RC), and whether a buyer-supplier match is a long-term collaboration (LTC). This follows directly from Proposition 2 above: Not every LTC is a RC, but there is a positive connection since the decision to stop searching can cause the start of a relational agreement. In this final section, we empirically explore this key predictions of our theoretical framework. Moreover, we also test the prediction of Proposition 3 that more cost dispersion across the potential suppliers reduces both RCs and LTCs, essentially because the firm has a higher incentive to search for a better partner.

In Section 2.5.1 we first describe how we construct proxies for the two main variables, LTCs and RCs. In Section 2.5.2 we explore their correlation while controlling for many other firm- and product-level characteristics. Finally, in Section 2.5.3 we evaluate how they are affected by cost dispersion.

2.5.1 Data and variables

Exports of fresh Chinese suppliers. Our empirical evaluation exploits customs data from the Chinese General Administration of Customs for the period 2000 to 2006. For each year, the data allow us to identify exports of Chinese firms disaggregated by HS6-product and destination. For simplicity, we call these firm-product-destination combinations an export *transaction*. 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 exclude all product categories j with fewer than 10 active firms, and those where control variables are not available. The final sample

¹⁷As 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.

then includes 16,150 fresh Chinese exporters, starting a total of 63,580 new export transactions, and spans 1,004 different HS6-digit product categories.

Match durations. Following the fresh Chinese exporters 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 product 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 product j , which we use as our main explanatory variable: the share of LTCs.

Relational contracts. Measuring the dependent variable – the share of RCs – is more challenging, because the various explicit and implicit match-specific contractual arrangements between the American buyer and the Chinese supplier are, almost by definition, not observable to the researcher. To still make some progress, we exploit the Chinese custom data which provides some further information about the type of arrangement between the partners. In particular, as our measure for the prevalence of RCs, we use the transactions realized using processing trade arrangements – a data feature that has been previously utilized in somewhat different, but related contexts.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

¹⁸Ideally, 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%).

¹⁹See Feenstra and Hanson (2005), Fernandes and Tang (2012), Kee and Tang (2016), Manova and Yu (2016).

featured in our model. Fixed-fee agreements are also associated with another interesting feature: the headquarter provides the intermediate inputs to its supplier. As a consequence, the supplier does not have to advance the money for buying these inputs. Using evidence from Rwanda’s Coffee Mills, this feature has previously been identified by Marchiavello and Morjaria (2014) as a major indicator for relational contracts (in their paper, the coffee farmer receive inputs from the coffee mill). If the relational contract is destroyed, the supplier has to pay the entire costs of inputs upfront. These funds must be borrowed, and a credit constrained Chinese supplier might not be able to obtain the necessary liquidity. Moreover, another related feature of these processing trade 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 buyer-supplier relationship.

Summing up, in our empirical analysis we consider two specific proxies for the prevalence of RCs. The first and somewhat wider measure is the overall share of all transactions realized using processing trade arrangements in a product j . Then, as a second and narrower measure, we only consider processing arrangements with fixed-fee 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 product. Both proxies are, arguably, only imperfect measures. Still, we believe that they capture the essence of RCs fairly well in at least some dimensions, so that our approach is a useful first step to gauge the empirical relevance of our theoretical results in the broad, but highly relevant context of US input sourcing from China.

2.5.2 Relational contracts and long-term collaborations

Descriptive overview. Table 2.1 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 last for 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, in which case the foreign party also provides the intermediate inputs, or one with supplied equipment by the US partner; around 13% of them actually have both features. Table 2.2 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 14–18 % of those transactions last for less than a year, while 55–60 % seem to be resilient long-

term collaborations. The share of LTCs is, thus, rather similar to the one among all processing trade transactions.

Table 2.1: Duration of processing and non-processing transactions

Duration of transactions	With Processing		Without Processing	
	# 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 2.2: Duration of processing transactions with fixed-fees and supplied equipments

Duration of transactions	Processing with fixed-fees		Processing with supplied equipments	
	# 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 2.3 and 2.4 are analogous to Tables 2.1 and 2.2, but exclude all transactions from Chinese foreign affiliates which are dependent subsidiaries or part of a joint venture. 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

²⁰This 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, 2016). 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.

Table 2.3: Duration of processing and non-processing transactions – excluding multinational affiliates

Duration of transactions	With Processing		Without Processing	
	# 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 2.4: Duration of processing transactions with fixed-fees and supplied equipments – excluding multinational affiliates

Duration of transactions	Processing with fixed-fees		Processing with supplied equipments	
	# 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

trade transactions drops dramatically by 70.3% (from 7,773 to 2,308), although this drop is less pronounced for the two sub-types. 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 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 no impact on the duration of processing transactions with fixed fees and/or with supplied equipments.

More importantly, dropping the foreign affiliates 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 2.4: 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 yields a positive correlation between match durations (LTCs) 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.

Correlations at the transaction level. Next, we present conditional correlations on the link between RCs and LTCs at the individual transaction-level. We use all new transactions by the fresh Chinese exporter i in product j , and construct a dummy variable which equals one if the transaction was carried out under a processing trade arrangement. This dummy we then regress on another 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.²¹

The results are reported in Table 2.5. Column 1 indicates that, when an export transaction is a continuing long-term collaboration (LTC), it has a 8.9 percentage points higher probability of being a RC. In column 2, we introduce a set of firm and HS6 digit product fixed effects to the regression. The rationale is that the use of processing trade may potentially vary across products within a firm, as argued by Manova and Yu (2016). Yet, while the magnitude of the coefficient associated with LTCs decreases substantially in column 2, we still find a strong and significant correlation between LTCs and RCs even after taking into account any potential firm- or product-specific effects. Finally, in columns 3 and 4, we use the narrower proxy for RCs and the same picture emerges: If the relationship is a LTC, it is more likely a RC.²²

Correlations at the product level. In Table 2.6, 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 product j on the share of long-term transactions (lasting 3 years or more), and in columns 3-4 we use the respective share of processing transactions with fixed-fee and/or supplied equipment. In all regressions, we add the share of export transactions realized by a foreign affiliate as an additional control variable. Moreover, in column 2 and 4, we introduce a set of further product-level control variables, namely the capital-labor ratio, the human capital intensity, and the R&D-sales ratio, which are standard measures

²¹Robust standard-errors are clustered at the HS6-digit level. In order to obtain transaction-level estimates that are directly comparable to the product-level estimations, we weight each transaction such that all products j receive the same overall weight in the regression.

²²Observe that the coefficient associated with foreign ownership is much smaller. In fact, foreign ownership seems to be a good predictor for the use of processing arrangements, but mainly when the processing transactions do not make use of fixed-fees and supplied equipment.

for the headquarter-intensity η .²³ Moreover, we include the HS6-digit version of the relationship-specificity index by Nunn (2007), and add demand elasticities from Broda and Weinstein (2006) to capture α .²⁴

Table 2.6 shows that there is a robust positive correlation between the share of LTCs and RCs also at the product level. The coefficients range between 5 and 9 percentage points, which is consistent with our previous results from Table 2.5. Adding the product-level control variables in columns 2 and 4 leaves the main coefficient basically unchanged.

Robustness checks. In Table 2.7 we perform a battery of robustness checks of our product-level regressions. First, in Columns 1 and 6, we exclude all the export transactions of foreign affiliates before aggregating the data. This leads to a decreasing number of observations as we still exclude products with less than 10 firms. The main coefficient for LTCs on RCs is almost identical to the one obtained in Table 2.6, however, and remains significant at the 1% level. Overall, this indicates that our results remain robust also in a smaller sample only consisting of independent Chinese-owned manufacturers.

Our second robustness check excludes all transactions which are not part of the firm's main sector of activity. Since many firms export multiple products, we define this main sector as the 2-digit HS product for which we observe the largest number of exported HS6 products. In columns 2 and 7, we only keep the transactions that belong to this main 2-digit sector of each firm, and again we obtain coefficients close to our baseline estimates.

The third 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 3 and 8 show that the main coefficient for the LTC share increases when those cases are excluded, but the qualitative results remain the same.

Forth, we introduce a set of HS2-digit fixed effects to exploit the intra-sectoral variation with each 2 digit sector. The inclusion of these sectoral fixed effects gen-

²³Those variables have been constructed by Nunn and Treffer (2013) using information for the US, hence, we implicitly assume that they are correlated with the factor-intensity of production in China. The information to construct the capital-labor ratio and the human capital intensity are from the U.S. Census Bureau's 2005 Annual Survey of Manufactures, while the data used to calculate the R&D ratio are from the Orbis database, constructed by Bureau van Dijk Electronic Publishing. In both cases, the data are available at the 6-digit NAICS industry level and reclassified into the HS6-digit product classification. Similarly, skill intensity is the ratio of non-production worker wages to total worker wages. R&D intensity is the R&D expenditures divided by firm sales in each product.

²⁴As data are disaggregated at the HS 10 digit level, we use the median value for each HS6 product.

erate results of the same magnitude to the one obtained in our baseline estimation, and the coefficient of interest remains significant at the 1% level (see columns 4 and 9).

Finally, we evaluate the impact of clustering the standard errors at different level of aggregation. In particular, we now cluster the standard errors at the more aggregate HS2-digit level. The coefficients in column 5 and 10 are, thus, identical as in columns 2 and 4 of Table 2.6 but the standard errors are larger. Still, the coefficient in column 5 and 10 of Table 2.7 remains significant at the 5% level.²⁵

Summing up, all results from Tables 2.1–2.7 empirically support our main theoretical prediction from Proposition 2: There is a positive correlation between RCs and LTCs in the Chinese data, and this conclusion is very similar both for the wider and for the narrower proxy of relational contracting.

2.5.3 The impact of cost dispersion

In this last section, we evaluate the impact of cost dispersion on LTCs and RCs, and thereby empirically explore Proposition 3 which predicts a negative effect in both cases.

To construct an empirical proxy for the cost dispersion across all Chinese exporters of a product j , we approximate the transaction-specific marginal cost of the respective supplier by the unit-value of each transaction. The cost dispersion measure within product j is then given by the standard deviation of the log of unit-values within the respective HS6 category 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

²⁵We have also experiment clustering the standard errors at yet another level of aggregation using the corresponding 6-digit NAICS code associated with each HS-6 digit. In fact, the information used by Nunn and Treffer (2013) to construct the factor intensity of production are collected at the 6-digit NAICS level. Using a correspondence table, the authors link the NAICS 6-digit classification with HS6-digit product classification. This level of clustering also leaves the coefficient associated with the LTC Share significant at least at the 5% level. To save in space, we did not report these results.

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.²⁶

The results in Table 2.8 show that cost dispersion is, indeed, negatively correlated with the share of RCs²⁷ (column 1 and 2) as well as with the share of LTCs (column 3 and 4). This is the case for the baseline specifications in columns 1 and 3, and also when the other product-level control variables are included (columns 2 and 4). Quantitatively, the results suggest that, an increase of the standard deviation of the log of unit values by 1 percent is associated with a sizable decrease of RC arrangements by around 2 percentage points, and LTCs even decrease by more than 4 percentage points.

Robustness checks. In Table 2.9 we first conduct the same robustness checks as in Table 2.7 above and find that our main coefficients remain very stable and highly statistically significant in all specifications. Table 2.10 then considers additional robustness checks, in particular with respect to the construction of our cost dispersion measure, and with respect to the destination market for the Chinese exports.

First, in columns 1 and 6 of Table 2.7 we exclude all transactions with processing arrangements when constructing our dispersion variable. Hence, it 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 measure captures differences in pricing rules associated with processing arrangements, because the duration of the transactions and the type of arrangement could both impact on the level of investments and, ultimately, on the unit values. Yet, columns 1 and 6 show that all results remain unchanged.

Second, we exclude all short-term transactions (less than 3 years) when constructing the dispersion, as those may also be priced differently which in turn could affect the unit values. Yet again, columns 2 and 7 yield very similar results as before.

Third, in columns 3 and 8, we use the unit values of export transactions to Japan instead of the US to construct the cost dispersion across Chinese exporters. 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

²⁶Manova and Zhang (2012) argue that differences in unit-values could also reflect quality variations. In this case, our empirical proxy would be capturing heterogeneity in suppliers' quality. However, under iso-elastic demand and monopolistic competition, a marginal cost or a quality difference 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.

²⁷To save space, we only present results for the narrower proxy for RCs, namely processing agreement with fixed-fees and/or supplied equipments. Very similar results are obtained when using the wider RC proxy.

environment, which makes the two countries comparable in this sense. It turns out that the results remain very similar also when using this alternative proxy for product-level wide cost dispersion.

In our fourth robustness check, we use the export transactions to Japan to build our dependent variables (RC and LTC share), while using the transactions in the US market to build the dispersion measure. The results in column 4 and 9 are again very similar and the negative correlations remain highly significant. Our findings thus do not seem to be specific to US-China trade, but also apply to Chinese exports to other destinations.

In our last exercise we experiment with the use of firms' sales data to proxy for cost dispersion, similar as in Helpman et al (2004) and Yeaple (2006). More specifically, Nunn and Trefler (2013) provide information on the standard deviation of Chinese firms' sales at the HS6-digit level based on plant sales from the Orbis database, and we consider these data as our alternative proxy. Column 5 and 10 of Table 10 present the correlations of this sales dispersion with the share of RCs and LTCs, respectively. The coefficients remain negative, and in the former case significant at the 10% level. In the latter case the estimate becomes insignificant, however, which likely reflects the fact that this information is only available for fewer HS6 products so that the coefficient is estimated with less precision.

Overall, we obtain empirical results in Tables 2.8–2.10 that are firmly in line with Proposition 3 of our theoretical model. In particular, we find that higher cost dispersion across all Chinese firms exporting a particular product is negatively correlated with the share of LTCs, and with the share of RCs in that product category. In the light of our model, we may put forward one explanation for this pattern in the data: US headquarter corporations have an incentive to search harder when there is more heterogeneity across potential suppliers in China, and this higher search propensity postpones the formation of relational contracts as US firms have not yet found their ideal match.

2.6 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 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. Furthermore, the data also confirm our theoretical prediction that more cost dispersion is negatively associated with match durations and relational contracts.

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.

Table 2.5: Relational-contracts and long-term collaborations: Transaction-level

Dependent variable:	Processing transactions	Processing with Fixed-fees and/or supplied equipments
LTC dummy	0.089*** (0.006)	0.031*** (0.003)
Foreign ownership	0.286*** (0.009)	0.018*** (0.004)
Firm fixed-effects	No	No
HS6 product fixed-effects	No	No
Observations	63,580	63,580
R^2	0.162	0.007

Notes: Dependent variable: In column 1-2, the dummy variable 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. Columns 2 and 4 include firm and HS6 digit fixed effects. Robust standard-errors in brackets using clustered standard errors at the HS6 level.

Table 2.6: Relational-contracts and long-term collaborations: Product-level

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

Table 2.7: Relational-contracts and long-term collaborations – robustness checks

Dependent variable:	Share of processing transactions					Share with fixed-fees and/or supplied equipments				
	Drop foreign	Main sector	Drop zero processing	HS2 fixed-effects	HS2 clusters	Drop foreign	Main sector	Drop zero processing	HS2 fixed-effects	HS2 clusters
LTC Share	0.083*** (0.022)	0.089*** (0.030)	0.088*** (0.029)	0.104*** (0.026)	0.083** (0.033)	0.054*** (0.017)	0.053*** (0.016)	0.065** (0.028)	0.054*** (0.015)	0.051** (0.023)
Foreign ownership		0.511*** (0.032)	0.471*** (0.037)	0.426*** (0.034)	0.492*** (0.054)		0.068*** (0.019)	0.068** (0.033)	0.051** (0.020)	0.085*** (0.020)
ln Capital/Worker	-0.020*** (0.005)	-0.027*** (0.008)	-0.014* (0.008)	0.026** (0.012)	-0.020* (0.012)	-0.017*** (0.005)	-0.015*** (0.005)	-0.001 (0.008)	0.018** (0.008)	-0.009 (0.007)
ln Skilled/Worker	-0.069*** (0.019)	-0.121*** (0.025)	-0.081*** (0.021)	-0.095*** (0.026)	-0.081** (0.034)	-0.062*** (0.016)	-0.053*** (0.012)	-0.054*** (0.017)	-0.043*** (0.016)	-0.044** (0.018)
ln R&D Sales	-0.004 (0.004)	-0.001 (0.005)	-0.000 (0.004)	-0.008* (0.005)	-0.003 (0.006)	-0.004 (0.003)	-0.011*** (0.003)	-0.010*** (0.003)	-0.004 (0.003)	-0.011* (0.006)
Nunn measure	0.135*** (0.033)	0.225*** (0.052)	0.233*** (0.049)	0.110** (0.048)	0.212*** (0.058)	0.123*** (0.025)	0.118*** (0.023)	0.157*** (0.039)	0.073*** (0.027)	0.132*** (0.046)
ln demand elasticity	0.017*** (0.006)	0.015* (0.009)	0.024*** (0.009)	0.009* (0.005)	0.014** (0.007)	0.011** (0.005)	0.009** (0.005)	0.018*** (0.007)	0.001 (0.004)	0.009** (0.004)
Observations	809	765	773	1,004	1,004	809	765	509	1,004	1,004
R ²	0.092	0.386	0.326	0.529	0.363	0.094	0.118	0.095	0.364	0.131

Notes: Dependent variable: In column 1-5, the dependent variable is the share of exports transactions using processing agreement on the overall number of transactions within a HS6 product. In column 6-10, the dependent variable is the share of exports transactions using processing agreement with fixed-fees and/or supplied equipments. Robust standard-errors in brackets.

Table 2.8: The impact of cost dispersion

Dependent variable:	Share with fixed-fees and/or supplied equipments		Share of long-term transactions	
Unit value SD	-0.020*** (0.004)	-0.024*** (0.005)	-0.045*** (0.010)	-0.060*** (0.011)
Foreign ownership	0.113*** (0.017)	0.109*** (0.016)	0.234*** (0.034)	0.231*** (0.034)
ln Capital/Worker		-0.008** (0.004)		0.003 (0.010)
ln Skilled/Worker		-0.059*** (0.011)		-0.032 (0.022)
ln R&D Sales		-0.010*** (0.002)		0.008* (0.004)
Nunn measure		0.137*** (0.020)		0.044 (0.058)
ln demand elasticity		0.006 (0.004)		-0.006 (0.008)
Observations	1,004	1,004	1,004	1,004
R^2	0.054	0.142	0.058	0.070

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

Table 2.9: The impact of cost dispersion – robustness checks (I)

Dependent variable:	Share with fixed-fees					Share of long-term collaborations				
	Drop foreign	Main sector	Drop zero processing	HS2 fixed-effects	HS2 clusters	Drop foreign	Main sector	Drop zero processing	HS2 fixed-effects	HS2 clusters
Unit value SD	-0.018*** (0.006)	-0.024*** (0.006)	-0.030*** (0.007)	-0.014*** (0.005)	-0.024*** (0.009)	-0.047*** (0.013)	-0.055*** (0.013)	-0.059*** (0.014)	-0.093*** (0.012)	-0.060*** (0.015)
Foreign ownership		0.080*** (0.019)	0.110*** (0.029)	0.070*** (0.020)	0.109*** (0.019)		0.104*** (0.034)	0.350*** (0.046)	0.257*** (0.037)	0.231*** (0.062)
ln Capital/Worker	-0.017*** (0.005)	-0.013*** (0.005)	-0.000 (0.009)	0.018** (0.008)	-0.008 (0.006)	-0.009 (0.010)	-0.015 (0.012)	-0.016 (0.014)	-0.015 (0.013)	0.003 (0.022)
ln Skilled/Worker	-0.073*** (0.017)	-0.069*** (0.014)	-0.073*** (0.019)	-0.045*** (0.017)	-0.059*** (0.015)	0.012 (0.025)	-0.014 (0.027)	-0.028 (0.026)	0.015 (0.029)	-0.032 (0.038)
ln R&D Sales	-0.003 (0.003)	-0.009*** (0.003)	-0.010*** (0.003)	-0.004 (0.003)	-0.010* (0.005)	0.004 (0.005)	0.024*** (0.006)	0.003 (0.005)	-0.003 (0.007)	0.008 (0.007)
Nunn measure	0.127*** (0.027)	0.117*** (0.025)	0.133*** (0.039)	0.067** (0.027)	0.137*** (0.046)	0.008 (0.063)	-0.035 (0.073)	-0.023 (0.098)	-0.073 (0.072)	0.044 (0.079)
ln demand elasticity	0.008* (0.005)	0.006 (0.005)	0.014* (0.007)	0.000 (0.004)	0.006* (0.003)	-0.010 (0.008)	-0.009 (0.010)	-0.002 (0.011)	-0.005 (0.008)	-0.006 (0.009)
Observations	809	765	509	1,004	1,004	809	765	509	1,004	1,004
R ²	0.093	0.123	0.108	0.360	0.142	0.025	0.057	0.127	0.303	0.070

Notes: Dependent variable: In column 1-5, the dependent variable is the share of exports transactions using processing agreement with fixed-fees and/or supplied equipments on the overall number of transactions within a HS6 product. In column 6-10, the dependent variable is the share of long-term collaborations. Robust standard-errors in brackets.

Table 2.10: The impact of cost dispersion – robustness checks (II)

Construction of unit value SD	No processing	Only LTC	Japan unit value SD	Japanese exports	Chinese sales	No processing	Only LTC	Japan unit value SD	Japanese exports	Chinese sales
Dependent variable:	Share with fixed-fees and/or supplied equipments					Share of long-term collaborations				
Unit value SD	-0.021*** (0.004)	-0.025*** (0.005)	-0.025*** (0.005)	-0.063*** (0.009)	-0.013* (0.007)	-0.062*** (0.010)	-0.050*** (0.011)	-0.051*** (0.010)	-0.078*** (0.012)	-0.006 (0.015)
Foreign ownership	0.109*** (0.016)	0.108*** (0.016)	0.109*** (0.016)	0.342*** (0.028)	0.102*** (0.019)	0.237*** (0.034)	0.221*** (0.035)	0.225*** (0.033)	0.188*** (0.028)	0.164*** (0.040)
ln Capital/Worker	-0.008*** (0.004)	-0.008*** (0.004)	-0.009*** (0.004)	-0.019*** (0.006)	-0.018*** (0.004)	0.004 (0.010)	0.002 (0.010)	0.001 (0.010)	0.041*** (0.011)	-0.021* (0.012)
ln Skilled/Worker	-0.058*** (0.011)	-0.060*** (0.011)	-0.057*** (0.011)	-0.128*** (0.021)	-0.043*** (0.010)	-0.035 (0.022)	-0.024 (0.022)	-0.020 (0.022)	-0.117*** (0.029)	0.003 (0.024)
ln R&D Sales	-0.010*** (0.002)	-0.010*** (0.002)	-0.009*** (0.002)	-0.027*** (0.003)	-0.013*** (0.002)	0.008* (0.004)	0.008* (0.004)	0.009*** (0.004)	0.008 (0.005)	0.008 (0.005)
Numm measure	0.136*** (0.020)	0.136*** (0.021)	0.141*** (0.020)	0.246*** (0.032)	0.130*** (0.022)	0.043 (0.058)	0.041 (0.060)	0.051 (0.059)	-0.124** (0.060)	0.020 (0.072)
ln elasticity	0.006* (0.004)	0.006 (0.004)	0.005 (0.004)	0.024*** (0.007)	0.010** (0.004)	-0.006 (0.008)	-0.005 (0.008)	-0.006 (0.007)	-0.007 (0.009)	0.001 (0.009)
Observations	1,004	1,004	1,004	860	740	1,004	1,004	1,004	860	740
R ²	0.137	0.143	0.150	0.401	0.183	0.073	0.060	0.068	0.117	0.039

Notes: Dependent variable: In column 1-5, the dependent variable is the share of exports transactions using processing agreement with fixed-fees and/or supplied equipments on the overall number of transactions within a HS6 product. In column 6-10, the dependent variable is the share of long-term collaborations. Robust standard-errors in brackets.

2.7 References

- ABREU, D. (1988): “On the Theory of Infinitely Repeated Games with Discounting,” *Econometrica*, 56(2), 383–396.
- AEBERHARDT, R., I. BUONO, AND H. FADINGER (2014): “Learning, Incomplete Contracts and Export Dynamics: Theory and Evidence from French Firms,” *European Economic Review*, 68, 219–249.
- ANTRÀS, P. (2003): “Firms, Contracts, and Trade Structure,” *The Quarterly Journal of Economics*, 118(4).
- ANTRÀS, P. (2015): *Global Production: Firms, Contracts and Trade Structure*.
- ANTRÀS, P., AND A. COSTINOT (2011): “Intermediated Trade,” *The Quarterly Journal of Economics*, 126, 1319–1374.
- ANTRÀS, P., AND E. HELPMAN (2004): “Global Sourcing,” *Journal of Political Economy*, 112(3), 552–580.
- ARAUJO, L., G. MION, AND E. ORNELAS (2016): “Institutions and Export Dynamics,” *Journal of International Economics*, 98, 2–20.
- BAKER, G., R. GIBBONS, AND K. J. MURPHY (2002): “Relational Contracts and the Theory of the Firm,” *The Quarterly Journal of Economics*, 117(1), 39–84.
- BERNARD, A. B., A. MOXNES, AND K. H. ULLTVEIT-MOE (2014): “Two-Sided Heterogeneity and Trade,” Working Paper 20136, National Bureau of Economic Research.
- BESEDEŠ, T. (2008): “A Search Cost Perspective on Formation and Duration of Trade,” *Review of International Economics*, 16(5), 835–849.
- BOARD, S. (2011): “Relational Contracts and the Value of Loyalty,” *American Economic Review*, 101(7), 3349–67.
- BRODA, C., AND D. E. WEINSTEIN (2006): “Globalization and the Gains from Variety,” *The Quarterly Journal of Economics*, 121(2), 541–585.
- BURDETT, K. (1978): “A Theory of Employee Job Search and Quit Rates,” *The American Economic Review*, 68(1), 212–220.
- CAJAL-GROSSI, J. (2016): “Searching for Trade Patterns in Developing Countries: Testing Firms in the ‘Fast Fashion’ Industry,” Mimeo, University of Warwick.
- CARBALLO, J., G. OTTAVIANO, AND C. VOLPE-MARTINCUS (2013): “The Buyer Margins of Firms’ Exports,” Discussion Paper 1234, CEP.

- CORTS, K. S., AND J. SINGH (2004): “The Effect of Repeated Interaction on Contract Choice: Evidence from Offshore Drilling,” *Journal of Law, Economics, and Organization*, 20(1), 230–260.
- EATON, J., M. ESLAVA, D. JINKINS, C. J. KRIZAN, AND J. TYBOUT (2014): “A Search and Learning Model of Export Dynamics,” Working Paper.
- EGAN, M. L., AND A. MODY (1992): “Buyer-Seller Links in Export Development,” *World Development*, 20(3), 321–334.
- FEENSTRA, R. C., AND G. H. HANSON (2005): “Ownership and Control in Outsourcing to China: Estimating the Property-Rights Theory of the Firm,” *The Quarterly Journal of Economics*, 120(2), 729–761.
- FERNANDES, A. P., AND H. TANG (2012): “Determinants of Vertical Integration in Export Processing: Theory and Evidence from China,” *Journal of Development Economics*, 99(2), 396 – 414.
- GREEN, E. J., AND R. H. PORTER (1984): “Noncooperative Collusion under Imperfect Price Information,” *Econometrica*, 52(1), 87–100.
- GROSSMAN, S. J., AND O. D. HART (1986): “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *The Journal of Political Economy*, 94(4), 691–719.
- HAAG, M., AND R. LAGUNOFF (2007): “On the Size and Structure of Group Cooperation,” *Journal of Economic Theory*, 135, 68–89.
- HARRINGTON, J. E. (1989): “Collusion Among Asymmetric Firms: The Case of Different Discount Factors,” *International Journal of Industrial Organization*, 7(2), 289–307.
- HART, O. D., AND J. MOORE (1990): “Property Rights and the Nature of the Firm,” *Journal of Political Economy*, 98(6), 1119–58.
- HELPMAN, E., M. J. MELITZ, AND S. R. YEAPLE (2004): “Export versus FDI with Heterogeneous Firms,” *The American Economic Review*, 94(1), 300–316.
- JOHNSON, R. C., AND G. NOGUERA (2012): “Accounting for Intermediates: Production Sharing and Trade in Value Added,” *Journal of International Economics*, 86(2), 224–236.
- KAMAL, F., AND H. TANG (2015): “Relational Trade: Evidence from the United States,” Mimeo.
- KEE, H. L., AND H. TANG (2016): “Domestic Value Added in Exports: Theory and Firm Evidence from China,” *American Economic Review*, 106(6), 1402–1436.

- KUKHARSKYY, B. (2016): “Relational Contracts and Global Sourcing,” *Journal of International Economics*, 101, 123 – 147.
- KUKHARSKYY, B., AND M. PFLÜGER (2016): “Time is On My Side: Relational Contracts and Aggregate Welfare,” Mimeo.
- LEHRER, E., AND A. PAUZNER (1999): “Repeated Games with Differential Time Preferences,” *Econometrica*, 67(2), 393–412.
- MACCHIAVELLO, R., AND A. MORJARIA (2014): “Competition and Relational Contracts: Evidence from Rwanda’s Coffee Mills,” Working Paper.
- MANOVA, K., AND Z. YU (2016): “How Firms Export: Processing vs. Ordinary Trade with Financial Frictions,” *Journal of International Economics*, 100, 120–137.
- MANOVA, K., AND Z. ZHANG (2012): “Export Prices Across Firms and Destinations,” *The Quarterly Journal of Economics*, 127, 379–436.
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71(6), 1695–1725.
- MELITZ, M. J., AND S. J. REDDING (2014): “Heterogeneous Firms and Trade,” *Handbook of International Economics*, 4, 1–54.
- MONARCH, R. (2015): “It’s Not You, It’s Me: Breakups in U.S.-China Trade Relationships,” Working paper.
- NUNN, N. (2007): “Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade,” *The Quarterly Journal of Economics*, 122(2), 569–600.
- NUNN, N., AND D. TREFLER (2013): “Incomplete Contracts and the Boundaries of the Multinational Firm,” *Journal of Economic Behavior & Organization*, 94(1), 330–344.
- RAUCH, J. E. (1999): “Networks versus Markets in International Trade,” *Journal of International Economics*, 48(1), 7–35.
- RAUCH, J. E., AND J. WATSON (2003): “Starting Small in an Unfamiliar Environment,” *International Journal of Industrial Organization*, 21(7), 1021–1042.
- SCHWARZ, C., AND J. SUEDEKUM (2014): “Global Sourcing of Complex Production Processes,” *Journal of International Economics*, 93, 123 – 139.
- YEAPLE, S. (2006): “Offshoring, Foreign Direct Investment, and the Structure of U.S. Trade,” *Journal of the European Economic Association*, 4(2-3), 602–611.

2.A Appendix

2.A.1 Proof of Proposition 1

We prove this Proposition in two steps. First, consider the special case where the headquarter and the supplier have a common discount factor $\delta_H = \delta_M = \delta$. In that case, we may rephrase Proposition 1 to state that the RC inducing first best investment levels (h^*, m^*) can be supported as a SPNE of the repeated game for all

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

with a corresponding optimal bonus payment $B^*(\delta) = m^*c_m + [\delta\pi_M^N + (1-\delta)\pi_M^D]$. To show this, use (IC-H') and (IC-M') to solve $B_H(\underline{\delta}) = B_M(\underline{\delta})$ 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 (2.7), and comparing it to (2.5) and (2.9), 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 thus $\underline{\delta} \in (0, 1)$. Plugging in the expressions for π_i^N , π_i^{RC} , and π_i^D from (2.5), (2.7), and (2.9) yields the explicit expression for the critical discount factor $\underline{\delta}$ that solely depends on α , β and η , but not on A , c_h , c_m or τ :

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

Second, turning now to the case with heterogeneous discount factors ($\delta_H \neq \delta_M$), we can rewrite the agents' incentive constraints as

$$\frac{1}{1 - \delta_M} [B - m^*c_m] \geq \pi_M^D + \frac{\delta_M}{1 - \delta_M} \pi_M^N, \quad (\text{IC-M})$$

$$\frac{1}{1 - \delta_H} [R^* - h^*c_h - B] \geq \pi_H^D + \frac{\delta_H}{1 - \delta_H} \pi_H^N. \quad (\text{IC-H})$$

The firm will set the bonus B such that (IC-M) binds with equality. Solving it for B yields the optimal bonus payment, $B^*(\delta_M) = m^*c_m + [\delta_M\pi_M^N + (1 - \delta_M)\pi_M^D]$. Inserting this into (IC-H) and rearranging gives

$$\delta_M(\pi_M^D - \pi_M^N) + \delta_H(\pi_H^D - \pi_H^N) \geq \pi_M^D + \pi_H^D - \pi^{JFB},$$

and dividing this expression by $\sum_i (\pi_i^D - \pi_i^N)$ we obtain condition (2.10) as stated above:

$$\tilde{\delta} \equiv \frac{\sum_i \delta_i (\pi_i^D - \pi_i^N)}{\sum_i (\pi_i^D - \pi_i^N)} \geq \frac{\sum_i \pi_i^D - \pi^{JFB}}{\sum_i (\pi_i^D - \pi_i^N)} = \underline{\delta}(\alpha, \beta, \eta).$$

The RHS of this inequality is the same as for the case of homogeneous discount factors, i.e., the same critical level $\underline{\delta}$ applies as before. The LHS is a weighted

average of δ_H and δ_M , with weights given by $(\pi_i^D - \pi_i^N)$. Plugging in the terms from (2.5) and (2.9) we can rewrite the LHS as

$$\tilde{\delta}(\delta_H, \delta_M, \alpha, \beta, \eta) = \frac{\delta_H \cdot \beta(1 - \alpha\eta) \left[\beta^{\frac{\alpha\eta}{1-\alpha\eta}} - \hat{\beta} \right] + \delta_M \cdot (1 - \beta)(1 - \alpha(1 - \eta)) \left[(1 - \beta)^{\frac{\alpha(1-\eta)}{1-\alpha(1-\eta)}} - \hat{\beta} \right]}{\beta(1 - \alpha\eta) \left[\beta^{\frac{\alpha\eta}{1-\alpha\eta}} - \hat{\beta} \right] + (1 - \beta)(1 - \alpha(1 - \eta)) \left[(1 - \beta)^{\frac{\alpha(1-\eta)}{1-\alpha(1-\eta)}} - \hat{\beta} \right]}$$

where $\hat{\beta} \equiv (\beta^\eta(1 - \beta)^{(1-\eta)})^{\frac{\alpha}{1-\alpha}}$. Notice that $\tilde{\delta}$ also does not depend on A , c_h , c_m and τ , and that $\tilde{\delta} = \delta$ if $\delta_H = \delta_M = \delta$. \square

2.A.2 Preventing “cheat-and-run”

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

$$\frac{1}{1 - \delta_H} \pi_H^{RC,l} \geq E[\pi_H^{CAR} \mid c_m^0 = c_m^l] \quad (2.18)$$

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:

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

and solving for V_0 gives

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

With this, we can rewrite condition (2.18) as follows:

$$F \geq (1 - \delta_H(1 - P)) \pi_H^{D,l} + \delta_H(1 - P) \pi_H^{D,h} - \pi_H^{RC,l}. \quad (2.19)$$

Since the RHS of (2.19) is increasing in P , “cheat-and-run” can be precluded for all P by requiring the following sufficient condition:

$$F > \pi_H^{D,l} - \pi_H^{RC,l} \equiv \tilde{F}. \quad (2.20)$$

For consistency it is necessary to show that the lower bound \tilde{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 $\tilde{F} < \bar{F}^2(\tilde{\delta})$ can hold for every relevant patience level δ_H such that $\tilde{\delta} \in (\underline{\delta}, 1)$. In the following we derive a sufficient condition for $\tilde{F} < \bar{F}^2(\underline{\delta})$, which then automatically implies $\tilde{F} < \bar{F}^2(\tilde{\delta})$ since $\bar{F}^2(\tilde{\delta})$ is increasing in $\tilde{\delta}$. Rearranging $\tilde{F} < \bar{F}^2(\underline{\delta})$ gives:

$$P > \frac{(1 - \underline{\delta})(\pi_H^{D,l} - \pi_H^{RC,l}(\underline{\delta}) + \pi_H^{RC,h}(\underline{\delta}) - \pi_H^{N,h})}{\underline{\delta}(\pi_H^{RC,l}(\underline{\delta}) - \pi_H^{RC,h}(\underline{\delta}))} \equiv \tilde{P} \quad (2.21)$$

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

$$(1 - \underline{\delta}) \left[\pi_H^{D,l} - \pi_H^{N,h} \right] < \pi_H^{RC,l} - \pi_H^{RC,h} \quad \Leftrightarrow \quad (1 - \underline{\delta}) \pi_H^{D,l} - \pi_H^{RC,l} < (1 - \underline{\delta}) \pi_H^{N,h} - \pi_H^{RC,h} \quad (2.22)$$

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

$$\frac{c_m^h}{c_m^l} > C(\alpha, \beta, \eta), \quad (2.23)$$

where C is a positive term that depends on α , β , and η . Summing up, if the cost dispersion c_m^h/c_m^l is sufficiently large, condition (2.23) is satisfied and we have $\tilde{F} < \bar{F}^2(\tilde{\delta})$. Notice that this does not automatically ensure $\tilde{F} < \bar{F}^1(\tilde{\delta})$. However, even if $\bar{F}^1(\tilde{\delta}) < \tilde{F}$, this does not alter the validity of our prediction of a positive effect of match duration on RC propensity, since impatient agents never engage in RCs.

2.B Supplementary appendix

2.B.1 Heterogeneous supplier costs: Continuous case

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 $[c_m, \bar{c}_m]$.

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 re-match the current supplier, it must be the case that the expected gains must be larger than the continuation payoffs with the current match:

$$E[\pi_H^{RM,i} \mid c_m^0 = \widetilde{c}_m] > \frac{\widetilde{\pi}_H^i}{1 - \delta_H}, \quad (2.24)$$

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.

The term $E[\pi_H^{RM,i} \mid c_m^0 = \widetilde{c}_m]$ can be calculated as follows:

$$E[\pi_H^{RM,i} \mid c_m^0 = \widetilde{c}_m] = \widetilde{\pi}_H^N - F + \frac{\delta_H}{1 - \delta_H} E[\pi_H^i]$$

Note, that (2.24) has to be re-evaluated for every supplier encountered in the re-matching procedure.

Analogously as in section 2.4, from (2.24) we can derive the fixed cost thresholds $\bar{F}^k(\widetilde{c}_m)$, 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 $F < \bar{F}^k(\widetilde{c}_m)$ re-matching is incentive compatible, while otherwise it is not. We get the following thresholds:

$$\begin{aligned}\bar{F}^1(\widetilde{c}_m) &= \frac{\delta_H}{1-\delta_H} \left[E[\pi_H^N] - \widetilde{\pi}_H^N \right] \\ \bar{F}^2(\widetilde{c}_m) &= \widetilde{\pi}_H^N + \frac{1}{1-\delta_H} \left[\delta_H E[\pi_H^{RC}] - \widetilde{\pi}_H^{RC} \right]\end{aligned}$$

As in the main text, for the case where a RC is feasible (i.e. $\widetilde{\delta} \geq \underline{\delta}$) we must rule out any “cheat-and-run” (CAR) incentives of the firm, i.e. that it repeatedly deviates from the RC and re-matches to a new supplier. For this, search costs must not be too small and it must hold:

$$\frac{\widetilde{\pi}_H^{RC}}{1 - \delta_H} \geq E[\pi_H^{CAR} \mid c_m^0 = \widetilde{c}_m],$$

where $E[\pi_H^{CAR} \mid c_m^0 = \widetilde{c}_m] = \widetilde{\pi}_H^D - F + \frac{\delta_H}{1-\delta_H}(E[\pi_H^D] - F)$ is the firm’s off-equilibrium payoff when it always deviates from the RC. For CAR never to occur a sufficient condition is $E[\pi_H^{RM,RC} \mid c_m^0 = \widetilde{c}_m] > E[\pi_H^{CAR} \mid c_m^0 = \widetilde{c}_m]$, or:

$$F > \frac{1 - \delta_H}{\delta_H} \left(\widetilde{\pi}_H^D - \widetilde{\pi}_H^N \right) + E[\pi_H^D] - E[\pi_H^{RC}] \equiv \widetilde{F}(\widetilde{c}_m).$$

By construction, the condition also guarantees that $\bar{F}^2(\widetilde{c}_m) > \widetilde{F}(\widetilde{c}_m)$ for all \widetilde{c}_m . This gives the following result:

Proposition 4.

Suppose the headquarter is initially matched to a supplier with costs $c_m^0 = \widetilde{c}_m$.

- a) Impatient agents ($\widetilde{\delta} < \underline{\delta}$): Consider any $F > 0$. Impatient agents start re-matching if $F < \bar{F}^1(\widetilde{c}_m)$ and continue re-matching until, in period t , they find a supplier with costs c_m^{t+1} for which $F \geq \bar{F}^1(c_m^{t+1})$. The firm will engage in a LTC with any supplier for which $F \geq \bar{F}^1(c_m^{t+1})$ holds. The RC can never be implemented.
- b) Patient agents ($\widetilde{\delta} \geq \underline{\delta}$): Consider any $F > \widetilde{F}(\widetilde{c}_m)$. Patient agents start re-matching if $F < \bar{F}^2(\widetilde{c}_m)$ and continue until, in period t , they find a supplier with costs c_m^{t+1} for which $F \geq \bar{F}^2(c_m^{t+1})$. The RC forms with any supplier for which $F \geq \bar{F}^2(c_m^{t+1})$ holds. Otherwise, the RC cannot be implemented.

Notice, that for both, patient and impatient suppliers, the respective search cost threshold \bar{F}^k decreases in the cost level of the current supplier, \widetilde{c}_m . Thus, for given search costs it follows necessarily that the re-matching process stops at some level of supplier efficiency and a LTC is established. Proposition 4 confirms, that the logic of Proposition 2 extends to the case where the distribution of supplier costs is continuous.

2.B.2 Separated search and re-matching

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. In the following, we will first discuss the case where supplier costs are drawn from a two-point distribution as in the main text and then extend the analysis to the continuous case.

Two-point distribution of supplier costs

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* ($\widetilde{\delta} < \underline{\delta}$) 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 $F < \bar{F}^1$ from (2.13) therefore still applies.

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

$$\frac{1}{1 - \delta_H} \pi_H^{RC,l} \geq E[\pi_H^{CAR} \mid c_m^0 = c_m^l] \quad (2.25)$$

Compared to (2.18), 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:

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

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:

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

Using the equivalent steps from the baseline model, from (2.25) we can derive a lower threshold on search costs, \tilde{F}' , where $F > \tilde{F}'$ 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 (2.14), the expected payoff when engaging in search must be higher than the continuation payoff with the initial high-cost supplier, i.e.,

$$E[\pi_H^{search,RC} \mid c_m^0 = c_m^h] > \frac{1}{1 - \delta_H} \pi_H^{RC,h}. \quad (2.26)$$

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 $\tilde{\delta} > \underline{\delta}$. If condition (2.26) 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

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

and for M_0 the RC is incentive compatible if

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

These incentive compatibility constraints thus boil down to

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

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 $\tilde{\delta} \geq \frac{1}{1-P} \frac{\pi_H^{JFB} - \pi_H^D - \pi_M^D}{\pi_H^N + \pi_M^N - \pi_H^D - \pi_M^D} = \underline{\delta}/(1 - P) > \underline{\delta}$, where $\underline{\delta}$ is the previously derived critical discount factor given in (2.10). Put differently, very patient agents with an average 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 $\underline{\delta} < \tilde{\delta} < \underline{\delta}/(1 - P)$ 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 \bar{F}^2). For the very patient agents who always engage in RCs ($\tilde{\delta} > \underline{\delta}/(1 - P)$), the search decision can be formalized as

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

which yields these expected profits

$$E[\pi_H^{search, RCstrong} \mid c_m^0 = c_m^h] = \frac{1}{1 - \delta_H(1 - P)} \left[\pi_H^{RC,h} - F + \frac{\delta_H P}{1 - \delta_H} \pi_H^{RC,l} \right] \quad (2.27)$$

Plugging (2.27) into (2.26) and rearranging we obtain the critical search cost level \bar{F}^{2*} for the very patient agents case:

$$F < \frac{\delta_H P}{1 - \delta_H} \left[\pi_H^{RC,l} - \pi_H^{RC,h} \right] \equiv \bar{F}^{2*}. \quad (2.28)$$

Comparing (2.13), (2.16) and (2.28), it can be verified that $\bar{F}^{2*} > \bar{F}^2$ always holds. Moreover, as \bar{F}^1 and \bar{F}^2 , also \bar{F}^{2*} is increasing in δ_H (for given δ_M). Finally, we can make a similar consistency argument as in Appendix 2.A.2 to guarantee $\bar{F}^2 > \tilde{F}'$ which in turn implies $\bar{F}^{2*} > \tilde{F}'$ since $\bar{F}^{2*} > \bar{F}^2$.

We summarize our results under separated search and re-matching in the following Propositions 5 and 6, 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 5. *a) Suppose the headquarter is initially matched with a low-cost supplier ($c_m^0 = c_m^l$). Patient agents (with $\tilde{\delta} > \underline{\delta}$) 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 ($F > \tilde{F}'$). Impatient agents (with $\tilde{\delta} < \underline{\delta}$) will form a long-term collaboration of repeated Nash bargainings with that supplier.*

b) With $c_m^0 = c_m^h$, impatient agents with $\tilde{\delta} < \underline{\delta}$ search if $\tilde{F} < F < \bar{F}^1$ 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 $\underline{\delta} < \tilde{\delta} < \underline{\delta}/(1 - P)$ search if $\tilde{F} < F < \bar{F}^2$ 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 $\tilde{\delta} > \underline{\delta}/(1 - P)$ search if $\tilde{F} < F < \bar{F}^{2}$ 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 6. *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 \bar{F}^1 , \bar{F}^2 and \bar{F}^{2*} , 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 2.3, which is comparable to Figure 2.2. 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.




Figure 2.3: Contractual and search/re-matching decision ($c_m^0 = c_m^h$)

A MPS on the distribution of supplier costs has the same effect on \bar{F}^{2*} as it has on \bar{F}^2 and \bar{F}^1 and shifts the \bar{F}^{2*} -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 . As for Figure 2.2 in the main text, the concave form of the \bar{F}^k -functions obtains whenever we consider changes in δ_H while holding δ_M constant.

Continuous distribution of supplier costs

In the following, we extend the analysis to the case of a continuous distribution of supplier costs. These get drawn from a continuously differentiable cdf $G(c_m)$, and we denote by $g(c_m)$ the corresponding pdf with support $[\underline{c}_m, \overline{c}_m]$. The analysis is tightly related to that of Supplementary Appendix 2.B.1, where search and re-matching were no separable decisions. However, as already pointed out above, the separability of search and re-matching generates a class of very patient agents that engage in a RC while being on the search for a better partner.

Let \widetilde{c}_m denote the cost realization of the firm's current supplier, and let $\widetilde{\pi}_H^i$ be the firm's payoff at this cost realization given the contract structure $i = \{N, RC\}$. In order that the firm starts searching it must be the case that the expected gains are larger than the continuation payoffs with the current match:

$$E[\pi_H^{RM,j} \mid c_m^0 = \widetilde{c}_m] > \frac{\widetilde{\pi}_H^i}{1 - \delta_H}, \quad (2.29)$$

where the three cases of agents, $j \in \{\text{impatient, mildly patient, very patient}\}$, must be distinguished. Ultimately, once search has stopped impatient agents will engage in contract structure $i = N$ and both types of patient agents in $i = RC$. The term $E[\pi_H^{RM,j} \mid c_m^0 = \widetilde{c}_m]$ can be calculated as follows:

$$\begin{aligned} E[\pi_H^{RM, \text{impatient}} \mid c_m^0 = \widetilde{c}_m] &= \frac{1}{1 - \delta_H(1 - G(\widetilde{c}_m))} \left[\widetilde{\pi}_H^N - F + \frac{\delta_H G(\widetilde{c}_m)}{1 - \delta_H} E[\pi_H^N \mid c_m < \widetilde{c}_m] \right] \\ E[\pi_H^{RM, \text{mildly patient}} \mid c_m^0 = \widetilde{c}_m] &= \frac{1}{1 - \delta_H(1 - G(\widetilde{c}_m))} \left[\widetilde{\pi}_H^N - F + \frac{\delta_H G(\widetilde{c}_m)}{1 - \delta_H} E[\pi_H^{RC} \mid c_m < \widetilde{c}_m] \right] \\ E[\pi_H^{RM, \text{very patient}} \mid c_m^0 = \widetilde{c}_m] &= \frac{1}{1 - \delta_H(1 - G(\widetilde{c}_m))} \left[\widetilde{\pi}_H^{RC} - F + \frac{\delta_H G(\widetilde{c}_m)}{1 - \delta_H} E[\pi_H^{RC} \mid c_m < \widetilde{c}_m] \right] \end{aligned}$$

Note, that (2.29) has to be re-evaluated for every supplier that the firm encounters in the re-matching procedure. By rearranging (2.29), we can obtain the critical search costs \overline{F}^k for all the three patience levels that now each depend on the current cost realization \widetilde{c}_m . Using the same notation as in the previous subsection we get:

$$\begin{aligned} \overline{F}^1(\widetilde{c}_m) &= \frac{\delta_H G(\widetilde{c}_m)}{1 - \delta_H} \left[E[\pi_H^N \mid c_m < \widetilde{c}_m] - \widetilde{\pi}_H^N \right] \\ \overline{F}^2(\widetilde{c}_m) &= - \left(\widetilde{\pi}_H^{RC} - \widetilde{\pi}_H^N \right) + \frac{\delta_H G(\widetilde{c}_m)}{1 - \delta_H} \left[E[\pi_H^{RC} \mid c_m < \widetilde{c}_m] - \widetilde{\pi}_H^{RC} \right] \\ \overline{F}^{2*}(\widetilde{c}_m) &= \frac{\delta_H G(\widetilde{c}_m)}{1 - \delta_H} \left[E[\pi_H^{RC} \mid c_m < \widetilde{c}_m] - \widetilde{\pi}_H^{RC} \right] \end{aligned}$$

By using the same steps as in Supplementary Appendix 2.B.1, we can derive a lower bound on search costs $\widetilde{F}(\widetilde{c}_m)$ under which CAR can be ruled out. Because we have $\overline{F}^{2*}(\widetilde{c}_m) > \overline{F}^2(\widetilde{c}_m)$ for all \widetilde{c}_m it is sufficient to construct the lower bound for mildly

patient agents in order to rule out CAR for all agents. Hence we can formulate the following proposition:

Proposition 7.

Suppose the headquarter is initially matched to a supplier with costs $c_m^0 = \widetilde{c}_m$.

- a) Impatient agents ($\widetilde{\delta} < \underline{\delta}$): Consider any $F > 0$. Impatient agents start searching if $F < \overline{F}^1(\widetilde{c}_m)$ and continue searching until, in period t , they find a supplier with costs c_m^{t+1} for which $F \geq \overline{F}^1(c_m^{t+1})$. The firm will engage in a LTC with any supplier for which $F \geq \overline{F}^1(c_m^{t+1})$ holds. The RC can never be implemented.
- b) Mildly patient agents ($\underline{\delta}/(1-P) > \widetilde{\delta} \geq \underline{\delta}$): Consider any $F > \widetilde{F}(\widetilde{c}_m)$. Mildly patient agents start searching if $F < \overline{F}^2(\widetilde{c}_m)$ and continue until, in period t , they find a supplier with costs c_m^{t+1} for which $F \geq \overline{F}^2(c_m^{t+1})$. The RC forms with any supplier for which $F \geq \overline{F}^2(c_m^{t+1})$ holds. Otherwise, the RC cannot be implemented.
- b) Very patient agents ($\widetilde{\delta} \geq \underline{\delta}/(1-P)$): Consider any $F > \widetilde{F}(\widetilde{c}_m)$. Very patient agents start searching if $F < \overline{F}^{2*}(\widetilde{c}_m)$ and continue until, in period t , they find a supplier with costs c_m^{t+1} for which $F \geq \overline{F}^{2*}(c_m^{t+1})$. The RC forms with any supplier, in STCs and LTCs.

Moreover, the MPS result from above can be extended to continuous distributions:

Proposition 8. A mean-preserving spread on the continuous distribution of supplier costs increases the critical search cost levels \overline{F}^1 , \overline{F}^2 and \overline{F}^{2*} , and thereby expands the parameter range where the headquarter engages in search.

Proof. Suppose $\hat{G}(c_m)$ is a MPS on $G(c_m)$ and denote by \hat{F}^k , $k \in \{1, 2, 2^*\}$, the associated search cost thresholds. Then in order to show that the Proposition holds it is sufficient to show that, for all k :

$$\hat{F}^k - \overline{F}^k > 0.$$

Plugging in the expression from above this can be simplified to:

$$\hat{G}(\widetilde{c}_m) \left[E_{\hat{G}}[\pi_H^i | c_m < \widetilde{c}_m] - \widetilde{\pi}_H^i \right] - G(\widetilde{c}_m) \left[E_G[\pi_H^i | c_m < \widetilde{c}_m] - \widetilde{\pi}_H^i \right] > 0,$$

where $i \in \{N, RC\}$. Integrating the conditional expected values by parts, we can simplify this expression to:

$$\int_{\underline{c}_m}^{\widetilde{c}_m} \frac{\partial \pi_H^i}{\partial c_m} G(c_m) dc_m > \int_{\underline{c}_m}^{\widetilde{c}_m} \frac{\partial \pi_H^i}{\partial c_m} \hat{G}(c_m) dc_m \quad (2.30)$$

Observing that $\frac{\partial \pi_H^i}{\partial c_m} < 0$, we can conclude that the Proposition holds. \square

2.B.3 Heterogeneous discount factors: Continuous case

In this appendix, we analyze the question of supplier search and re-matching where the supplier's discount factor δ_M is heterogeneous and drawn from the continuous interval $(0, 1)$. We focus on the case of stationary bonus payments which ensures that the firm will set the same bonus in every period such that the supplier's IC constraint exactly binds. After every round of output realization the firm can decide whether to re-match with a new supplier. Search and re-matching cannot be separated. For this appendix, suppliers are homogeneous w.r.t. their costs c_m . The re-matching stage of the game is adjusted as follows.

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

We assume that the cumulative distribution function $G(\delta_M)$ is continuously differentiable on its entire support, $\delta_M \in (0, 1)$. Clearly, depending on search costs F and the discount factor δ_M^t of its current match, the firm will have an incentive to search for a more patient supplier. We assume δ_H to take a fixed value and to make the exposition interesting we restrain the values of it on the interval $\delta_H \in (\underline{\delta}_H, \overline{\delta}_H)$ introduced in the following Corollary to Proposition 1. This restriction guarantees that there will always be some impatient suppliers which which a RC will be impossible, while for patient ones it will be possible. Given that $\delta_H < \underline{\delta}_H$ the RC will always fail no matter how patient the supplier is. On the other hand, if $\delta_H > \overline{\delta}_H$ the RC will always work no matter how impatient the supplier is.

Corollary 1. *Suppose that $\delta_H \in (\underline{\delta}_H, \overline{\delta}_H)$, where $\underline{\delta}_H = \max \left\{ 0, \frac{\pi_H^D + \pi_M^N - \pi^{JFB}}{\pi_H^D - \pi_H^N} \right\}$ and $\overline{\delta}_H = \min \left\{ 1, \frac{\pi_H^D + \pi_M^D - \pi^{JFB}}{\pi_H^D - \pi_H^N} \right\}$. Then there exists $\underline{\delta}_M \in (0, 1)$ such that Nash will be played with all suppliers for which $\delta_M < \underline{\delta}_M$, and a RC with (h^*, m^*) is implementable for all $\delta_M \geq \underline{\delta}_M$.*

Proof. First, consider $\underline{\delta}_H$. If δ_H is “too small”, then it will never be possible to start a relational contract no matter how patient the supplier is, i.e. even when $\delta_M \rightarrow 1$. Suppose that indeed $\delta_M = 1$ and plug it into (1) which thus can be simplified to $\delta_H \leq \frac{\pi_H^D + \pi_M^N - \pi^{JFB}}{\pi_H^D - \pi_H^N}$. Whenever this condition is satisfied, a RC will not be possible no matter how patient the supplier is.

On the other hand, consider $\overline{\delta}_H$. If δ_H gets “too large”, then it will always be possible to start a relational contract no matter how impatient the supplier is, i.e.

even when $\delta_M \rightarrow 0$. For $\delta_M = 0$, (1) can be rearranged to $\delta_H \geq \frac{\pi_H^D + \pi_M^D - \pi^{JFB}}{\pi_H^D - \pi_H^N}$. Whenever this condition is satisfied, a RC will be possible no matter how impatient the supplier is.

Thus, because $\underline{\delta}_M \in (0, 1)$ and δ_M is continuous on $(0, 1)$ for all $\delta_H \in (\underline{\delta}_H, \overline{\delta}_H)$ there will exist a range of δ_M , namely $(0, \underline{\delta}_M)$, where Nash will be played and a range $[\underline{\delta}_M, 1)$ where a RC is implementable. \square

Now suppose that the initially matched supplier has discount factor $\tilde{\delta}_M$. Then in order that it is profitable for the firm to re-match, the expected gains $E[\pi_H^{RM}]$ must be larger than the continued payoffs from the current relationship, i.e.

$$E[\pi_H^{RM}] > \frac{\pi_H^i(\tilde{\delta}_M)}{1 - \delta_H}, \text{ where } i = \begin{cases} RC & \text{if } \tilde{\delta}_M > \underline{\delta}_M \\ N & \text{otherwise} \end{cases}. \quad (2.31)$$

Note, that for every match that the firm encounters in the process of re-matching equation (2.31) must be re-evaluated with the new realization of δ_M . Note, that the firm's per period payoff π_H^N with any impatient supplier, i.e. with $\delta_M < \underline{\delta}_M$, is independent of δ . We will therefore drop the current realization index for the impatient supplier case in the following. This fact bears the additional conclusion that whenever the firm engages in re-matching an impatient supplier it will only stop the re-matching process whenever it finds a patient supplier that allows for an RC. The expected payoffs from re-matching the current supplier can be expressed as follows:

$$E[\pi_H^{RM}] = \pi_H^N - F + \frac{\delta_H}{1 - \delta_H} [(1 - G(\underline{\delta}_M))E[\pi_H^{RC} \mid \delta_M \geq \underline{\delta}_M] + G(\underline{\delta}_M)\pi_H^N]$$

With this expression at hand, we can solve both cases of (2.31) for F and obtain:

$$\begin{aligned} F &\leq \frac{\delta_H(1-G(\underline{\delta}_M))}{1-\delta_H} [E[\pi_H^{RC} \mid \delta_M \geq \underline{\delta}_M] - \pi_H^N] && \equiv \overline{F}^3, \text{ if } i = N \\ F &\leq \frac{\delta_H(1-G(\underline{\delta}_M))}{1-\delta_H} [E[\pi_H^{RC} \mid \delta_M \geq \underline{\delta}_M] - \pi_H^N] - \frac{\pi_H^{RC} - \pi_H^N}{1-\delta_H} && \equiv \overline{F}^4(\tilde{\delta}_M), \text{ if } i = RC \end{aligned}$$

The expressions state that whenever $F < \overline{F}^k(\tilde{\delta}_M)$, $k \in \{3, 4\}$, is fulfilled the firm will re-match its current supplier and not do so otherwise. Observe, that while the search cost threshold \overline{F}^3 is constant for all $\tilde{\delta}_M < \underline{\delta}_M$, \overline{F}^4 is monotonically decreasing in $\tilde{\delta}_M$ for all patient agents. Note also, that for all patience levels $\overline{F}^4(\tilde{\delta}_M) < \overline{F}^3$ holds.

As before, in order to characterize the re-matching behavior of firms we need to rule out off-equilibrium “cheat-and-run” (CAR) behavior. For this, suppose that we

are matched with some $\tilde{\delta}_M > \underline{\delta}_M$ and (2.31) does not hold, i.e. we do not re-match the current supplier. Formally:

$$\frac{\pi_H^{RC}(\tilde{\delta}_M)}{1 - \delta_H} \geq E[\pi_H^{RM}] \quad (2.32)$$

Simultaneously, for CAR not to occur we must have:

$$\frac{\pi_H^{RC}(\tilde{\delta}_M)}{1 - \delta_H} \geq E[\pi_H^{CAR}], \quad (2.33)$$

where $E[\pi_H^{CAR}]$ is the solution to:

$$\begin{aligned} V_0 &= \pi_H^D - F + \delta_H V_1 \\ V_1 &= (1 - G(\underline{\delta}_M))\pi_H^D + G(\underline{\delta}_M)\pi_H^N - F + \delta_H V_1 \Leftrightarrow V_1 = \frac{(1 - G(\underline{\delta}_M))\pi_H^D + G(\underline{\delta}_M)\pi_H^N - F}{1 - \delta_H} \end{aligned}$$

Inspection of (2.32) and (2.33) reveals that for CAR not to occur it is sufficient that $E[\pi_H^{RM}] > E[\pi_H^{CAR}]$ holds, which can be reduced to the condition of a lower bound \tilde{F} on F :

$$F > \pi_H^N + (1 - G(\underline{\delta}_M))(\pi_H^D - E[\pi_H^{RC} \mid \delta_M \geq \underline{\delta}_M]) \equiv \tilde{F}$$

With this bound at hand we can formulate the following proposition.

Proposition 9. Suppose the headquarter is currently matched to a supplier with patience level $\tilde{\delta}_M \in (0, 1)$, for which $\tilde{\delta} < \underline{\delta}$ ($\tilde{\delta} \geq \underline{\delta}$) holds.

a) Then, whenever she faces search costs $F \geq \bar{F}^3(\tilde{\delta}_M)$ ($F \geq \bar{F}^4(\tilde{\delta}_M)$) she will engage in a LTC with this supplier in the form of repeated Nash bargainings (a relational contract).

b) On the other hand, suppose that $\tilde{F} < F < \bar{F}^3(\tilde{\delta}_M)$ ($\tilde{F} < F < \bar{F}^4(\tilde{\delta}_M)$). Then the firm will play Nash with the current supplier and re-match him at the end of the initial period. The firm will continue re-matching until, in period t , she finds a supplier for which $\delta_M^{t+1} \geq \underline{\delta}$ and $F \geq \bar{F}^4(\delta_M^{t+1})$ hold. With this supplier, the firm will start a long-termed RC.

This immediately leads to the following Corollary:

Corollary 2. Whenever the firm decides to start re-matching, any future LTC will involve a relational contract. However, if she engages in a LTC right away, repeated Nash play cannot be ruled out.

2.B.4 “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. For this section we assume that $\delta_H = \delta_M = \delta$. 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.

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

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

First, consider (IC_{*i*}-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_{*i*}-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 $\hat{\delta}$:

$$\hat{\delta}^T \pi^{JFB} - \hat{\delta}(\pi_H^D + \pi_M^D) + (\pi_H^D + \pi_M^D - \pi^{JFB}) = 0 \quad (2.34)$$

From (2.34) 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:

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

The solution to this system of equations is given by:

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

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.

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

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

$$(1 - \delta^{T+1})V^{+H} \geq (1 - \theta)\pi_H^D \quad (\text{IC}_H)$$

$$(1 - \delta^{T+1})V^{+M} \geq B, \quad (\text{IC}_M)$$

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

$$\begin{aligned} \psi\pi_H^{RC} - (1 - \theta)\pi_H^D &\geq 0 \\ \psi\pi_M^{RC} - B &\geq 0 \end{aligned} \quad (2.35)$$

Plugging π_M^{RC} into (2.35) the latter can be rearranged as $B \geq \frac{\psi}{1 - \psi}\tau c_m m^*$, 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 (2.35), we can compute $\tilde{\delta}$ implicitly from the following equation, which inter alia depends also on θ and T :

$$\psi \left(R^* - h^* c_h - \frac{\psi}{1 - \psi} \tau c_m m^* \right) - (1 - \theta) \pi_H^D = 0 \quad (2.36)$$

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 (2.36). However, the more important observation for our purpose is that c_m can be factored out of the LHS of (2.36). 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.

Chapter 3

Supplier search and re-matching in global sourcing – theory and evidence from China

Co-authored with Jens Südekum and Fabrice Defever

3.1 Introduction

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

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

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

A starting point of the global sourcing literature is the canonical incomplete contracts model by Antràs (2003) and Antràs and Helpman (2004), which places the seminal property rights theory of Grossman and Hart (1986) and Hart and Moore (1990) in an international trade context. It features hold-up and underinvestment problems as the key elements of the buyer-supplier relationship, which can be fine-tuned by the firm's decision how to organize the production process and the ownership structure along the value chain. As it is well known, these baseline models by Antràs (2003) and Antràs and Helpman (2004) are static models of a one-shot interaction and, therefore, not suitable to study long-term collaborations. However,

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

several authors have recently embedded this model into a dynamic setting with repeated interactions, similarly as in Baker, Gibbons and Murphy (2002). In particular, this includes Kukharskyy (2016), Kamal and Tang (2015) and Kukharskyy and Pflüger (2016), who show that the inefficiencies caused by contract incompleteness can be overcome by *relational contracts*, i.e., by informal, non-binding and purely trust-based agreements to cooperate in an enduring relationship. These papers study the crucial ownership choice of integration versus outsourcing in such a dynamic global sourcing model.² However, they consider a setting with a fixed partner, and therefore cannot capture the evidence by Rauch and Watson (2003) and Besedes (2008) on trial-and-error search for efficient and reliable component manufacturers.

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

The aim of this chapter is to complement the analysis by Defever et al. (2016) in some important dimensions. In particular, we start from a stripped-down version of our model and completely assume away contract incompleteness. The investment levels of the two partners are fully verifiable and contractible, and hence, relational

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

contracts are not needed to induce efficiency in this environment. This simplification allow us to put our focus exclusively on the firm's search and matching problem. The trade-off that we study in this chapter is simple: The headquarter is currently matched with a supplier of a certain type, more specifically, with a certain unit cost level that determines his efficiency. The firm has an interest to collaborate with a supplier that is as efficient as possible, because this also raises her own profits. But supplier switching is subject to a fixed cost for search and re-matching. The firm then essentially weighs this fixed costs against the expected gains from continued search, and will stop searching if it has – at some point – found a supplier that is efficient enough, i.e., with unit costs below a certain threshold level.

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

In the final section of this chapter, we then aim to address the empirical relevance of this novel theoretical prediction. The main challenge we are facing here is data availability. Our theory is about particular matches of domestic buyers and foreign suppliers, but current data only very rarely allows observing such matches. A notable exception is the paper by Eaton, Eslava, Jinkins, Krizan and Tybout (2014) who are able to construct pairs of Colombian exporting firms and their importers in the 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

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

co-existence of long-term and one-shot collaborations in that market. More recent studies which also use proprietary Census data on matches of US importers and foreign exporters include Monarch (2015) and Kamal and Tang (2015).

In this paper, we exploit Chinese customs data that encompass all export transactions to the US, Japan, UK and Germany between 2000 and 2006. We build a sample of firms that start exporting a particular product (6-digit industry level) to one of those destinations, but unfortunately we cannot identify the buyer. Still, following these fresh Chinese exporters over time, we observe which firms still export the same product to the same destination after a few years, as opposed to those which have terminated that exporting activity in the meantime. This allows us to build a proxy whether a particular export transaction can be thought of as a long-term collaboration or rather a one-shot deal, and to construct the share of short-term relationships across Chinese HS6 industries. We correlate this share with an industry-wide measure of cost dispersion across Chinese firms, which we construct from unit-value data.

Our model tells us that this correlation between cost dispersion and the share of short-term collaboration should be positive, because headquarter firms tend to search more, and will thus exhibit more supplier turnover and one-shot deals in a given time frame. And indeed, the empirical evidence from our Chinese data is firmly in line with this theoretical prediction. More precisely, by aggregating the Chinese export transactions at the HS6 product level, we consistently find a positive correlation between the share of one-shot transactions and the variance in the log of unit-value within an industry. This result is robust for exports across different destination markets, for independent Chinese firms as well as for Chinese foreign affiliates, as well as for different types of export transactions. Moreover, apart from that, this chapter also establishes some interesting stylized facts about the resilience of Chinese exporters that may be useful independently from the empirical test of our model.

The rest of this chapter is organized as follows. Section 3.2 describes our model, and Section 3.3 introduces the data and present the empirical results. Section 3.4 provides a brief conclusion.

3.2 The model

3.2.1 Environment

Consider an industry where a single headquarter firm (labelled H) sources a customized intermediate input to produce a final good. There are many suppliers in that industry who could potentially manufacture this component, and those suppliers differ in their efficiency. In particular, we assume that there is a continuum of potential suppliers, which are all identical except for the constant marginal costs c

that they would incur in the production of the input. Costs are distributed across potential suppliers according to some probability density function $g(c)$ with support $[\underline{c}, \bar{c}]$, where $\underline{c} > 0$. The corresponding cumulative density function $G(c)$ is continuously differentiable in c .

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

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

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

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

and revenue sharing. Those elements could be easily introduced, but they are not necessary to derive the main theoretical predictions that we later take to the data.

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

Finally, in the baseline version of our model studied in Section 3.2.2, we assume that the firm has the option to stick to her old supplier if she is not satisfied with the candidate supplier that she encounters during the search in any given period. On the one hand, this is realistic since suppliers have small (zero) outside options and would therefore always prefer to stay inside the relationship, even if the firm actively seeks for a better partner. On the other hand, there are also arguments in favour of an alternative assumption, namely that the firm *must* re-match if it decides to engage in searching. Specifically, this alternative scenario can be rationalized in terms of the information that is available to the firm about the candidate that she encounters in the search process. In the baseline version, the firm can perfectly observe his cost-efficiency before actually re-matching to him. This allows her to *separate* the steps of search and re-matching into two distinct decisions. In the second scenario, however, we may assume that the firm can only observe the costs of the newly drawn supplier *ex post*, i.e., only after the first period of collaboration. When the candidate's costs are uncertain, the firm will always re-match when engaging in search (otherwise she would not pay F), and the search and re-matching decisions then become *non-separable*. We study this alternative setting below in Section 2.3.

3.2.2 Separable search and re-matching

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

supplier search incentive constraint:

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

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

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

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

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

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

$$\begin{aligned} V_0 &= \tilde{\pi} - F + \delta V_1 \\ V_1 &= \frac{1}{1 - \delta(1 - G(\tilde{c}))} \left[(1 - G(\tilde{c}))(\tilde{\pi} - F) + \frac{G(\tilde{c})E[\pi \mid c < \tilde{c}]}{1 - \delta} \right] \end{aligned} \quad (3.1)$$

Using (3.1) in (IC-Search), we can then solve for F and obtain a search cost threshold $\bar{F}(\tilde{c})$ that determines whether or not the headquarter firm will engage in supplier search:

$$F < \frac{\delta}{1 - \delta} G(\tilde{c}) [E[\pi \mid c < \tilde{c}] - \tilde{\pi}] \equiv \bar{F}(\tilde{c}) \quad (3.2)$$

Expression (3.2) demonstrates that the fixed costs F must be low enough in order for search to occur. Notice that the critical search cost level $\bar{F}(\tilde{c})$ depends positively on δ , i.e., search is more attractive the more patient agents are. The intuition is that future payoffs matter more, and hence it becomes more important to find an efficient supplier with low costs. Moreover, $\bar{F}(\tilde{c})$ depends positively on $G(\tilde{c})$ and on the term $[E[\pi \mid c < \tilde{c}] - \tilde{\pi}]$. This shows that search is more likely the higher is the chance to find a more efficient supplier, and the larger is the headquarter's expected per-period payoff difference with a better and the current supplier.

We summarize the result in the following proposition:

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

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

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

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

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

$$\hat{F}(\tilde{c}) > \bar{F}(\tilde{c}). \quad (3.3)$$

Substituting (3.2) into (3.3) and simplifying gives

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

Plugging in the conditional expected values and simplifying gives

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

Integrating by parts allows us to simplify further:

$$\int_{\underline{c}}^{\tilde{c}} \frac{\partial \pi}{\partial c} G(c) dc > \int_{\underline{c}}^{\tilde{c}} \frac{\partial \pi}{\partial c} \hat{G}(c) dc. \quad (3.4)$$

Because $\frac{\partial \pi}{\partial c} < 0$, it follows from the definition of the MPS that expression (3.4) is true for any value of \tilde{c} . From this we can conclude that a more dispersed cost distribution in an industry increases the headquarter's search and re-matching propensity.

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

3.2.3 Non-separable search and re-matching

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

- 2'. **Re-matching stage:** The firm decides whether to re-match with a new supplier. When deciding to re-match, she incurs a fixed cost $F > 0$. Let c^t be the unit cost level of her current supplier, and c^{t+1} the unit cost of the new supplier that she has encountered during her search. The cost c^{t+1} is randomly drawn from the distribution function $G(c)$ which is i.i.d. over periods. The firm cannot observe c^{t+1} before the start of a transaction relationship and hence re-matches whenever she searches.

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

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

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

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

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

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

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

$$F < \frac{\delta}{1 - \delta} \left[\left(E[\pi] - \tilde{\pi} \right) - (\delta(1 - G(\tilde{c})) \left(E[\pi \mid c > \tilde{c}] - \tilde{\pi} \right) \right] \equiv \tilde{F}(\tilde{c}) \quad (3.5)$$

Expression (3.5) demonstrates that search costs F have to be low enough in order for re-matching to occur. Again, $\tilde{F}(\tilde{c})$ depends positively on the patience level δ . Notice that because $(E[\pi \mid c > \tilde{c}] - \tilde{\pi})$ is always negative, $(E[\pi] > \tilde{\pi})$ is a sufficient condition for $\tilde{F}(\tilde{c}) > 0$, i.e., having a below-average supplier is always enough for re-matching to occur if search costs are small enough.

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

Proposition 1'. *Suppose the headquarter is currently matched with a supplier with marginal input production costs \tilde{c} . When the headquarter can infer the cost-efficiency*

of a candidate supplier only ex-post to contracting, she will engage in re-matching whenever $F < \tilde{F}(\tilde{c})$ and stick to the current match otherwise.

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

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

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

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

$$\dot{F}(\tilde{c}) > \tilde{F}(\tilde{c}). \quad (3.6)$$

Substituting (3.5) into (3.6) and simplifying gives:

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

We now plug in the conditional expected values and simplify:

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

Integrating by parts we arrive at:

$$\int_{\tilde{c}}^{\bar{c}} \frac{\partial \pi}{\partial c} \dot{G}(c) dc > \int_{\tilde{c}}^{\bar{c}} \frac{\partial \pi}{\partial c} G(c) dc \quad (3.7)$$

Because $\dot{G}(c)$ is a MPS of $G(c)$, and $\frac{\partial \pi}{\partial c} < 0$, it follows that expression (3.7) is true for all values of \tilde{c} . This concludes the proof.

3.2.4 Summary of theoretical results and main prediction

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

Figure 3.1 gives a graphical summary. Making the additional structural assumption that $\frac{\partial^2 \pi}{\partial \tilde{c}^2} < 0$, we can draw \bar{F} as a function that is increasing and convex in \tilde{c} . Then, the model predicts that for all (\tilde{c}, F) -combinations that lie below this function the firm will engage in search for a better supplier and hence be in a short-term collaboration (STC) with its current partner (blue area of the Figure). Otherwise, for all (\tilde{c}, F) -combinations above that convex curve, the firm will not search but establish a long-term collaboration (LTC) with its current match. Performing a MPS on the cost distribution $g(c)$ shifts the \bar{F} -function upwards to \hat{F} , and thereby enlarges the parameter range for which the firm will engage in supplier search / re-matching.




Figure 3.1: Illustration of search threshold and MPS

The fact that our results do not differ qualitatively between the two cases underscores that our predictions are robust with respect to the quality with which information on supplier costs is transmitted to the headquarter. In particular, it indicates that the prediction that a larger cost dispersion within an industry leads to a higher search and re-matching propensity is not crucially influenced by how fast firms learn about the supplier's efficiency. This gives us confidence in our industry-level prediction of a positive relationship between the cost-dispersion of potential suppliers and firms' propensity to engage in supplier search and re-matching.

3.3 Empirical analysis

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

3.3.1 Data and variables

Export transactions 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, Japan, UK and Germany, which are the top four main destinations for Chinese exports besides Hong-Kong and Macao. Specifically, our sample is composed of Chinese firms i which export a product j to a country c 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. We exclude all industries j with fewer than 10 active firms, and those where control variables are not available.

As it is shown in the upper panel of Table 3.1, an important observation is that quite a few of the new Chinese exporters in our data are owned by multinational corporations. This includes Chinese foreign affiliates which are dependent subsidiaries

or part of a joint venture. Altogether, the final sample includes 28,370 fresh Chinese exporters (15,017 domestic firms and 13,353 foreign affiliates), starting a total of 166,032 new export transactions, and spans 1,192 different HS6-digit product categories. In our baseline estimations we include all these observations, but we also conduct robustness checks that focus only on independent Chinese exporters.

Regarding the destination markets, while our full sample pools transaction with US, Japan, UK and Germany, we also present results for each market independently, and the lower panels of Table 3.1 show the number of Chinese exporters, exported products and export transactions in those four markets. In addition, for our baseline estimations we keep all types of goods, and in the robustness checks we limit this sample to complex goods with a relatively high degree of relationship-specificity, as this might better reflect the context of our theoretical model.

Table 3.1: Firms, products and destinations

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

Supplier turnover: Following the fresh Chinese exporters over time, we observe which firms that started an exporting activity in the country c and product j in year t , terminate its export activity in year $t + 3$, as opposed to those that still export the *same* product j to the *same* country c after three years. Ideally, we would like to observe if the Chinese exporter still deals with the same importer, but the data do not allow us to do so. Nevertheless, looking specifically at the US as the country of destination in Table 3.3, 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 (43.4%).

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 have stopped in year $(t + 3)$. This gives us our measure on the supplier turnover in industry j , which we use as our main explanatory variable: the share of short-term transactions.

Industry-wide cost dispersion: In order to construct an empirical proxy for the industry-wide cost dispersion, we approximate the transaction-specific marginal cost of the respective supplier by the unit-value of each transaction. The cost dispersion measure for industry j is then given by the standard deviation of the log of unit-values within the respective 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 Proposition 1 of our model, where re-matching is more likely the higher are the marginal costs of the current supplier.

Manova and Zhang (2012) argue that differences in unit-values could also reflect quality variations. In this case, our empirical proxy would be capturing heterogeneity in suppliers' quality. However, under iso-elastic demand and monopolistic competition, a marginal cost or a quality difference 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 decrease of supplier turnover. Hence, the interpretation of our results would be robust to the way we interpret differences in unit-values.

3.3.2 Descriptive evidence

Table 3.2 reports the duration of the transactions for independent firms and for foreign affiliates. It is important to take these ownership structures into account, as Chinese foreign affiliates may differ systematically in their economic behaviour from independent contractors. While the ownership dimension does not feature in our theoretical model, the key trade-off formalized by our model is – in principle – not limited to arm's length outsourcing relationships, but can also arise with a similar intuition in the context of intra-firm trade.

Looking at Table 3.2, of the 122,707 transactions made by domestically-owned firms, it turns out that 59,361 (48.4%) are terminated after less than one year, while 34,175 (27.9%) of them endure for more than 3 years. While transactions involving foreign affiliates tend to be more resilient, still 28.2% of the transactions stop after the first year. Digging deeper into the different countries of destinations (US, Japan, UK and Germany), Table 3.3 shows that the share of one-year transactions within these four subgroups is rather similar across destinations (from 40.4% for Japan to 45.8% for UK).

Table 3.2: Duration of export transactions by ownership

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

Table 3.3: Duration of export transactions by export destination

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

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

Ordinary least squares (OLS) regressions. In Table 3.4, we present some OLS results at the industry level. To do this, we first aggregate the data to the HS6 product-destination level. Specifically, we use all new transactions by fresh Chinese exporting firms i in industry j to country c . We create a dummy variable indicating if the transaction is a short-term transaction (value 1) or if it is still observed in $t + 3$ (value 0), and may thus be classified as a long-term collaboration. We then obtain our dependent variable by aggregating the data at the product-destination level by computing the share of short-term transactions among all the transactions within a HS6 product. All our baseline regressions include a set of destination fixed effects δ_c .

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

In column 1, we regress our dependent variable on the within-product standard deviation of the log of unit prices, while controlling for the ownership status of the Chinese firm. We find a positive correlation between the share of short-term transactions and the within-industry spread in unit value, as expected from Proposition 2 of our theoretical model. In column 2, we introduce a set of additional industry-level control variables, namely the capital-labor ratio, the human capital intensity, and the R&D-sales ratio, which are standard measures for headquarter-intensity. Moreover, we include the relationship-specificity index by Nunn (2013), and add demand elasticities from Broda and Weinstein (2006).⁴ Adding these control variable leaves the main coefficient unchanged.

In Columns 3 and 4 of Table 3.4, we exclude all the export transactions of foreign affiliates before aggregating the data at the product level. This leads to a decreasing number of industries as we still exclude all products with less than 10 observations. The main coefficient remains very similar and remains significant at the 1% level. Overall, this indicates that our results remain robust also in a smaller sample only consisting of independent Chinese-owned component manufacturers.

Next, in Table 3.5 we report results separately for transactions realized in each of the top-four export destinations, that is US, Japan, UK and Germany. Our main coefficient of interest remains stable in all specifications, and is always significant at the 1% level, regardless of whether we include or exclude foreign affiliates or whether we consider the additional industry-level control variables.

Taken together, we find that the results in Tables 3.4 and 3.5 are firmly in line with the predictions of our theoretical model, more specifically, with Proposition 2 which states higher cost dispersion among the suppliers leads to more search and re-matching in that industry, and thus negatively affects observed match durations.

⁴As data are disaggregated at the HS 10 digit level, we use the median value for each HS6 product.

Table 3.4: Relational-contracts and long-term collaborations: Transaction-level

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

Notes: Dependent variable: Share of export transactions which are one-shot on the overall number of transactions within a HS6 product. Robust standard-errors in brackets using clustered standard errors at the HS6 level.

3.3.3 Robustness checks and discussion

Table 3.6 provides a battery of robustness checks. The first one consists of dropping all the transactions which include processing trade. As shown in Defever et al. (2016), processing transactions are usually associated with high buyer-supplier interrelationship and relational contracts. Column 1 shows that the main coefficient for the within sector price dispersion remains extremely similar compared to our baseline results.

In column 2 we only keep industries where the Nunn-measure is above the median, and in column 3 we limit the sample even further to industries in the upper quartile of specificity. In fact, our baseline estimates are for all 1,004 product categories, some of which are rather standardized goods or final products. Since our model features the duration of a relationship between buyer and supplier when producing a specific intermediate input, we now limit the sample and – rather than just adding the Nunn measure as a control variable – focus on product categories with a high degree of relationship-specificity. As can be seen, all results remain quali-

tatively unchanged, even if the point estimates become slightly smaller in Column 3. Even when focussing on highly relationship-specific goods, we thus find that the within-sectoral cost-spread increases the proportion of short-term transactions.

In column 4, since many firms export multiple products, we exclude all the transactions 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. We only keep the transactions that belong to this main 2-digit sector of each firm, and again we find that the coefficients are very similar to our baseline estimates.

We have also conducted all these robustness checks separately for the different destination markets, along the lines of Table 3.5, and found that the pattern remains very stable in all these specifications. That is, in line with Proposition 2 of our model, we find a robust positive correlation between the industry-wide measure of cost dispersion and the share of short-term export collaborations in that industry, indicating that headquarters tend to engage in more supplier search when there is more cost dispersion among potential suppliers.

3.4 Conclusion

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

This key prediction of our model is empirically supported in the context of export transactions of fresh Chinese suppliers to the US, Japan, UK and Germany, the top four main destinations for Chinese exports besides Hong-Kong and Macao. By aggregating the Chinese export transactions at the HS6 product level, we show that the share of short-term transactions increase with an within-industry cost dispersion, which we construct from unit-value data. This result is robust for exports across different destination markets, for independent Chinese firms, as well as for Chinese foreign affiliates.

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

Table 3.5: Relational-contracts and long-term collaborations: Industry-level

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

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

Notes: Dependent variable: Share of export transactions which are one-shot on the overall number of transactions within a HS6 product. Robust standard-errors in brackets using clustered standard errors at the HS6 level.

Table 3.6: Relational-contracts and long-term collaborations: Transaction-level

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

Notes: Dependent variable: Share of export transactions which are one-shot on the overall number of transactions within a HS6 product. Robust standard-errors in brackets using clustered standard errors at the HS6 level.

3.5 References

- AEBERHARDT, R., I. BUONO, AND H. FADINGER (2014): “Learning, Incomplete Contracts and Export Dynamics: Theory and Evidence from French Firms,” *European Economic Review*, 68, 219–249.
- ANTRÀS, P. (2003): “Firms, Contracts, and Trade Structure,” *The Quarterly Journal of Economics*, 118(4).
- ANTRÀS, P. (2015): *Global Production: Firms, Contracts and Trade Structure*.
- ANTRÀS, P., AND A. COSTINOT (2011): “Intermediated Trade,” *The Quarterly Journal of Economics*, 126, 1319–1374.
- ANTRÀS, P., AND E. HELPMAN (2004): “Global Sourcing,” *Journal of Political Economy*, 112(3), 552–580.
- ARAUJO, L., G. MION, AND E. ORNELAS (2016): “Institutions and Export Dynamics,” *Journal of International Economics*, 98(1), 2–20.
- BAKER, G., R. GIBBONS, AND K. J. MURPHY (2002): “Relational Contracts and the Theory of the Firm,” *The Quarterly Journal of Economics*, 117(1), 39–84.
- BERNARD, A. B., A. MOXNES, AND K. H. ULLTVEIT-MOE (2014): “Two-Sided Heterogeneity and Trade,” Working Paper 20136, National Bureau of Economic Research.
- BESEDEŠ, T. (2008): “A Search Cost Perspective on Formation and Duration of Trade,” *Review of International Economics*, 16(5), 835–849.
- BOARD, S. (2011): “Relational Contracts and the Value of Loyalty,” *American Economic Review*, 101(7), 3349–67.
- BRODA, C., AND D. E. WEINSTEIN (2006): “Globalization and the Gains from Variety,” *The Quarterly Journal of Economics*, 121(2), 541–585.
- CARBALLO, J., G. OTTAVIANO, AND C. VOLPE-MARTINCUS (2013): “The Buyer Margins of Firms’ Exports,” Discussion Paper 1234, CEP.
- CORTS, K. S., AND J. SINGH (2004): “The Effect of Repeated Interaction on Contract Choice: Evidence from Offshore Drilling,” *Journal of Law, Economics, and Organization*, 20(1), 230–260.
- DEFEVER, F., C. FISCHER, AND J. SUEDEKUM (2016): “Relational Contracts and Supplier Turnover in the Global Economy,” *Journal of International Economics*, forthcoming.

- DWYER, R. F., P. H. SCHURR, AND S. OH (1987): “Developing Buyer-Seller Relationships,” *The Journal of Marketing*, 51(2), 11–27.
- EATON, J., M. ESLAVA, D. JINKINS, C. J. KRIZAN, AND J. TYBOUT (2014): “A Search and Learning Model of Export Dynamics,” Working Paper.
- FERNANDES, A. P., AND H. TANG (2015): “Scale, Scope, and Trade Dynamics of Export Processing Plants,” *Economics Letters*, 133(1), 68–72.
- GROSSMAN, S. J., AND O. D. HART (1986): “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *The Journal of Political Economy*, 94(4), 691–719.
- HART, O. D., AND J. MOORE (1990): “Property Rights and the Nature of the Firm,” *Journal of Political Economy*, 98(6), 1119–58.
- KALWANI, M. U., AND N. NARAYANDAS (1995): “Long-Term Manufacturer-Supplier Relationships: Do They Pay off for Supplier Firms?,” *The Journal of Marketing*, 59(1), 1–16.
- KAMAL, F., AND H. TANG (2015): “Relational Trade: Evidence from the United States,” Mimeo.
- KUKHARSKYY, B. (2016): “Relational Contracts and Global Sourcing,” *Journal of International Economics*, 101, 123 – 147.
- KUKHARSKYY, B., AND M. PFLÜGER (2016): “Time is On My Side: Relational Contracts and Aggregate Welfare,” Mimeo.
- MACCHIAVELLO, R., AND A. MORJARIA (2014): “Competition and Relational Contracts: Evidence from Rwanda’s Coffee Mills,” Working Paper.
- MANOVA, K., AND Z. ZHANG (2012): “Export Prices Across Firms and Destinations,” *The Quarterly Journal of Economics*, 127, 379–436.
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71(6), 1695–1725.
- MELITZ, M. J., AND S. J. REDDING (2014): “Heterogeneous Firms and Trade,” *Handbook of International Economics*, 4, 1–54.
- MONARCH, R. (2015): “It’s Not You, It’s Me: Breakups in U.S.-China Trade Relationships,” Working paper.
- NUNN, N. (2007): “Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade,” *The Quarterly Journal of Economics*, 122(2), 569–600.

-
- RAUCH, J. E. (1999): “Networks versus Markets in International Trade,” *Journal of International Economics*, 48(1), 7–35.
- RAUCH, J. E., AND J. WATSON (2003): “Starting Small in an Unfamiliar Environment,” *International Journal of Industrial Organization*, 21(7), 1021–1042.

Chapter 4

Payment contracts and trade finance in export relationships

4.1 Introduction

Shipping goods internationally is risky and takes time. Incomplete cross-border contract enforcement exposes trading partners to an omnipresent risk of expropriation (Thomas and Worrall 1994). At the same time, shipping and finally selling goods abroad is time-consuming which challenges firms with a decision of how to finance the working capital necessary for their export transactions. Taken together, these aspects point at the importance of selecting an appropriate payment contract such that risks and costs of exporting goods are managed optimally.¹

Which trade-offs do trade partners face in their decision on the optimal payment contract? Existing studies show that the (relative) costs of capital of exporter and importer as well as the (relative) quality of contracting institutions in their respective countries are crucial. When bank finance is considered the efficiency of the relevant banking market also matters (see Schmidt-Eisenlohr 2013). Most of the available literature frames firms' payment contract choice in a static context, i.e. as a situation where exporters and importers transact only once. However, a substantial share of businesses worldwide interact with their export partners on a repeated basis, suggesting that the mechanics of payment contract selection in the dynamic context of business relationships deserve more study.² It is well-known from a literature on incomplete relational contracts that repeated interaction can allow for self-enforcing agreements between firms that can substitute for well-functioning legal institutions (see e.g. Baker, Gibbons, and Murphy 2002). As a consequence, existing static models on payment contract choice may fundamentally mispredict the effective costs of financing trade and therefore also the role of external bank or insurer finance that can help to facilitate export activities.

In this paper I develop a repeated-game model of payment contract choice that provides highly tractable predictions on the optimal sequence of payment contracts in export relationships. At the same time it is general enough to cover any quality of contracting institutions and applies to a wide range of refinancing costs and demand characteristics that firms may face. The model shows that, when choosing the payment contract type, the exporter faces a trade-off between maximizing the payoffs from the current transaction and making use of the screening opportunities that differ inherently between format types. Such screening can help to attain information about the reliability of importers which is crucial to develop efficient and intensive export relationships.

¹To finance the time gap between production and sale of traded goods the following categories of payment contracts exist: open account (exporter finance), cash in advance (importer finance), and different types of bank-intermediated finance. For a comprehensive overview of payment methods in international trade, see ITA (2007).

²There exists plenty of evidence that international business transactions frequently involve long-term collaboration. See e.g. Egan and Mody (1992), Rauch and Watson (2003), Besedeš (2008), and Defever, Fischer, and Suedekum (2016).

Section 4.2 introduces the model. It considers an exporter (he) who sells a product through an importer (she) abroad. The importer's type is her private information and acts as a measure of her costs of capital. In every period, the exporter proposes a contract to his importer specifying an export quantity and a payment contract for the current period. This contract can only be enforced with a certain probability that depends on the quality of the contracting institutions of the participants' countries. The question of how to finance trade is relevant in the model because goods shipped abroad can only be sold in the destination market in the subsequent period.

Section 4.3 then contrasts cash in advance and open account as diametrically opposed payment contract forms. I derive equilibria that study each of these payment types in isolation and that highlight the fundamental differences between the contracts. Most importantly, I show that while under cash in advance exporters can offer a contract to importers that perfectly picks out reliable types, any open account contract necessarily involves pooling of importers.

With this finding at hand, in Section 4.4 the exporter can freely choose between these two payment contracts leading to the trade-off between information acquisition and payoff maximization sketched above. As a corollary, the result allows to give sharp predictions on the optimal sequence of payment contracts within an export relationship. Most importantly, this sequence depends on the distribution of importer types and the quality of contracting institutions, among other factors.

Section 4.5 introduces bank finance into the model and shows, how the availability of a trade credit insurance affects and augments the selection of payment contracts in the dynamic context. Such an insurance can help to speed up export growth as it can substitute for weak legal institutions. However, the benefits of such payment guarantees are only temporary because in later periods an importer's reputation for reliability may make a costly insurance redundant. Section 4.6 concludes.

This paper is related to several lines of literature. A first strand studies payment contract selection and trade finance in international trade. Schmidt-Eisenlohr (2013) was first to study the optimal choice between payment contracts in a strategic environment, however his analysis is constrained to a static setting. Antràs and Foley (2015) characterize the exporter's payment contract choice with a similar setup. Their analysis is motivated by stylized, empirical facts emerging from one specific export industry in the United States. They also offer a dynamic version of the model that gives predictions on how firms decide between cash in advance and open account payment contracts in the course of their export relationships. However, their analysis cannot account for the fact that the two types of payment contracts allow for very different learning possibilities about the importer's reliability. Furthermore, while their paper does not consider the role of bank finance in the dynamic context, my model incorporates this possibility in the exporter's choice set.

The role of banks and insurance firms in international trade finance has been the subject of several studies. For the case of export credit insurances that I consider

in my model, most of the existing work has an empirical focus. Van der Veer (2015) uses panel data from private trade credit insurers from 1992 to 2006 to show a positive effect of the existence of such insurances on exports. Auboin and Engemann (2014) use data on the subsequent period of the financial crisis to demonstrate that this positive effect remained stable over this time frame, i.e. it did not vary between crisis and non-crisis periods. Felbermayr and Yalcin (2013) study the effect of export credit guarantees on exports in Germany and find qualitatively comparable effects. These consistent findings are very much in line with Proposition 5 of this paper where I show that the exporter can benefit from insuring transactions to certain export destinations. Additionally, the result predicts that this positive effect might be particularly driven by early transactions within export relationships. Beyond export credit insurances, there exist several other types of bank-mediated payment contracts. It is frequently documented, that letters of credit and documentary collections are of particular importance and Schmidt-Eisenlohr and Niepmann (2016) provide an insightful empirical study on the usage of these contract forms. Olsen (2015) studies in a general equilibrium framework how banks' reputation for payment can be a valuable substitute for the trading partners' own reputations and weak legal institutions.³

Furthermore, the paper is related to a literature on learning and export dynamics in trade relationships. Araujo, Mion, and Ornelas (2016) study how contract enforcement and export experience shape firm export dynamics when information about importers is incomplete. The learning mechanics on the importer type under the open account contract in my model are inspired by their setup, however the authors do not use it to study questions related to trade finance. It is due to this learning that exporters increase export volumes over time, a pattern that the authors confirm empirically using a panel of Belgian exporters. Aeberhardt, Buono, and Fadinger (2014) provide comparable results for French exporters. Building on Araujo et al. (2016), Monarch and Schmidt-Eisenlohr (2016) document that there exists substantial heterogeneity in how export-relationships mature in different countries. This is consistent with my model as it predicts that country- and/or industry-specific characteristics (through the quality of legal institutions and the type distribution of importers) matter for the selection of payment contracts.

Finally, my model is related to a literature on self-enforcing relational contracts with incomplete information and adverse selection that were first studied by Levin (2003). More specifically, it adds to a growing literature on non-stationary relational contracts with adverse selection, in which contractual terms vary with the length of relationships. While in my paper learning about the importer induces switching between different payment contracts, previous work has studied non-stationarities in other contexts. Chassang (2010) studies how agents with conflicting interests

³A survey of this literature in the broader context of corporate finance in international trade can be found in Foley and Manova (2015).

can develop successful cooperation when details about cooperation are not common knowledge. Halac (2012) studies optimal relational contracts when the value of a principal-agent relationship is not commonly known and, also, how information revelation affects the dynamics of the relationship. Yang (2013) studies firm-internal wage dynamics when worker types are private information. Board and Meyer-ter Vehn (2015) analyze labor markets in which firms motivate their workers through relational contracts and study the effects of on-the-job search on employment contracts. Moreover, Defever, Fischer, and Suedekum (2016) study buyer-supplier relationships in which learning about the quality of a supplier can cause switches in the contractual nature of these relationships.

4.2 The model

The model builds on Araujo et al. (2016) and extends their setting to study the dynamic selection of payment contracts. It considers the problem of an exporter from country Home that markets a product in a foreign market (Foreign). The exporter cannot access foreign consumers directly and needs to contract with an importer in order to make products available to consumers. In Foreign, there is a continuum of agents with the ability to internally distribute goods produced by the exporter. Each exporter is a monopolist and has constant marginal costs c for the production of his output. The revenue from selling q units of the product in Foreign generates revenue R , which we assume to be a strictly increasing and concave function of q , i.e.

$$R = R(q), \text{ with } R'(q) > 0, R''(q) < 0, \text{ and } R(0) = 0. \quad (4.1)$$

Whether the concavity of the revenue function stems from technology, preferences or market structure is not important for the analysis below. The revenue is realized by the importer.

We model the export relationship as a repeated game, where in every period $t = 0, 1, 2, \dots$ an export transaction is performed. The exporter can engage in only one partnership at the same time. Figure 4.1 gives an intuitive summary of the stage game. In every period t , the exporter first decides whether to start a relationship with a new importer or to keep his previous one. He then proposes a contract to his importer specifying export volume, a payment from importer to exporter, and the point in time when this payment is made. Depending on the agreed payment timing, the exporter will receive the transfer either before he ships out his goods (Cash in Advance, $F_t = CIA$) or after the importer sells them in Foreign (Open Account, $F_t = OA$). The timing of the transfer is directly payoff relevant because shipping the goods from Home to Foreign takes time. We model this aspect by assuming that the goods can be sold in Foreign only at the beginning of the subsequent period $t+1$.

As the game shifts from one period to the next players discount their expected payoffs, potentially at different rates. Such an asymmetry may be interpreted as




Figure 4.1: Payment contract determines the stage game timing

interest rates and thus costs of capital that differ between Home and Foreign, but may also be the result of other institutional differences. Formally, we denote the exporter's discount factor as $\delta_E \in (0, 1)$. In order that payment contract selection is non-trivial we assume two types of importers that the exporting firm may encounter. Her type is the importer's private information. A share $\hat{\theta}$ of importers is myopic with discount factor $\delta_M = 0$, i.e. these importers only give positive value to payoffs that are realized in the current period. The remaining $1 - \hat{\theta}$ importers are patient with discount factor $\delta_I \in (0, 1)$. Absent additional information, the exporter holds the prior belief $\theta_0 = \hat{\theta}$ that a randomly chosen importer is myopic.

The quality of the contracting institutions in Home and Foreign come into the model through an exogenous country-specific probability that a contract is enforced when firms do not want to fulfill it voluntarily, a concept originating from Schmidt-Eisenlohr (2013). More specifically, in every period of the game a spot contract C_t , which is specified below, can be enforced in country Foreign (Home) with the i.i.d. probability λ (λ^*). In the model, λ is the probability that the importer is forced make the agreed-upon payment to the exporter for the period t transaction while λ^* is the probability that the exporter is being forced to produce and ship abroad the agreed-upon volume of exports.

In the following a formal representation of the stage game that introduces further notation and assumptions:

Stage game.

1. **Revenue realization.** The q_{t-1} units shipped in the previous period arrive in Foreign. The importer generates revenue $R(q_{t-1})$.
2. **Payment (if $F_{t-1} = OA$).** The importer decides whether to transfer $T_{t-1} = \beta R(q_{t-1})$, $0 < \beta < 1$, to the exporter. He (costlessly) finds an opportunity to only pay a transfer $T_{t-1} = \gamma R(q_{t-1})$, $0 \leq \gamma < \beta$, with probability $1 - \lambda$. Upon contract non-compliance the match is permanently dissolved.
3. **Matching decision.** Whenever unmatched, the exporter starts a new partnership. Otherwise, the exporter chooses whether to stick to the current importer or to re-match with a new partner.

4. Contracting.

- The exporter proposes a one-period spot contract $C_t = \{q_t, F_t, T_t\}$ to the importer. The contract specifies the units q_t of the exporter's product to be sold by the importer, and a trade finance arrangement F_t for period t , where $F_t \in \mathcal{F} = \{CIA, OA\}$. The choice of F_t determines the timing of the transfer payment $T_t = \beta R(q_t)$ from the importer to the exporter that is a fraction of the revenue. We make the restriction to spot contracts to reduce the screening possibilities in the model to a tractable minimum.
- The importer decides whether to accept or reject the contract. Upon rejection, the match is permanently dissolved and a new match is formed in the following period.

5. **Payment (if $F_t = CIA$).** The importer decides whether to transfer T_t to the exporter. Upon non-payment the match is permanently dissolved.

6. **Production and Shipment.** The exporter decides whether to produce and ship q_t units of his goods to the importer, as specified in the contract. He (costlessly) finds an opportunity to deviate from the contracted terms with probability $1 - \lambda^*$. Upon non-shipment of the agreed quantity the match is permanently dissolved.

4.3 The nature of payment contracts

Before discussing how a payment contract is selected in the presence of several alternatives and how this choice varies over time, this section studies each of the two alternatives in isolation. We explore for every payment format separately, how information about the importer is revealed and how the export intensity changes over time. The payment type selection in the dynamic context of the model is discussed in the subsequent section.

In the following, we derive equilibria for both cases, CIA and OA, that are highly comparable but highlight important differences resulting from the nature of the two payment contracts.

4.3.1 Cash in advance

In this section we study the case where the exporter is restricted to cash in advance, i.e. $F_t = CIA$ for all t . Consider the following strategy profile that we will show to be part of a sequential equilibrium of the repeated game.

CIA strategy profile.

- **Exporter strategy:** The exporter forms a new partnership whenever unmatched. He terminates an existing partnership if and only if the importer defaults on the contract. The exporter proposes a contract C_t that maximizes his current period expected payoffs.
- **Importer strategy:** A myopic importer will reject any CIA contract. A patient importer never defaults from the proposed contract.

The importers' participation constraints for period t are:

$$(\delta_i - \beta)R(q_t) \geq 0, \quad \text{where } i = M, I. \quad (\text{PC}_i^A)$$

The constraints state, that tomorrow's revenue $R(q_t)$ attained by the sale of today's exports q_t must be larger than the the share β of revenue that the importer has to transfer to the exporter before shipping. Because the goods arrive in Foreign only in the following period revenue has to be discounted. Observe that because $\delta_M = 0$, (PC_M^A) cannot be fulfilled for any $\beta \in (0, 1)$. Hence the myopic importer will never accept any CIA contract. The exporter will therefore offer a separating contract to the patient importer and set β to extract all rents from her.

With this in mind, suppose that we are in the initial period $k = 0$ of a newly matched export relationship. At the contracting stage, the exporter will set the export quantity q_0 and a transfer payment T_0 that maximizes his spot payoffs subject to the participation constraint of the patient importer:

$$\max_{q_0} (1 - \theta_0)(\beta R(q_0) - cq_0) \quad \text{s.t.} \quad (\delta_I - \beta)R(q_0) \geq 0.$$

The exporter chooses the exported quantity q_0 such that it maximizes his share of the revenue $\beta R(q_0)$ minus the costs of production, and subject to (PC_I^A) . The payoff is weighted by the probability of facing a patient importer $1 - \theta_0$ since non-participation of the myopic type leads to zero payoffs. The exporter sets $\beta^A = \delta_I$ to make the patient importer indifferent between accepting and rejecting the contract and to receive the maximally feasible share of the revenue. This simplifies the maximization problem to:

$$\max_{q_0} (1 - \theta_0)(\delta_I R(q_0) - cq_0).$$

The producer's optimal export quantity under cash in advance, Q^A , is determined by the related first-order condition (FOC):

$$R'(q_0) = \frac{c}{\delta_I}.$$

Once an initial CIA export transaction is successful the exporter knows for sure that he is matched with a patient type. We can write the belief that the importer is myopic given k previous successful transactions under cash in advance as:

$$\theta_k^A = \begin{cases} \theta_0 & \text{if } k = 0, \\ 0 & \text{if } k > 0. \end{cases}$$

Because the FOC for the initial transaction is already belief-independent it follows that Q^A is the same for every cash in advance transaction. The expected payoffs in the initial period of a CIA relationship thus are:

$$\pi_0^A \equiv \pi^A(\theta_0^A) = (1 - \theta_0)(\delta_I R(Q^A) - cQ^A).$$

In any subsequent transaction the exporter expects to receive the following payoffs:

$$\bar{\pi}^A \equiv \pi^A(\theta_k^A = 0) = \delta_I R(Q^A) - cQ^A. \quad (4.2)$$

Because the production and shipment decision is made after the transfer payment under a CIA contract, we need to ensure that the exporter does not have an incentive to deviate and not ship out his goods as agreed upon. The following Lemma gives a simple condition that rules out any such deviation.

Lemma 1. *For every $\theta_0 \in (\frac{cQ^A}{\bar{\pi}^A}, 1)$ there exists a unique $\underline{\delta}_E = \frac{cQ^A}{\theta_0 \bar{\pi}^A}$ such that the exporter never deviates at the Production and Shipment stage if and only if $\delta_E \geq \underline{\delta}_E$.*

Proof. See Appendix.

Lemma 1 shows that if the mass of myopic types in the population of potential importers is not too small, it is always optimal for sufficiently patient exporters (those with $\delta_E \geq \underline{\delta}_E$) to ship the contracted goods to the importer who paid for them beforehand. If however the mass of myopic types is too small breach of contract by the exporter cannot be avoided because he finds it too likely to draw again a patient type in the following period where, again, he could cash the transfer, refrain from shipping and re-match.

The following Proposition summarizes our findings on the cash in advance equilibrium.

Proposition 1. *Suppose only cash in advance payments are possible. The exporter starts a partnership whenever he finds a match, maintains the partnership as long as he does not observe a default, and exports Q^A in each period k in which the partnership is active. A myopic importer never participates under cash in advance. A patient importer never defaults and never terminates a partnership. This strategy profile together with the belief updating rule θ_k^A is a sequential equilibrium.*

There are several points noteworthy about the equilibrium. First, a cash in advance contract is very demanding with respect to the financial capabilities of the importer. This aspect is stressed in our model by the fact that impatient importers are fully myopic which precludes them from accepting the contract. Therefore any CIA contract in the present model is a separating contract that perfectly screens out patient importers. As a consequence, learning about the importer's type under CIA is immediate. After the initial period of interaction, the exporter knows with certainty the type of his current match.

Certainly, assuming full myopia of the impatient type is a strong assumption and immediate separation of types the extreme consequence. However, the result is very illustrative when contrasted with the screening mechanism available under open account introduced in the following section.

4.3.2 Open account

We now study the case where the exporter is restricted to open account contracts, i.e. $F_t = OA$ for all t . Essentially, this version of our model is a stripped-down variant of the setup by Araujo et al. (2016). Consider the following strategy profile that we will show to be part of a sequential equilibrium of the repeated game.

OA strategy profile.

- **Exporter strategy:** The exporter forms a new partnership whenever unmatched. He terminates an existing partnership if and only if the importer defaults on the contract. The exporter proposes a contract C_t that maximizes his current period expected payoffs.
- **Importer strategy:** Both, myopic and patient importer accept the OA contract. While the patient importer never defaults, the myopic importer will default whenever she finds the opportunity to do so.

With this strategy in mind we can write the participation constraints of the two types as:

$$\begin{aligned} (1 - \beta)R(q_t) &\geq 0, & (\text{PC}_I^\Omega) \\ (1 - \lambda\beta)R(q_t) &\geq 0, & (\text{PC}_M^\Omega) \end{aligned}$$

where (PC_I^Ω) is the participation constraint of the patient importer and (PC_M^Ω) that of the myopic importer, respectively. While it is possible to construct a separating contract that picks up only patient importers under CIA, this is not possible for OA. Myopic importers only make a transfer when exporters can enforce their contract (which happens with probability λ). This makes their PC more lenient. Because $0 < \beta < 1$, it follows that every feasible OA transaction involves pooling of importer

types. Besides, discounting does not affect the importer's participation decision since revenue realization and payment for a period t contract are both made in $t + 1$.

Suppose for the moment that importers behave as prescribed by the strategy profile and consider belief formation. In the initial period of a partnership, $k = 0$, the exporter believes that he is matched with a myopic type with probability $\theta_0 = \hat{\theta}$. According to the strategy profile patient agents never deviate and myopic types always want to deviate but in every period this is only possible if contracts cannot be enforced. This happens with probability $1 - \lambda$. Hence, if no deviation occurs in the course of the first k transactions, the belief of facing a myopic type in period k is updated by Bayes' rule as follows:

$$\theta_k^\Omega = \frac{\hat{\theta}\lambda^k}{1 - \hat{\theta}(1 - \lambda^k)}.$$

The probability of payment in the initial period can be written as $\Lambda_0 = 1 - \theta_0 + \lambda\theta_0$. More generally, the probability of payment in period k is $\Lambda_k = \frac{1 - \hat{\theta}(1 - \lambda^{k+1})}{1 - \hat{\theta}(1 - \lambda^k)}$.

Observe that in order to make the patient importer behave as described by the strategy profile, it is not enough to merely consider her participation constraint as we did in the CIA case. Under open account, she must be granted a payoff such that she does not seize the revenue realized by her instead of making the transfer payment. We denote the maximal share of revenue that the importer is able to seize by $1 - \gamma$, where $0 \leq \gamma < \beta$ is the exogenous revenue share that the exporter can still claim under any importer deviation. We can formulate the following Lemma:

Lemma 2. *There exists a unique $\delta_I = \frac{\beta - \gamma}{1 - \gamma}$ such that the importer never deviates at the Payment stage of any period of an export relationship if and only if $\delta_I \geq \delta_I$. Equivalently, a patient importer will never deviate from the contract if and only if*

$$\beta \leq \delta_I + \gamma(1 - \delta_I) \equiv \beta^\Omega. \quad (\text{IC}_I^\Omega)$$

Proof. See Appendix.

The Lemma states, that as long as the share of revenue that has to be transferred to the exporter is below the threshold level β^Ω , the patient importer has no incentive to deviate from the accepted contract. Obviously, payment is never incentive compatible for a myopic importer who will refrain from transferring β^Ω whenever possible.

Let us now turn to the exporter's maximization problem. In any period k of the export relationship he chooses q_k to maximize:

$$\max_{q_k} \delta_E \Lambda_k \beta R(q_k) - cq_k \quad \text{s.t.} \quad \beta \leq \beta^\Omega.$$

While the exporter has to bear the costs of production cq_k already today, he will receive the expected transfer $\Lambda_k \beta R(q_k)$ only in the following period which is therefore

discounted by δ_E . The exporter wants to extract the maximum possible transfer from the importer and he will set β such that (IC_I^Ω) binds with equality, i.e. $\beta = \beta^\Omega$. The unconstrained maximization problem of the exporter for transaction k of an export relationship thus is:

$$\max_{q_k} \delta_E \Lambda_k \beta^\Omega R(q_k) - cq_k$$

The FOC to this problem can be written as:

$$R'(q_k) = \frac{c}{\delta_E \Lambda_k \beta^\Omega},$$

implying that the producer's optimal export quantity under OA, Q_k^Ω , depends on his belief θ_k and is increasing in k (since the probability of payment Λ_k is increasing in k). We can write the expected exporter payoff from an open account transaction in period k of the export relationship as:

$$\pi_k^\Omega \equiv \pi^\Omega(\theta_k^\Omega) = \delta_E \Lambda_k \beta^\Omega R(Q_k^\Omega) - cQ_k^\Omega.$$

Observe that with the export relationship growing mature, the probability of payment Λ_k converges to one. The exported quantity therefore also converges to a limit value that we denote by Q^Ω . We denote the payoffs that the exporter receives at this limit as:

$$\bar{\pi}^\Omega \equiv \pi^\Omega(\theta_k^\Omega = 0) = \delta_E \beta^\Omega R(Q^\Omega) - cQ^\Omega. \quad (4.3)$$

The following Proposition summarizes our findings on the open account equilibrium.

Proposition 2. *Suppose only open account payments are possible. The exporter starts a partnership whenever he finds a match, maintains the partnership as long as he does not observe a default, and exports Q_k^Ω in the k -th period of any active partnership. A myopic importer deviates from the contract whenever she has the opportunity. A patient importer never defaults. Both types never terminate a partnership. This strategy profile together with the belief updating rule θ_k^Ω is a sequential equilibrium.*

4.3.3 Comparing the payment contract equilibria

There are several noteworthy differences between export relationships with CIA and OA payment contracts. First, while learning about the importer's type under CIA is immediate (see updating rule θ_k^A) information acquisition under OA is only gradual (see θ_k^Ω). The reason is that while the exporter can design the CIA contract to only attract patient importers, OA necessarily involves pooling which makes it necessary for patient importers to build up a reputation for reliability.

Second, since under CIA there is no risk involved in any transaction for the exporter (since payment is made before production) he will export the ex-post optimal

quantity Q^A starting with the very first transaction. In contrast, under OA the export decision is risky because payment takes place after production. Additionally, since the exporter can only use pooling contracts the optimal export quantity Q_k^Ω is adjusted by the importer's reputation of being a patient type and therefore increases gradually over time (up to Q^Ω , for $k \rightarrow \infty$).

The first two points together imply the following for the exporter's expected payoff. Given that he is matched to a patient type, under CIA the payoff will after the initial period immediately jump from π_0^A to the maximum $\bar{\pi}^A$. Under OA, it will start from π_0^Ω increase gradually in the course of the relationship up to the maximum $\bar{\pi}^\Omega$.

4.4 Dynamic selection of payment contracts

In this section we relax the restriction to one specific payment contract and allow the exporter to freely choose in every period k of the export relationship whether to employ CIA or OA. At the Contracting stage of every period t the exporter thus decides on an element $F_t \in \mathcal{F} = \{CIA, OA\}$. Essentially, we explore how the exporter decides between the CIA and OA equilibrium derived in Propositions 1 and 2 and how this choice varies over time, as beliefs about the matched importer evolve.

For both payment formats the expected exporter stage game payoff is highest with all uncertainty about the importer's patience being resolved, i.e. when the exporter is certain to be facing a patient type. This is the case when $\theta_k^A = 0$ and $\theta_k^\Omega = 0$, respectively. Comparison of the related payoffs in equations (4.2) and (4.3) suggests that $\bar{\pi}^\Omega \geq \bar{\pi}^A$ whenever $\beta^\Omega \delta_E \geq \delta_I$, and $\bar{\pi}^\Omega \leq \bar{\pi}^A$ otherwise. In the main text, we focus on the scenario where under complete information OA is more profitable than CIA, i.e. the case where $\beta^\Omega \delta_E > \delta_I$. We thus capture the scenario where $\delta_E > \delta_I$, i.e. the exporter has relatively low costs of capital while the importer faces relatively high costs of capital. From a modeling perspective, this case is particularly interesting because while OA is not learning-optimal for the exporter it gives him the largest long-run payoffs. In the Appendix, we also briefly discuss the reverse situation.

4.4.1 Main result

The comparison of the two payment contracts in Section 4.3.3 hints at the latent trade-off in the design of the dynamically optimal series of spot contracts between acquiring information about the importer and using the payment contract that promises the largest spot payoff. In this section we study how the exporter should optimally design the series of spot export contracts, i.e. how he should over time select between the two possible payment contracts in the form of the two equilibria

derived in Propositions 1 and 2. To keep the analysis tractable, we assume that the exporter has to decide for one specific payment contract and cannot propose a whole menu of contracts to the importer.

In order to understand payment contract selection it is important to observe that i) open account is maximizing the exporter's payoff under full information, i.e. $\bar{\pi}^\Omega > \bar{\pi}^A$, ii) for each payment contract the exporter's payoff is maximal under full information, and iii) learning the importer's type under CIA is immediate. The three points together imply that if a CIA transaction is successfully conducted any further transaction with this importer will be on OA and involve export quantity Q^Ω and exporter payoffs $\bar{\pi}^\Omega$.

With this in mind, let us assume that the importer's type has not yet been learned, i.e. OA has been played up to period k of the relationship. In choosing the transaction's payment contract the exporter has to decide whether he wants to learn the true type of the importer today (through switching to CIA for the current transaction) or whether to continue with the OA format, at least until the following period $k+1$ in which he can reconsider. Formally, the exporter will conduct the k th transaction through open account if and only if

$$\pi_k^\Omega + \delta_E \pi^A(\theta_{k+1}^\Omega) + \frac{\delta_E^2}{1 - \delta_E} \bar{\pi}^\Omega \geq \pi^A(\theta_k^\Omega) + \frac{\delta_E}{1 - \delta_E} \bar{\pi}^\Omega, \quad (4.4)$$

where $\pi^A(\theta_k^\Omega)$ and $\pi^A(\theta_{k+1}^\Omega)$ denote the exporter's expected payoffs under CIA for beliefs θ_k^Ω and θ_{k+1}^Ω , respectively, that are derived from the open account updating rule. Using this expression we obtain the following Proposition.

Proposition 3. *Suppose that $\delta_E > \frac{\delta_I}{\beta^\Omega}$, i.e. OA is payoff-maximal with perfect information about the importer's type. Also suppose that $\theta_0 \geq \frac{cQ^A}{\delta_E \bar{\pi}^A} \equiv \underline{\theta}_0$, i.e. Lemma 1 holds. Consider the k th transaction in an export relationship. The exporter will conduct this transaction through Open Account if and only if*

$$\pi_k^\Omega - \pi^A(\theta_k^\Omega) \geq \delta_E [\bar{\pi}^\Omega - \pi^A(\theta_{k+1}^\Omega)] \quad (4.5)$$

and ask for Cash in Advance otherwise.

The inequality in (4.5) is obtained from rearranging (4.4) and provides a necessary and sufficient condition whether any transaction of an export relationship is conducted through either CIA or OA. Intuitively, it states that OA will be chosen for the current transaction if today's payoff gain from choosing OA instead of CIA is larger than tomorrow's payoff gain from OA when using the more screening-efficient CIA contract today. Note, that expression (4.5) is always fulfilled when the importer's type has been revealed in the past, e.g. when a CIA has been conducted. I provide a more extensive discussion of the result in the following subsection.

4.4.2 Discussion

The results up to now are formulated fairly generally and to make their discussion illustrative we make an additional but unrestrictive assumption for the revenue function. We let $R(q) = \frac{q^{1-\alpha}}{1-\alpha}$, where the parameter α , with $0 < \alpha < 1$, determines the concavity level of the revenue function. Note that this function captures all the properties that we had assumed in (4.1) and is general enough to capture arbitrary levels of concavity. This functional form allows us to obtain simple expressions for many of the previously derived results. Importantly for the discussion of Proposition 3, $\pi_k^\Omega = (\Lambda'_k)^{\frac{1}{\alpha}} \bar{\pi}^\Omega$ and $\bar{\pi}^\Omega = \left(\frac{\delta_E \beta^\Omega}{\delta_I} \right)^{\frac{1}{\alpha}} \bar{\pi}^A$. With $\Lambda'_k = 1 - \theta'_k(1 - \lambda)$ we denote the probability of payment for some belief θ'_k . The expressions allow us to rewrite (4.5) as:

$$\mathbf{I}(\theta'_k, \lambda, \delta_I, \delta_E, \alpha, \gamma) \equiv \Lambda'_k{}^{\frac{1}{\alpha}} - \delta_E + \left(\frac{\delta_I}{\delta_E \beta^\Omega} \right)^{\frac{1}{\alpha}} \frac{(\delta_E - \Lambda'_k)(1 - \theta'_k)}{\Lambda'_k} \geq 0, \quad (4.6)$$

which determines whether the k th transaction of an export relationship is conducted via OA (if $\mathbf{I} > 0$) or via CIA (if $\mathbf{I} < 0$). Observe that in $\mathbf{I}(\cdot)$ the variables $\lambda, \delta_I, \delta_E, \alpha$ and γ are all constants for a given export relationship. Within a relationship the value of $\mathbf{I}(\cdot)$ only changes with the belief θ'_k of facing a myopic importer. This belief can be characterized as:

$$\theta'_k = \begin{cases} \theta_k^\Omega & \text{if } \mathbf{I} > 0 \text{ for all previous transactions, or when } k = 0, \\ 0 & \text{if } \mathbf{I} < 0 \text{ in some previous transaction.} \end{cases} \quad (4.7)$$

Note also, that the expression for the minimal initial belief of facing a myopic type from Lemma 1 that is necessary for equilibrium existence simplifies to $\underline{\theta}_0 = \frac{1-\alpha}{\delta_E \alpha}$.

The predictions of Proposition 3 in the form of expressions (4.6) and (4.7) are best explained by making use of their graphical representations in Figure 4.2. The sub-figures (a) and (b) show for all possible combinations of initial belief θ_0 and contract enforceability in Foreign λ which payment contract will be used in the initial and (conditional on continuation) in the subsequent period of the export relationship, respectively.

First, consider the left panel of Figure 4.2 where agents are in their initial period of interaction, i.e. $k = 0$. The increasing and concave line represents all those points where the exporter is indifferent between the use of a CIA and an OA payment contract, i.e. $\mathbf{I} = 0$. For all parameter combinations in the blue-shaded area the exporter prefers CIA in the initial period, i.e. $\mathbf{I} < 0$, while OA is preferred above the indifference line, i.e. $\mathbf{I} > 0$. Note that for initial beliefs left from the dashed line at $\underline{\theta}_0$ we cannot make predictions since these are not covered by the CIA equilibrium.

The plot predicts, that exports to countries with higher contract enforceability are more likely to start out with OA from the beginning compared to countries with lower contract enforceability. Furthermore, the higher the share of unreliable

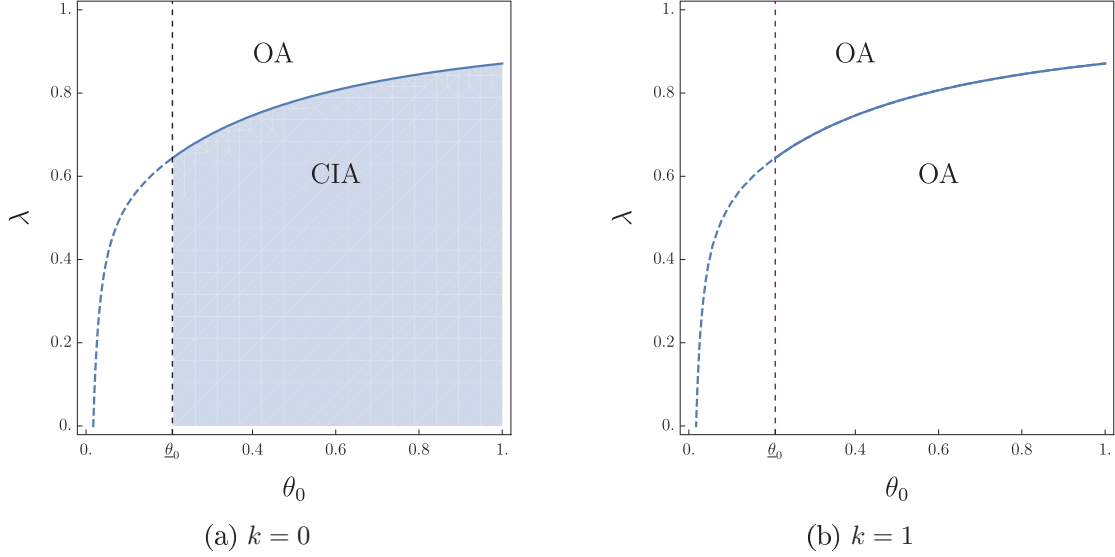


Figure 4.2: Payment contract choice over time (parameterization: $\delta_I = .6$, $\delta_E = .85$, $\alpha = .85$, $\gamma = .3$)

importers in the population of importers the more likely it is that CIA is used initially. These predictions are implied by the increasing, concave form of the $\mathbf{I} = 0$ -contour line. I have conducted this plotting exercise for a large variety of model parametrizations and this shape was robust over all of them. Analytical proof seems impossible since the function is only implicit and highly non-linear. I am however confident that this pattern generalizes and the following discussions are done under this conjecture.

Next, consider the right panel of Figure 4.2, where agents are in the second period of interaction, i.e. $k = 1$. The updating rule θ'_k prescribes that exporters who played CIA initially are, given continuation, sure in $k = 1$ to be facing a patient type and choose OA because under full information $\bar{\pi}^\Omega > \bar{\pi}^A$. In addition, because of the concave and increasing shape of the payment contract indifference line all exporters who played OA already in the initial period continue playing OA. This is because their belief is updated to $\theta'_1 = \theta_1^\Omega$ which represents a horizontal shift to the left in the Figure.

In sum, Proposition 3 predicts that if the exporter faces lower costs of capital than the importer CIA contracts can act as a useful tool in the initial phase of export relationships. This is particularly pronounced so when the quality of legal institutions in the destination country is low and the probability of facing unreliable importers is high. Once an importer has build up a reputation for being reliable OA contracts will be used across all destination markets. In the following section, we

will study how this result is enriched when exporters have the possibility to ensure the payment in their open account transactions.

4.5 Trade credit insurance

With a trade credit insurance (CI) the exporter can rule out the risk of importer non-payment in an open account transaction. For its service, the insurer charges a per-transaction fee that we denote by F_k . We assume a perfectly competitive insurance market and that the fee can be separated into a fixed and a variable component. It is given by

$$F_k = m + \delta_E(1 - \Lambda_k^{CI})T_k, \quad (4.8)$$

where the fixed (and time-invariant) component m covers setup and monitoring costs that the bank incurs for insuring the transaction. The second addend represents the variable component that depends on the size of the insured transfer in period k , T_k , which is weighted by the probability of nonpayment $1 - \Lambda_k^{CI}$, where Λ_k^{CI} denotes the probability of payment in the k th period under insurance.⁴ Finally, because potential payment default occurs only in $t + 1$ the variable component is discounted by one period.⁵

Because the insurer has a vital interest that the importer does not default it will engage in importer screening itself before granting a credit insurance.⁶ We model this aspect by assuming that starting an export relationship with a trade credit insurance decreases the proportion of myopic types in the population to $\hat{\theta}^{CI} = \phi\hat{\theta}$, where $\phi \in (0, 1)$ is an inverse measure of the insurer's ability to screen out myopic types.

The remainder of the section is structured as follows. In Section 4.5.1 we study export relationships when the exporter can only choose insured open account as payment contract. The results bear much resemblance with the open account case in Section 4.3.2. In part 4.5.2 we then study how trade credit insurance interacts with the non-intermediated payment contracts. It thus complements the analysis from Section 4.4. Finally, in Section 4.5.3 we conduct comparative static exercises.

⁴The formalization of the insurance fee is inspired by Schmidt-Eisenlohr and Niepmann (2016), in particular by their model of the letter of credit contract. However, they do not study this contract in a dynamic setting.

⁵For the discounting we use the exporter's discount factor. It would be more appropriate to use the insurer's own time preference rate here. Because this difference is not of central importance for the study of the exporter's insurance decision we abstract from this complicating detail here.

⁶This assumption is endorsed by the fact that credit insurers such as Euler Hermes and AIG advertise their insurance services with their expertise in monitoring the reliability of transaction counterparts.

4.5.1 Insured payment contracts in isolation

We now study the export relationship when restricted to credit insured payment contracts. In order to achieve comparability of results we derive an equilibrium that is based on the same strategy profile as in the open account scenario in Section 4.3.2.

The participation constraints of the importers are the same as in the open account scenario. Also, the payment incentive constraint for the patient importer is the same as under open account. The exporter therefore asks the importer to transfer the revenue share $T_k = \beta^\Omega R(q_k)$. The exporter thus maximizes in period t :

$$\max_{q_t} \delta_E \beta^\Omega R(q_t) - cq_t - F_t, \quad (4.9)$$

which by plugging in F_t can be rewritten as

$$\max_{q_t} \delta_E \Lambda_k^{CI} \beta^\Omega R(q_t) - cq_t - m. \quad (4.10)$$

Observe that even though the CI eliminates the risk of non-payment the probability of payment Λ_k^{CI} still indirectly affects the exporter's maximization problem through the LC-fee. This can be nicely seen when rewriting (4.9) to (4.10). For the calculation of Λ_k^{CI} the same mechanics as for the calculation of Λ_k under open account apply. However, as described above the exporter beliefs to initially face a myopic type with smaller probability $\hat{\theta}^{CI}$ and the belief of facing a myopic type in period k is determined via Bayes' rule as:

$$\theta_k^{CI} = \frac{\hat{\theta}^{CI} \lambda^k}{1 - \hat{\theta}^{CI} (1 - \lambda^k)}.$$

Consequently, the probability of payment in the k th relationship period under CI is $\Lambda_k^{CI} = \frac{1 - \hat{\theta}^{CI} (1 - \lambda^{k+1})}{1 - \hat{\theta}^{CI} (1 - \lambda^k)}$. With this in mind we can solve the maximization problem in (4.10). The FOC is:

$$R'(q_t) = \frac{c}{\delta_E \Lambda_k^{CI} \beta^\Omega},$$

and we denote the resulting, optimal export quantity as Q_k^{CI} . The payoff in a CI relationship in the k th transaction then is

$$\pi_k^{CI} \equiv \pi^{CI}(\theta_k^{CI}) = \delta_E \Lambda_k^{CI} \beta^\Omega R(Q_k^{CI}) - cQ_k^{CI} - m. \quad (4.11)$$

The following Proposition summarizes our findings on the letter of credit equilibrium.

Proposition 4. *Suppose only credit insured open account payment contracts are possible. The exporter starts a partnership whenever he finds a match, maintains the partnership as long as he does not observe a default, and exports Q_k^{CI} in the*

k -th period of any active partnership. A myopic importer deviates from the contract whenever she has the opportunity. A patient importer never defaults. Both types never terminate a partnership. This strategy profile together with the belief updating rule θ_k^{CI} is a sequential equilibrium.

When all uncertainty about the importer's type is resolved, equation (4.11) can be written as

$$\bar{\pi}^{CI} \equiv \pi^{CI}(\theta_k^{CI} = 0) = \delta_E \beta^\Omega R(Q^\Omega) - cQ^\Omega - m = \bar{\pi}^\Omega - m$$

which shows that $\bar{\pi}^\Omega > \bar{\pi}^{CI}$. This gives the following Corollary.

Corollary 1. *Under complete information, insured open account payment contracts are strictly dominated by non-insured open account contracts.*

The corollary hints at a result that we will derive in the following when studying the interaction of credit insurances with the other contract forms: Once the exporter attains enough information about the importer's type it will be better for him to conduct his transactions through open account instead of insuring it.

4.5.2 When to employ trade credit insurance

In this section, we study the payment contract decision when the choice set now also includes the possibility to insure the OA transaction. We denote the extended choice set by $\mathcal{F}^+ \equiv \{CIA, OA, CI\}$. Because of the independence axiom from expected utility theory we can analyze the payment contract decision in two separate cases: i) We study under which conditions CI is preferred to OA (and vice versa) *given that* OA is preferred to CIA, i.e. expression (4.5) holds. ii) We study under which conditions CI is preferred to CIA (and vice versa) *given that* CIA is preferred to OA, i.e. expression (4.5) does not hold.

First, consider case i) and let us determine under which conditions we will choose CI instead of OA. The respective belief updating rules θ_k^{CI} and θ_k^Ω are very similar. Their only difference is that the prior probability of being matched with a myopic type under CI is downsized by the factor ϕ . Note that we can write the belief under OA for period $k + 1$ as $\theta_{k+1}^\Omega = \frac{\theta_k^\Omega \lambda}{1 - \theta_k^\Omega (1 - \lambda)}$ and that this is an increasing and strictly convex function of θ_k^Ω . As a consequence, given (4.5), the incentives to employ a credit insured contract are always largest in the initial period of interaction because the exporter's learning effect through the screening activity of the insurer is largest. Hence, whenever CI is employed instead of OA it will for sure be employed in the initial transaction. Note also, that since the learning about the importer through insurer screening is permanent, the exporter will not employ CI for more than the initial period.⁷ Consequently, the exporter decides for a credit insurance in an export

⁷This is implied by our assumption of only two types of importers and that their type is constant.

relationship if and only if the payoffs from using a CI in the initial period and OA forever after are larger than the payoffs obtained from using OA in every period. Formally,

$$\pi_0^{CI} + \sum_{k=1}^{\infty} \delta^k \pi^{\Omega}(\theta_k^{CI}) > \sum_{k=0}^{\infty} \delta^k \pi^{\Omega}(\theta_k^{\Omega}).$$

This expression can be rearranged to:

$$\sum_{k=0}^{\infty} \delta^k [\pi^{\Omega}(\theta_k^{CI}) - \pi^{\Omega}(\theta_k^{\Omega})] > m, \quad (4.12)$$

which intuitively means that CI is employed (and only so in the initial period) if the continued benefits from importer screening through the insurer are larger than the fixed costs of the one-period insurance contract m .

Next, consider case ii) and let us determine under which conditions we will choose CI instead of CIA. The problem simplifies by recognizing that after one period of CI it is dominated by OA. In addition, switching from CIA to CI is never profitable because the exporter will have complete knowledge of the importer type after the first CIA transaction (which leads to OA by Corollary 1). Hence, if the exporter uses CI in the present scenario he will do it only in the initial transaction. CI is preferred to CIA in the initial period if and only if:

$$\pi_0^{CI} + \delta_E \pi^A(\theta_1^{CI}) + \frac{\delta_E^2}{1 - \delta_E} \bar{\pi}^{\Omega} \geq \pi^A(\theta_0^{\Omega}) + \frac{\delta_E}{1 - \delta_E} \bar{\pi}^{\Omega}.$$

The equation can be simplified to:

$$\pi_0^{CI} - \pi^A(\theta_0^{\Omega}) > \delta_E [\bar{\pi}^{\Omega} - \pi^A(\theta_1^{CI})] \quad (4.13)$$

The following Proposition gives a complete summary for every period and beliefs in the export relationship which payment contract will be used:

Proposition 5. *An exporter will consider to conduct a transaction through credit insurance only if in the initial period of an export relationship. This transaction will be conducted with CI if and only if either of the following is true:*

- a) (4.5) holds and (4.12) holds, or
- b) (4.5) does not hold and (4.13) holds.

All later periods of the export relationship are captured entirely through expression (4.5).

Because it is more illustrative, in analogy to Section 4.4.2, we conduct the discussion of the results for the case of a parameterized revenue function. The following

Corollary restates the differences in Proposition 5 when using the revenue function from before.

Corollary 2. *Suppose that $R(q) = \frac{q^{1-\alpha}}{1-\alpha}$. Then the exporter conducts the initial transaction with CI if and only if either of the following is true:*

- a) $\mathbf{I} > 0$ holds and $\mathbf{J} \equiv \sum_{k=0}^{\infty} \delta^k \left[(\Lambda_k^{CI})^{\frac{1}{\alpha}} - (\Lambda_k)^{\frac{1}{\alpha}} \right] - \kappa > 0$, or
- b) $\mathbf{I} < 0$ and $\mathbf{K} \equiv (\Lambda_0^{CI})^{\frac{1}{\alpha}} - \delta_E - \kappa + \left(\frac{\delta_I}{\delta_E \beta^\Omega} \right)^{\frac{1}{\alpha}} [\delta_E(1 - \theta_1^{CI}) - (1 - \theta_0^\Omega)] > 0$ holds.⁸

The implications of the Corollary are best studied the graphs in Figure 4.3. For the same parameter values that we also use in Figure 4.2, sub-figures (a) and (b) show for all possible combinations of initial belief θ_0 and contract enforceability in Foreign λ which payment contract will be used in the initial and (conditional on continuation) in the second period of the export relationship, respectively.

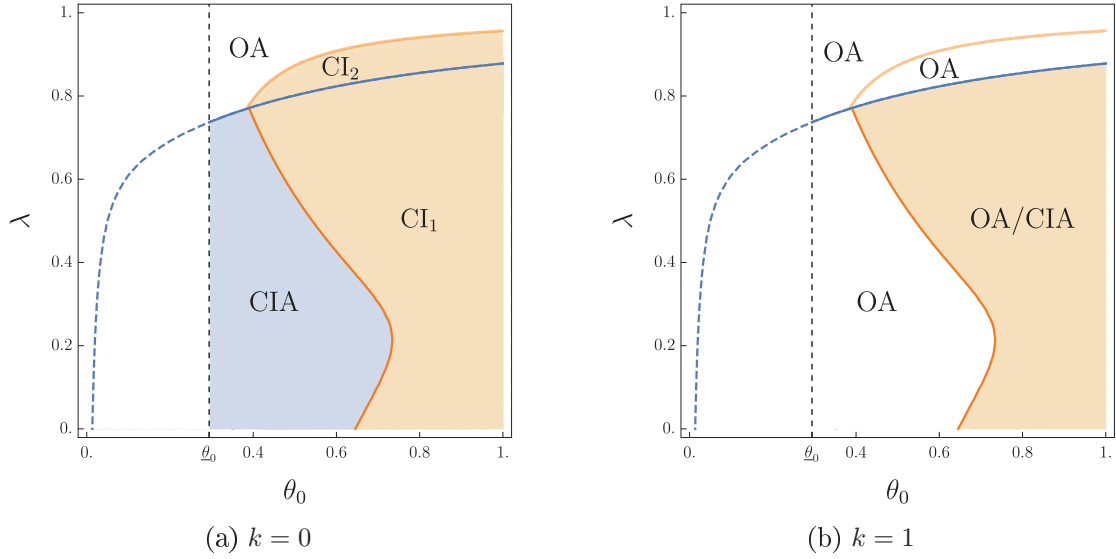


Figure 4.3: Payment contract choice over time (parameterization: $\delta_I = .6$, $\delta_E = .85$, $\alpha = .8$, $\gamma = .3$, $\phi = .65$, $\kappa = .14$)

First, consider the left panel in Figure 4.3 where agents are in the initial period of interaction. As in Figure 4.2 the blue line represents the contour line where $\mathbf{I} = 0$ and the exporter is indifferent between OA and CIA. For the area above this line,

⁸Without loss of generality we define and use $m \equiv \kappa \bar{\pi}^\Omega$, $\kappa > 0$, since $\bar{\pi}^\Omega = (\delta_E \beta^\Omega)^{\frac{1}{\alpha}} c^{\frac{\alpha-1}{\alpha}} \frac{\alpha}{1-\alpha}$ is a constant.

part a) of the Corollary applies and if $\mathbf{J} > 0$ CI will be chosen (the orange area denoted by CI_2 in the graph), while OA is preferred otherwise.⁹ On the other hand, consider the area below the blue line where part b) of the Corollary applies. Then CI will be chosen whenever $\mathbf{K} > 0$ holds (the orange area denoted by CI_1 in the graph) and CIA otherwise. Note from the graph that exporters' credit insurance usage is particularly pronounced for an intermediate range of contract enforceability. This pattern emerges for many parameterizations and how it changes is discussed in the following comparative statics section in more detail.

Next, consider the right panel of Figure 4.3 where agents are in the second period of interaction. For those parameters where CIA was played initially, by the same logic as in Section 4.4.2, OA will be played for sure. For those parameter combinations resulting in the area CI_2 in the left panel we will surely move into OA in $k = 1$ since for the previously discussed reasons no further value can be attained from continuing to insure the transaction. The behavior of exporters starting a relationship in area CI_1 is more intricate. Surely, they will not engage in CI anymore for the same reasons as in the first case. However, depending on the belief θ_1^{CI} , they might either conduct the transaction through OA or CIA. If OA is employed the exporter will continue to use OA for any subsequent period since the $(\mathbf{I} = 0)$ -contour line is increasing and concave. If CIA is used then the exporter will switch to OA from $k = 2$ onwards. Generally, we can conclude that for any combination of parameters a transaction will be conducted via open account for $k \geq 2$.

To sum up, Proposition 5 predicts that trade credit insurances can act as a valuable tool in the initial phase of export relationships when the importer's expected reliability is low. The graphs in Figure 4.3 (as well as the robustness-checks in the following subsection) suggest that the usage of credit insurances is particularly pronounced for intermediate qualities of contracting institutions. An intuitive rationalization for this pattern is that credit insurance is particularly attractive when, on the one side, insurance is not too expensive. It could be too costly because of the high risk of expropriation at low values of λ . On the other side, its value is potentially marginalized when facing strong legal institutions at high λ . Empirically, this pattern finds traction through the results by Schmidt-Eisenlohr and Niepmann (2016) who report for the U.S. that letter of credit insurances matter most for exports to countries with intermediate levels of risk.¹⁰ The analysis shows that while exporters can initially benefit from the insurer's screening abilities once importers have build up a reputation for reliability, behavior will converge towards OA that in the medium- and long-run yields the largest flow payoffs.

⁹For the contour line at $\mathbf{J} = 0$ the infinite sum in part a) of the Corollary was numerically approximated for the given parameter values.

¹⁰Letters of credit have a similar insurance effect for the exporter when compared to trade credit insurances. The major difference is that letters of credit are typically issued by the importer's bank while trade credit insurances do not necessarily have a bank involvement.

4.5.3 Comparative statics

To gain some intuition on how the payment contract choice varies with parameter values, Figures 4.4 and 4.5 depict how the left panel of Figure 4.3 differs when conducting ceteris paribus changes on selected parameters. The exercise confirms the pattern that the usage of credit insurances is particularly pronounced at low levels of expected importer reliability and intermediate levels of institutional quality.

Figure 4.4 considers changes in the costs and quality of the trade credit insurance. The left panel shows that the range of parameters where CI is used is reduced when the insurance contract becomes more expensive through an increase in the fixed costs parameter κ . The usage of insurances now is strongly restricted to high levels of θ and intermediate levels of λ . A very similar pattern emerges when the insurer's ability to screen out the myopic types declines as can be seen in the right panel of the figure. Figure 4.5 then considers an increase in market demand as well as a variation in the exporter's costs of capital. Both scenarios confirm the highlighted pattern.

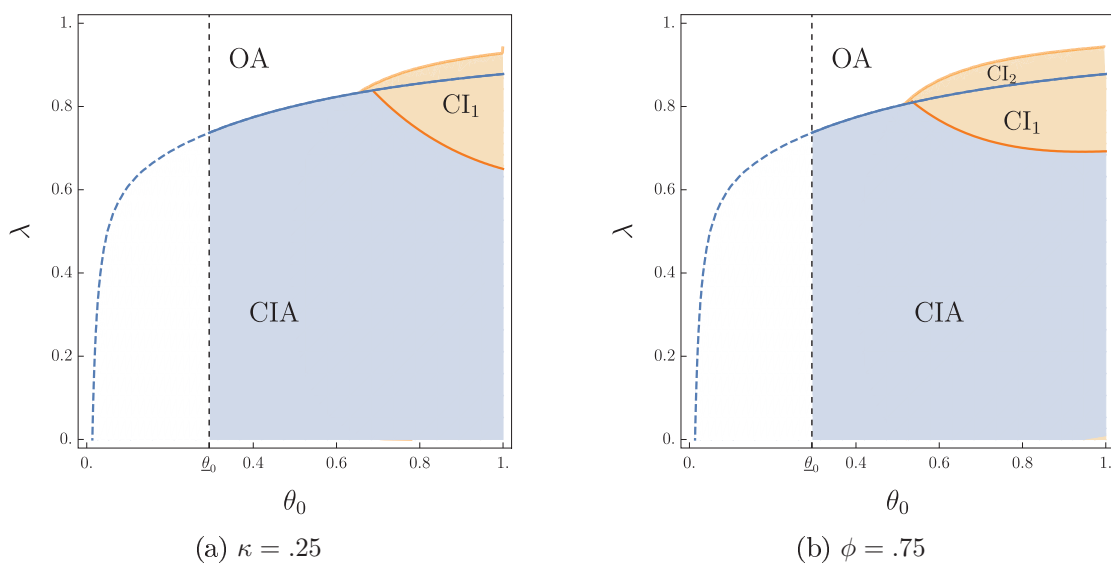


Figure 4.4: Ceteris paribus parameter changes to Figure 4.3 as indicated ($k = 0$)

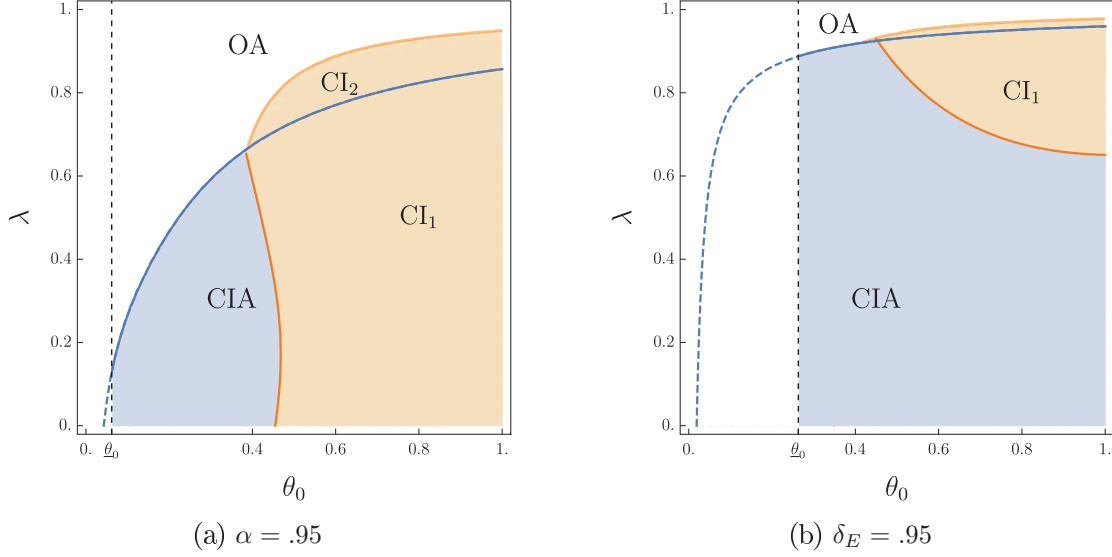


Figure 4.5: Ceteris paribus parameter changes to Figure 4.3 as indicated ($k = 0$)

4.6 Conclusion

In this paper, I have developed a dynamic model of payment contract choice in international trade where an exporter is challenged to choose the optimal form of trade finance in a sequence of export transactions. In doing so, he has to find the optimal balance between acquiring information about the quality of a potentially unreliable importer and handling the costs of capital and risk profiles of different payment contracts – a trade-off novel to the literature on payment contracts in international trade. To the best of my knowledge, this is also the first study to incorporate trade credit insurances into a dynamic model of payment contract choice.

The model can be extended in various directions. For example, I have focused on the case where all decision power on the payment contract is centralized on the exporter. While this may be a plausible assumption for situations where a large exporter can absorb all of the bargaining power on his side, other situations may come to mind in which importers have decisive influence. However, this does not impair the importance of the paper's main result that gives guidance on how to choose a payment contract that manages the potentially differing costs of capital between the contracting parties efficiently.

A central challenge for further research is the collection of data that can be used to discipline the predictions of a dynamic model of payment contract choice. Calibration of my model would require data that offers matched transaction-level information on the payment contracts that firms use in their export activities. Currently, such data is unavailable for a representative sample. Antràs and Foley (2015)

use highly detailed panel data to study payment contract choice over time, however it only covers one single exporter firm. This leaves plenty of scope for further investigations.

4.7 References

- AEBERHARDT, R., I. BUONO, AND H. FADINGER (2014): “Learning, Incomplete Contracts and Export Dynamics: Theory and Evidence from French Firms,” *European Economic Review*, 68, 219–249.
- ANTRÀS, P., AND C. F. FOLEY (2015): “Poultry in Motion: A Study of International Trade Finance Practices,” *Journal of Political Economy*, 123(4), 853–901.
- ARAÚJO, L., G. MION, AND E. ORNELAS (2016): “Institutions and Export Dynamics,” *Journal of International Economics*, 98, 2–20.
- AUBOIN, M., AND M. ENGEMANN (2014): “Testing the Trade Credit and Trade Link: Evidence from Data on Export Credit Insurance,” *Review of World Economics*, 150(4), 715–743.
- BAKER, G., R. GIBBONS, AND K. J. MURPHY (2002): “Relational Contracts and the Theory of the Firm,” *The Quarterly Journal of Economics*, 117(1), 39–84.
- BESEDEŠ, T. (2008): “A Search Cost Perspective on Formation and Duration of Trade,” *Review of International Economics*, 16(5), 835–849.
- BOARD, S., AND M. MEYER-TER VEHN (2015): “Relational Contracts in Competitive Labour Markets,” *The Review of Economic Studies*, 82(2), 490–534.
- CHASSANG, S. (2010): “Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts,” *American Economic Review*, 100(1), 448–465.
- DEFEVER, F., C. FISCHER, AND J. SUEDEKUM (2016): “Relational Contracts and Supplier Turnover in the Global Economy,” *Journal of International Economics*, forthcoming.
- EGAN, M. L., AND A. MODY (1992): “Buyer-Seller Links in Export Development,” *World Development*, 20(3), 321–334.
- FELBERMAYR, G. J., AND E. YALCIN (2013): “Export Credit Guarantees and Export Performance: An Empirical Analysis for Germany,” *The World Economy*, 36(8), 967–999.

- FOLEY, C. F., AND K. MANOVA (2015): “International Trade, Multinational Activity, and Corporate Finance,” *Annual Review of Economics*, 7(1), 119–146.
- HALAC, M. (2012): “Relational Contracts and the Value of Relationships,” *American Economic Review*, 102(2), 750–779.
- ITA (2007): “Trade Finance Guide: A Quick Reference for Exporters,” U.S. Department of Commerce: International Trade Administration.
- LEVIN, J. (2003): “Relational Incentive Contracts,” *American Economic Review*, 93(3), 835–857.
- MONARCH, R., AND T. SCHMIDT-EISENLOHR (2016): “Learning and the Value of Trade Relationships,” Working Paper.
- OLSEN, M. (2015): “How Firms Overcome Weak International Contract Enforcement: Repeated Interaction, Collective Punishment, and Trade Finance,” Working Paper.
- RAUCH, J. E., AND J. WATSON (2003): “Starting Small in an Unfamiliar Environment,” *International Journal of Industrial Organization*, 21(7), 1021–1042.
- SCHMIDT-EISENLOHR, T. (2013): “Towards a Theory of Trade Finance,” *Journal of International Economics*, 91(1), 96–112.
- SCHMIDT-EISENLOHR, T., AND F. NIEPMANN (2016): “International Trade, Risk and the Role of Banks,” Working Paper.
- THOMAS, J., AND T. WORRALL (1994): “Foreign Direct Investment and the Risk of Expropriation,” *The Review of Economic Studies*, 61(1), 81–108.
- VAN DER VEER, K. J. (2015): “The Private Export Credit Insurance Effect on Trade,” *Journal of Risk and Insurance*, 82(3), 601–624.
- YANG, H. (2013): “Nonstationary Relational Contracts with Adverse Selection,” *International Economic Review*, 54(2), 525–547.

4.A Appendix

4.A.1 Proof of Lemma 1

At the Production and Shipment stage of any period k of any export relationship the exporter will not deviate from the contract if and only if:

$$-cQ^A + \frac{\delta_E}{1-\delta_E}\bar{\pi}^A \geq (1-\lambda^*) \left(\delta_E \pi_0^A + \frac{\delta_E^2}{1-\delta_E}\bar{\pi}^A \right) + \lambda^* \left(-cQ^A + \frac{\delta_E}{1-\delta_E}\bar{\pi}^A \right) \quad (4.14)$$

Equation (4.14) states that making the effort to produce the contracted output plus the future payoff from continuing the current export relationship (where the exporter can be sure to interact with a patient type) must result in a higher payoff than deviating by not producing and shipping the agreed quantity Q^A . Note that deviation is possible only if contracts cannot be enforced which happens with probability $1-\lambda^*$. Simplification of (4.14) gives:

$$\delta_E \geq \frac{cQ^A}{\theta_0 \bar{\pi}^A} \equiv \underline{\delta_E} \quad (4.15)$$

Observe that the equilibrium does exist if and only if $\underline{\delta_E} < 1$. Consequently, a necessary and sufficient condition for existence is $\theta_0 > \frac{cQ^A}{\bar{\pi}^A}$.

Observe that the proof holds for every relationship period and not only for $k=0$ because deviation implies re-matching. This is the case, because the importer anticipates to again not receive the goods in the following transaction with again leads to a loss for her. It is therefore better for her to terminate the relationship. Hence, the exporter must re-match.

4.A.2 Proof of Lemma 2

At the Payment stage of any period $t+1$ it is incentive compatible for the patient importer to transfer the contracted share β of the revenue $R(q_t)$ if and only if:

$$-\beta R(q_t) + (1-\beta) \sum_{k=t+1}^{\infty} \delta_I^{k-t} R(q_k) \geq \lambda \left(-\beta R(q_t) + (1-\beta) \sum_{k=t+1}^{\infty} \delta_I^{k-t} R(q_k) \right) - (1-\lambda)\gamma R(q_t)$$

Simplifying gives:

$$(1-\beta) \sum_{k=t+1}^{\infty} \delta_I^{k-t} R(q_k) \geq (\beta-\gamma) R(q_t) \quad (4.16)$$

Note, that $\Lambda_k \rightarrow 1$ for $k \rightarrow \infty$ which implies that the optimal export quantity Q_k^Ω under open account and therefore revenue $R(Q_k^\Omega)$ are increasing with relationship

duration k . This observation allows us to derive a simple expression for the incentive compatible transfer share β from (4.16). Observe that the deviation incentive is largest in the limit (for relationship length $k \rightarrow \infty$) and denote the associated export quantity by Q^Ω . We can rewrite the IC-constraint to:

$$t \rightarrow \infty : \quad (1 - \beta) \frac{\delta_I}{1 - \delta_I} R(Q^\Omega) \geq (\beta - \gamma) R(Q^\Omega) \quad \Leftrightarrow \quad \beta \leq \delta_I + \gamma(1 - \delta_I) \quad (\text{IC}_I^\Omega)$$

If (IC_I^Ω) holds, the patient importer will never deviate from the contract no matter how large the exported volume under open account.

4.A.3 When the importer is more patient than the exporter

Noting that CIA now leads to larger flow payoffs than OA under full information we can derive by steps analogue to the main text a result equivalent to Proposition 3 but which now covers the case when $\delta_I > \delta_E \beta^\Omega$. This yields:

Proposition 3’. *Suppose that $\delta_I > \delta_E \beta^\Omega$. Consider the k th transaction in an export relationship. The exporter will conduct this transaction through Open Account if and only if*

$$\pi_k^\Omega - \pi^A(\theta_k^\Omega) \geq \delta_E [\bar{\pi}^A - \pi^A(\theta_{k+1}^\Omega)] \quad (4.17)$$

and ask for Cash in Advance otherwise.

The result gives the following behavioral predictions. Suppose that parameters are such that the initial transaction of the export relationship is conducted through CIA, i.e. expression (4.17) does not hold. Then any subsequent transaction with the same importer will be conducted through CIA as well since $\bar{\pi}^A > \bar{\pi}^\Omega$. Therefore no switching between contracts occurs. On the other side, suppose that OA is played initially which can be the case when $\pi_0^\Omega > \pi_0^A$ (and additionally, (4.17) holds). However, because $\bar{\pi}^A > \bar{\pi}^\Omega$ it must be the case that in some k behavior will switches to CIA and then stays there forever.

Chapter 5

Collusion and bargaining in asymmetric cournot duopoly – an experiment

Co-authored with Hans-Theo Normann

5.1 Introduction

How should firms collude in asymmetric Cournot oligopolies? There is no straightforward answer to this question. Joint-payoff maximization would require the inefficient firms to shut down; so when side payments are ruled out, joint-payoff maximization is entirely implausible. If, on the other hand, all firms produce positive amounts, the outcome is inefficient. Firms may produce quantities on the Pareto frontier but, as pointed out by Bishop (1960), Schmalensee (1987), and Tirole (1988), the bargaining frontier is convex due to the cost asymmetries and firms will disagree on the collusive price. Superior payoffs can be obtained if the firms alternately adopt the monopoly position in the market. But even with this approach, profits will fall short of the joint maximum and coordinating on the sequence and frequency of the alternating moves may be intricate and may raise antitrust suspicions.

In addition to the complex decision of how to produce the collusive output, firms must also decide how to divide the surplus. Schmalensee (1987) suggests using axiomatic theories of bargaining to solve this problem. Firms could agree on one of several possible bargaining solutions (Kalai-Smorodinski, equal relative gains, equal split, etc.) in order to determine the collusive outcome. Whereas firms may well coordinate on such outcomes, there are open questions. Schmalensee (1987) shows that different solutions imply different levels of joint payoffs and different gains (compared to static Nash) for asymmetric firms. So the bargaining problems firms face are more severe than in symmetric markets. Moreover, to even coordinate on one of the many solutions they would presumably need direct communication, so these solutions would only be available to explicit cartels.

Still in addition to these problems, asymmetric industries typically find it difficult to sustain a collusive agreement as a non-cooperative Nash equilibrium of a repeated game.¹ The repeated-game incentive constraint is more severe with cost asymmetries than in the symmetric case. A conventional wisdom maintains that asymmetries hinder collusion (see for example Ivaldi et al. 2003).

It seems fair to conclude that colluding in asymmetric Cournot markets is a formidable task. The Folk Theorem suggests many subgame-perfect equilibria exist in the repeated game and there are no focal points. Instead, firms face intricate efficiency, coordination, and bargaining problems. The incentive-compatibility and individual rationality of repeated game equilibria curbs the set of possible outcomes. So, can asymmetric firms collude successfully at all? And, if so, how?

We suggest a positive approach to answer these questions. We study how asymmetric Cournot duopolies collude in experiments. We believe that controlled laboratory settings with exogenously varied treatments nicely complement cartel studies. Laboratory studies do not suffer from the sample-selection problems that may affect data sets comprised of detected cartel cases. The experiment can control for

¹This problem is not specific to Cournot but can also occur for other types of oligopoly interaction. See Ivaldi et al. (2003).

cost and demand parameters, which allows an unambiguous interpretation of the outcomes.

Previous experiments on Cournot duopolies with asymmetric costs (Mason et al. 1992, Mason and Philips 1997, Selten et al. 1997, Fonseca et al. 2005, and Normann et al. 2014) have concentrated solely on implicit collusion, and they show that participants fail to reach supra-competitive payoffs throughout.^{2, 3, 4} In these experiments, subjects play asymmetric Cournot duopolies without any possibility to communicate. Behavior settles roughly around the static Nash equilibrium, but there are some discrepancies: high-cost firms produce more than predicted, low-cost firms less. Aggregate output is, if anything, above (rather than below) the level predicted by static Nash equilibrium. The failure of asymmetric Cournot duopolies to collude tacitly is in contrast to the symmetric duopolies reported in these papers where typically some level of tacit collusion can be maintained (Huck et al. 2004).

One of our main research questions is whether asymmetric Cournot duopolies can overcome the failure to collude by using explicit communication. Schmalensee (1987) emphasizes that his (theoretical) solutions apply to explicit cartels only. We also know that unrestricted communication usually leads to very effective cartels in symmetric oligopolies.⁵ It thus seems promising to extend the literature by studying how asymmetric Cournot duopolies operate when explicit communication is available.

A second main issue in our paper is motivated by the finding in the existing asymmetric Cournot experiments that quantities (and hence payoffs) are more equitable than predicted. When low-cost firms produce more and high-cost firms produce less than predicted, efficient firms earn correspondingly less. We expect this effect to

²Mason et al. (1992) were the first to observe that asymmetric Cournot markets are less cooperative and take longer to reach an equilibrium than symmetric markets. In a companion paper, Mason and Philips (1997) study the interaction of cost asymmetries and two different information environments: with private information, players know only their own payoffs. The results show that asymmetric markets are unaffected by private vs. public information but symmetric markets are more cooperative when payoffs are common knowledge. Fonseca et al. (2005) study endogenous timing in an asymmetric duopoly, that is, the endogenous emergence of a Stackelberg leader and follower. They also conduct control treatments with standard simultaneous-move (and cost-asymmetric) Cournot duopolies which can be compared to the previous experiments and this study. Finally, Normann et al. (2014) use symmetric and asymmetric Cournot duopolies to study the impact of the duration of an experiment.

³Selten et al. (1997) asked student participants of a game-theory seminar to program repeated-game strategies for asymmetric Cournot duopolies. The setup differs to the other experiments because of an additional asymmetry of a fixed-cost parameter. The significance of the results is somewhat difficult to assess because the group of subjects is relatively small and interacted repeatedly over the course of the entire term.

⁴Asymmetries among firms can alternatively be modeled as differences in production capacities. For recent experiments with capacity-asymmetric firms, see Harrington et al. (2016) and Fonseca and Normann (2008). In both studies, firms have the same preferences regarding prices, despite the asymmetry.

⁵Recently, see Cooper and Kühn (2014) and Fonseca and Normann (2012).

be less acute in the field because, there, firms engage in costly efforts to obtain a cost advantage, making them feel entitled to a correspondingly higher profit. This aspect is missing in existing experimental studies. We hence introduce treatments where the role of the low-cost firm is earned in a contest: before playing the Cournot part, subjects have to conduct a tedious real-effort task (taken from Charness et al. 2013).⁶ The best-performing participants then produce at low cost, the others at high cost. In this case, will the low-cost firm feel entitled to earn more? Will the discrepancies to static Nash disappear when the contest is introduced?

Our results are as follows. Without communication, firms fail to collude and essentially play the static Nash equilibrium, as suggested by the previous experiments. When express communication is available, joint profits increase but nearly all the gains from talking go to the inefficient firm. Bargaining solutions do not predict collusive outcomes well, except for the equal split solution. We employ coding analysis (Houser and Xiao 2011) to investigate the nature of the agreements reached in the treatments with communication and text-mining analysis (Möllers, Normann and Snyder 2016) to identify the language suitable for successful collusion. When the role of the efficient firm is earned in a contest, competitiveness is reduced without communication but increased with chat. Moreover, with earned roles and when firms can talk, they often collude by producing equal amounts – a collusive strategy unknown in the existing literature.

5.2 Model and benchmarks

We study Cournot duopoly markets with cost-asymmetric firms. Two firms, firm 1 and firm 2, choose non-negative quantities, q_i , $i = 1, 2$, as their actions. Their production costs are linear

$$C_i(q_i) = \theta_i q_i \quad i = 1, 2. \quad (5.1)$$

Assume $\theta_1 < \theta_2$, that is, let firm 1 be the low-cost firm. Firms face linear inverse demand

$$p(Q) = \max\{\alpha - Q, 0\} \quad (5.2)$$

where $Q = q_1 + q_2$ denotes aggregate output. In the unique static Nash equilibrium of this game, firms produce

$$q_i^* = \frac{\alpha - 2\theta_i + \theta_j}{3}, \quad i, j = 1, 2; \quad i \neq j \quad (5.3)$$

provided cost asymmetries are not too large. Nash equilibrium profits are $\pi_i^{NE} = (q_i^*)^2$, $i, j = 1, 2$. The monopoly output of firm i is

$$q_i^M = \frac{\alpha - \theta_i}{2}, \quad i, j = 1, 2 \quad (5.4)$$

⁶There are several experiments with “earned roles,” see, for example, Konow (2000), Fahr and Irlenbusch (2000), Gächter and Riedl (2005), and Oxoby and Spraggon (2008).

and the corresponding monopoly profit is $\pi_i^M = (q_i^M)^2$.

Consider now the three different ways of producing collusive outputs, as mentioned in the introduction. First, the joint-profit maximum would require $q_1 = q_1^M$ and $q_2 = 0$ and implies a joint profit of $(q_1^M)^2$. Only when side payments are feasible can this profit be freely allocated between firms. Secondly, firms may alternately produce their preferred q_i^M . The joint profit would be $\gamma(q_1^M)^2 + (1 - \gamma)(q_2^M)^2$ where γ denotes the likelihood or relative frequency of firm 1 being the monopolist.⁷ Thirdly, when both firms produce positive amounts in each period, they should from a normative perspective produce Pareto-efficient quantities. To derive the Pareto frontier, rewrite the payoff functions as $q_i = \pi_i / (p - \theta_i)$. Summing this expression up over both firms gives:

$$Q = \frac{\pi_1}{p - \theta_1} + \frac{\pi_2}{p - \theta_2} \quad (5.5)$$

Solving for π_2 and using $Q = \alpha - p$ yields the Pareto frontier:

$$\bar{\pi}_2(\pi_1) = \max_p \left(\alpha - p - \frac{\pi_1}{p - \theta_1} \right) (p - \theta_2) \quad (5.6)$$

Equation (5.6) states that for any given level of payoffs of firm 1, $\pi_1 \in [0, (q_1^M)^2]$, the payoffs of firm 2 on the Pareto frontier, $\bar{\pi}_2$, can be found by adjusting the market price p such that π_2 is maximized. This Pareto frontier is convex (Bishop 1960, Schmalensee 1987, Tirole 1988).

In addition, firms need to resolve the bargaining problem of how much either firm is going to earn. Schmalensee (1987) suggests bargaining theory at this point. The solution proposed by Kalai and Smorodinsky (1975), henceforth KS, maintains the ratios of the maximal payoffs (π_i^M) players can obtain in addition to Nash⁸:

$$\frac{\pi_1^{KS} - \pi_1^{NE}}{\pi_1^M - \pi_1^{NE}} = \frac{\pi_2^{KS} - \pi_2^{NE}}{\pi_2^M - \pi_2^{NE}}, \quad (5.7)$$

where π_i^{KS} , $i = 1, 2$, are the payoffs of firms 1 and 2 under the KS criterion. Roth's (1979) equal relative gains (ERG) suggests payoff increases proportional to the payoffs earned in the Nash equilibrium:

$$\frac{\pi_1^{ERG}}{\pi_1^{NE}} = \frac{\pi_2^{ERG}}{\pi_2^{NE}} \quad (5.8)$$

⁷Equivalently, firms could divide the market and allocate individual consumers to firms.

⁸The KS solution exists even though, as pointed out by Schmalensee (1987), the bargaining set is non-convex in the case of asymmetric, linear costs. As was shown by Conley and Wilkie (1991) it is sufficient for the existence of the KS solution that the bargaining set is *comprehensive*, which holds for our model. In a later paper, Conley and Wilkie (1996) also extend the Nash (1950) solution to non-convex but comprehensive bargaining sets. This solution coincides with the KS solution and we omit it here.

The equal split (ES) solution (Roth 1979)

$$\pi_1^{ES} = \pi_2^{ES} \quad (5.9)$$

has both firms earning the same absolute amount.

Figure 5.1 illustrates the three methods of producing collusive outputs and the bargaining solutions. The figure shows the Pareto frontier, the joint-profit maximum with side payments, and the alternating monopoly (AM) solution for $\gamma = .5$ (implying that firms literally take turns being the monopolist). Note that AM is not on the Pareto frontier but is (marginally) superior to it. The dashed lines indicate the outcomes that are Pareto superior to the static Nash equilibrium (NE). We also include the bargaining solutions (KS, ERG, ES) given firms are on the Pareto frontier (when both firms produce a positive amount).⁹




Figure 5.1: Feasible payoffs and bargaining solutions, exact plot for experimental parameters.

In the experiment, we use $\alpha = 91$, $\theta_1 = 13$, $\theta_2 = 25$ and $q_i \in \{0, 1, \dots, 45\}$. The corresponding benchmarks are given in Table 5.1. The table presents the quantities,

⁹The bargaining solutions could also be applied to JP and AM, however, we refrain from reporting these outcomes here. As argued in the introduction, they are often neither feasible nor realistic and turn out to be not empirically relevant in our experiment.

payoffs, and welfare of the Nash equilibrium,¹⁰ the joint-profit maximum, and the alternating monopoly solution (AM). For Pareto-efficient outcomes, we show the values for KS, ERG, and ES.

	Quantities			Profits			Welfare	
	q_1	q_2	Q	π_1	π_2	$\pi_1 + \pi_2$	CS	TW
NE	30	18	48.0	900	324	1224	1152	2376
JP	39	0	39.0	1521	0	1521	761	2282
AM	19.5	16.5	36.0	761	545	1306	653	1959
KS	23	13.3	36.3	958	396	1354	659	2013
ERG	24.2	12.3	36.5	1006	362	1368	666	2034
ES	14.7	20.4	35.1	631	631	1262	616	1878

Table 5.1: Benchmarks. Note that AM refers to an alternating monopoly with $\gamma = .5$ here, meaning that firms produce their preferred monopoly outputs (39 and 33) alternatingly which, on average, implies the quantities in the table.

Concluding, we comment on the incentive-compatibility of the collusive solutions. From Figure 5.1, it is clear that only KS and ERG can be sustained as subgame perfect equilibria in the repeated game. In a simple trigger-strategy equilibrium with Nash reversions, these outcomes would require a minimum discount factor of 0.603 and 0.663 for KS and ERG, respectively. JP with side payments is generally not feasible in our experiment, and ES and AM (with $\gamma = .5$) are not individually rational for firm 1. To make AM incentive-compatible for both firms would require $\gamma \in [100/169, 85/121]$, that is, firm 1 would need to be the monopolist between 59 and 70 percent of the time. These AM solutions would be Pareto superior to NE.

5.3 Experimental design

The experiments were implemented as a sequence of repeated games. Every participant had to play a total of five supergames. Each supergame was played with the same partner (fixed matching) but different supergames were played with a different player (absolute stranger re-matching, Dal Bó 2005). Players kept the role of either firm 1 or 2 for the entire experiment. At the end of every stage game, play would continue for another period with a probability of 75 percent. When the computer determined the end of a supergame, a new supergame with a different partner started. The actual number of periods in every supergame was determined ex ante and was the same in all sessions.

		Communication	
		No	Yes
Earned Roles	No	NoTalk-Random (66)	Talk-Random (68)
	Yes	NoTalk-Earned (62)	Talk-Earned (64)

Table 5.2: 2×2 factorial treatment design and the corresponding number of participants.

We have treatments with and without the opportunity to communicate (Talk and NoTalk), and with and without the possibility to earn the role of the efficient firm 1 (Earned and Random). This creates the 2×2 design in Table 5.2.

In the NoTalk-treatments, subjects had to post quantities in each period without being able to communicate with each other. In the Talk-treatments, subjects were allowed to communicate with the other firm once, before the beginning of each supergame. Communication was via typed messages, using an instant-messenger communication tool. Communication was unrestricted and subjects were allowed to exchange as many messages as they liked. However, they were not allowed to identify themselves. Unrestricted communication is important here because we expect players to discuss entire strategies, not just quantity targets (which could also be communicated through simple numeric announcements). The time to communicate was limited to three minutes in the first supergame, two minutes in the second supergame, and to 90 seconds in supergames three, four, and five.

Our second treatment variable varies how the role of the efficient firm was allocated. In the Random treatments, one subject was assigned the role of the more efficient firm by a random computer move. In the Earned treatments, winners of a pre-play contest were awarded that role. Here, subjects participated in a real-effort task at the beginning of the experiment to determine their role. During a period of five minutes subjects were instructed to translate letters into numbers using a translation table with one column of letters and a second column displaying the corresponding numbers (Charness et al. 2013). In each session, the better-performing half of the participants were assigned the role of firm 1 and the rest the role of firm 2. Participation in the effort task was voluntary and subjects could use their time as they wished as long as they did not interfere with the other subjects in the room. However, all our subjects decided to actively participate in the effort task.

¹⁰Further Nash equilibria close to this equilibrium may exist due to the discrete action space.

5.4 Procedures

We provided written experimental instructions which informed subjects of all the features of the market (the instructions are available in the Appendix). Subjects were told they were representing one of two firms in a market. The experiments were computerized, using z-Tree (Fischbacher 2007). Subjects learned their role (firm 1 or firm 2) from the computer screen.

In every period, subjects had to enter their quantity in a computer interface. On the decision screen, subjects also had access to a profit calculator which allowed them to compute firms' payoffs and the market price from the hypothetical quantity choices of firms 1 and 2. Once all subjects had made their decisions, the period ended and a screen displayed the quantity choices of both firms and the market price. On the screen was also displayed the individual payoff of the current period and the accumulated payoffs up to that point but not the payoffs of the other firm. When a supergame ended, a subsequent screen informed the subjects that they would now be re-matched with a new partner.

The experiment was conducted at the DICElab of Heinrich-Heine-University from April to November 2015. A total of 260 subjects participated in 12 sessions (three for each treatment). Subjects were mainly students and were randomly recruited (using ORSEE, Greiner 2015) from a pool of potential participants. Sessions lasted between 45 and 65 minutes.

Payments consisted of a show-up fee of € 5 plus the sum of payoffs attained during the course of the experiment. For payments, we used an experimental currency unit ("Taler"); 2,200 Taler were worth € 1. In order to keep the design balanced compared to Random, subjects assigned to the Earned treatments received a participation fee of 4,000 Taler for taking part in the effort task. Average earnings were € 11.39.

5.5 Results

5.5.1 Overview and behavior across supergames

We begin by reporting total output (Q) for each treatment across the five supergames (see Figure 5.2). We note a treatment effect due to communication: Talk strongly reduces output. The effect of Earned is less strong and depends on the communication mode: for NoTalk, the earned role reduces output but for Talk it increases it.

Figure 5.2 also shows that some learning is going on. In the Talk treatments, play becomes more collusive after the first supergame. Also, in the NoTalk treatments, outputs are somewhat volatile across the first couple of supergames. In order to account for the learning effects, we discard the first supergame from the data for the

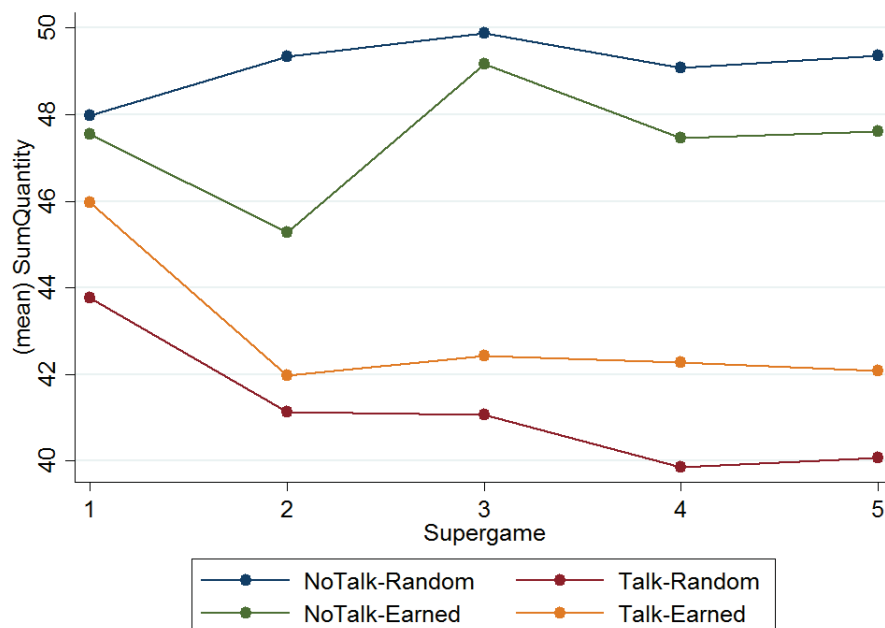


Figure 5.2: Aggregate quantities, Q , across supergames. For comparison, recall Q is 48 in the static Nash equilibrium and approximately 35 to 36 in the collusive solutions.

remainder of this section. Our results do not change qualitatively when we include the first supergame or when we exclude more supergames.

	Quantities			Payoffs			Welfare	
	Q	q_1	q_2	π_1	π_2	$\pi_1 + \pi_2$	CS	TS
NoTalk-Random	49.4 (0.15)	28.8 (0.74)	20.7 (0.88)	791 (28.7)	326 (17.2)	1117 (11.5)	1246 (6.00)	2363 (5.63)
NoTalk-Earned	47.8 (0.65)	28.6 (0.71)	19.2 (1.17)	820 (15.8)	312 (20.1)	1132 (10.4)	1182 (14.6)	2314 (24.9)
Talk-Random	40.4 (2.06)	18.8 (1.30)	21.6 (1.15)	673 (21.7)	537 (32.1)	1210 (31.2)	843 (90.1)	2053 (59.9)
Talk-Earned	42.4 (1.42)	22.5 (0.77)	19.9 (0.70)	756 (14.7)	457 (13.1)	1213 (25.3)	927 (63.5)	2140 (39.9)

Table 5.3: Summary of treatment averages for outputs, profits, and welfare measures (standard deviations in parenthesis, based on session averages), employing data from supergames 2 to 5.

Table 5.3 is the summary statistics of our study. It shows our main variables: aggregate and individual-level output, profit, consumer surplus, and welfare. We will repeatedly refer to this table in this results section.

Dep. variable	Quantities			Payoffs			Welfare	
	(1) Q	(2) q_1	(3) q_2	(4) π_1	(5) π_2	(6) $\pi_1 + \pi_2$	(7) CS	(8) TS
Talk dummy	-8.946*** (1.045)	-9.838*** (0.730)	0.892 (0.712)	-115.5*** (16.72)	207.8*** (17.72)	92.29*** (16.63)	-400.4*** (45.62)	-308.1*** (30.31)
Earned dummy	-1.673*** (0.333)	-0.0600 (0.471)	-1.613** (0.680)	32.47* (15.00)	-16.54 (11.98)	15.93* (7.383)	-63.53*** (8.086)	-47.61*** (12.99)
Talk \times Earned	3.404** (1.298)	3.595*** (0.889)	-0.191 (0.950)	51.70** (19.42)	-60.69** (20.95)	-8.988 (20.99)	138.4** (55.56)	129.4*** (38.47)
Constant	49.46*** (0.0721)	28.71*** (0.347)	20.75*** (0.418)	788.5*** (13.07)	327.9*** (7.851)	1,116*** (5.215)	1,246*** (2.813)	2,363*** (2.502)
Observations	1,820	1,820	1,820	1,820	1,820	1,820	1,820	1,820
R-squared	0.194	0.280	0.031	0.092	0.261	0.042	0.175	0.256

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 5.4: Ordinary least-squares regressions for output, payoff, and welfare variables, clustered at the session level.

5.5.2 The effect of communication

Our first main research question is whether firms collude at all, given the intricacies of cooperating in asymmetric Cournot. As can be seen in Figure 5.2 and Table 5.3, this strongly depends on the availability of Talk. Behavior in the NoTalk treatments fits well with the static Nash equilibrium. The prediction of $Q = 48$ is nearly perfectly matched: we observe in Table 5.3 $Q = 49.4$ for NoTalk-Random and $Q = 47.8$ for NoTalk-Earned. Clearly, there is no (tacit) collusion in the NoTalk-treatments. The Talk treatments are, in contrast, collusive. Table 5.3 shows that, compared to the NoTalk treatments, Talk reduces aggregate outputs by nine (Random) and 5.4 (Earned) quantity units, respectively. Whereas Q is greatly reduced in the Talk treatments ($Q = 40.4$ and $Q = 42.4$, respectively) compared to NoTalk, it does not reach the Q levels suggested by the KS, ERG, AM, and ES benchmarks (where Q is about 35 to 36).

We run linear regressions (clustered at the session level) in order to check the statistical significance of these effects. Table 5.4 reports the results. Regression (1) of Table 5.4 regresses Q and confirms that the treatment effect due to Talk is statistically highly significant.

Result 1. *The NoTalk treatments are not collusive as they match the predictions of the static Nash equilibrium. The Talk treatments are collusive in that aggregate output is significantly reduced compared to NoTalk.*

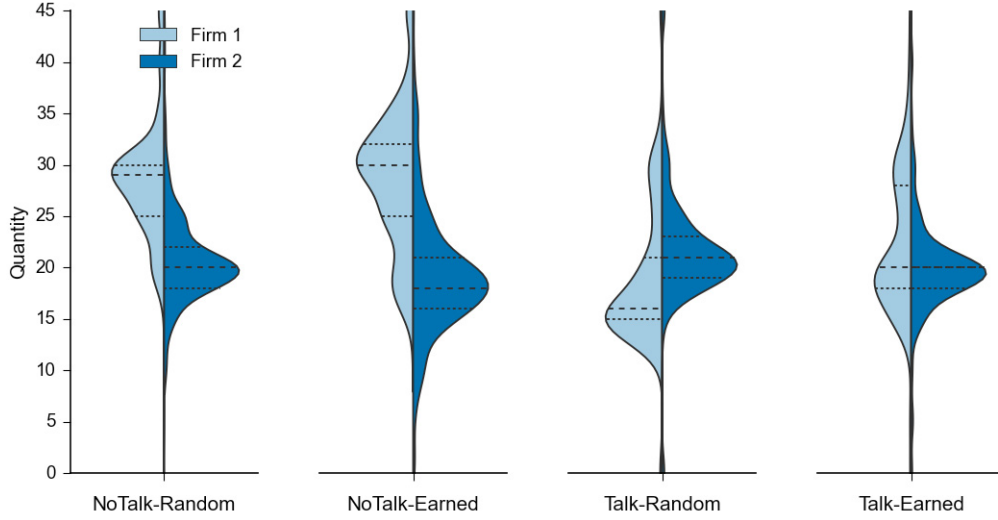


Figure 5.3: Kernel density estimates (KDE) for outputs by treatment and firm type. Medians are indicated by dashed lines, lower and upper quartiles by dotted lines.

We next turn to firm-specific behavior, looking at high- and low-cost firms separately. The Talk treatments show some surprising behavior. In Table 5.3, the outputs of firm 1 are greatly reduced compared to Nash; they read 18.8 and 22.5 for Talk-Random and Talk-Earned, respectively. This is below the level suggested by KS and ERG. Firm 2, in contrast, still produces output comparable to the Nash equilibrium. Comparing Talk and NoTalk, firm 2's output remains constant but firm 1 strongly reduces its quantity. These effects are statistically significant in regressions (2) and (3) of Table 5.4: the Talk dummy has a huge impact in (2) but there is nearly no effect in (3). The Kernel density estimates (KDE) of the distribution of outputs in Figure 5.3 illustrate the dramatic reduction of q_1 . In Talk-Random, the median outputs of firm 1 is even lower than for firm 2. We conclude that the collusive (Q -reducing) effect of communication is entirely due to the efficient firm reducing its output.

For NoTalk, we note some discrepancies with the static Nash equilibrium at the firm level. In Table 5.3 we see that the efficient firm 1 produces less than its NE quantity, firm 2's output is higher than its NE choice. Figure 5.3 illustrates and confirms these findings. For both NoTalk treatments, we see that the median firm

1 output is less than the predicted 30, for firm 2 the median is above 18. These results are in line with the experiments cited above.¹¹

Result 2. *The collusive effect of the Talk treatments is entirely due to the efficient firm reducing its output. In NoTalk, efficient firms produce less, inefficient firms more than in NE.*

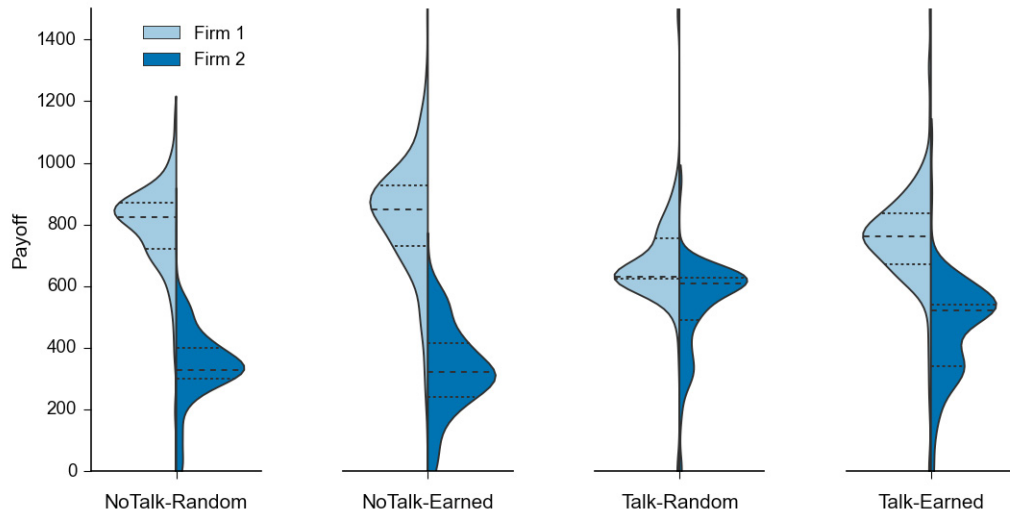


Figure 5.4: Kernel density estimates (KDE) for profits by treatment and firm type. Medians are indicated by dashed lines, lower and upper quartiles by dotted lines.

We finally turn to an analysis of profits, consumer surplus, and welfare. Industry profits increase significantly in Talk compared to NoTalk but they just about match the static Nash equilibrium levels. The efficient firms earn less than in the NoTalk treatments. At least on average play does not seem to be incentive-compatible. In terms of profits, firm 2 strongly benefits from the possibility to talk whereas firm 1 loses profits. Profit regressions (4) to (6) in Table 5.4 confirm that these effects are significant. The Kernel density estimates (KDE) of the distribution of profits in Figure 5.4 also show this: there is a substantial downward shift in the density of firm 1's profit and a corresponding upwards shift for firm 2.

¹¹Since all existing asymmetric Cournot experiments are also done with linear demand and cost, we can normalize the results by using the ratio of actual and Nash equilibrium output, q_i/q_i^{NE} , for a detailed comparison. For q_1 , we find that this ratio varies between 0.96 and 0.99 in previous studies compared to 0.95 (Earned) and 0.96 (Random) for our two NoTalk treatments. For q_2 , the ratio in previous studies varies between 1.03 and 1.09, and it is 1.07 (Earned) and 1.15 (Random) in our case. This confirms that our results are in line with previous results.

In our setup, consumer surplus is affected by aggregate output only, any reduction of Q also reduces CS. It follows that consumers lose when Talk is introduced. This effect is statistically confirmed by regression (7) in Table 5.4.

Total welfare depends not only on aggregate output but also on how efficiently this output is produced (that is, how much firm 1 produces). In any event, since the reduced aggregate output is due to the efficient firm producing less, there are no efficiency gains in production due to the collusion in Talk. Consequently, total welfare is reduced in the Talk treatments (see regression (8) in Table 5.4).

Result 3. *Comparing Talk to NoTalk, industry profits increase significantly whereas consumer surplus and total welfare decrease significantly.*

5.5.3 The effect of earned roles

Figure 5.2 and Table 5.3 shows the ambiguous effect the contest for the role of the efficient firm has: Q is smaller when comparing NoTalk-Earned to NoTalk-Random but is larger when comparing Talk-Earned to Talk-Random. In regression (1) of Table 5.4, the Earned dummy has a significant negative coefficient but $\text{Talk} \times \text{Earned}$ is positive significant.

Compared to NoTalk-Random, the NoTalk-Earned treatment leaves the efficient firm virtually unchanged but the inefficient firm produces significantly less (see regressions (2) and (3)). Taken together, these two effects imply the significant reduction of Q noted above.

The comparison of Talk-Random and Talk-Earned is also revealing. Regressions (2) and (3) of Table 5.4 and Figure 5.3 show that this increase is due to firm 1 increasing its output. Median outputs are the same for both firms in Talk-Earned but it is still not the case that the efficient firm produces more, as implied by KS and ERG. Kernel density estimates of the distributions of profits are shown in Figure 5.4. They suggest the following interpretation: whereas in Talk-Random firm 1's profits are very much centered around the equal split and have little density toward higher profits, the Talk-Earned treatment has a substantial amount of density for profits of 800 and above toward the NE level. Nevertheless, even in Talk-Earned a large fraction of firm 1 subjects still reach profits close to the equal split.

Result 4. *When the role of the efficient firm is earned in a contest, aggregate output decreases in NoTalk but increases in Talk. The effect in NoTalk-Earned is almost exclusively due to the inefficient firm reducing its output whereas the effect in Talk-Earned is almost exclusively due to the efficient firm expanding its output.*

5.5.4 Bargaining outcomes

Figure 5.5 shows at the session level the market outcomes in the π_1 - π_2 space. Each dot represents the average of one session. The figure shows that there is relatively little heterogeneity at the session level, suggesting our results are robust. The figure also confirms some of the previous results: NoTalk outcomes (green and blue dots) are close to the static Nash equilibrium, the communication in Talk shifts payoffs toward the equal split (ES) outcome in a way that is not incentive-incompatible for firm 1.

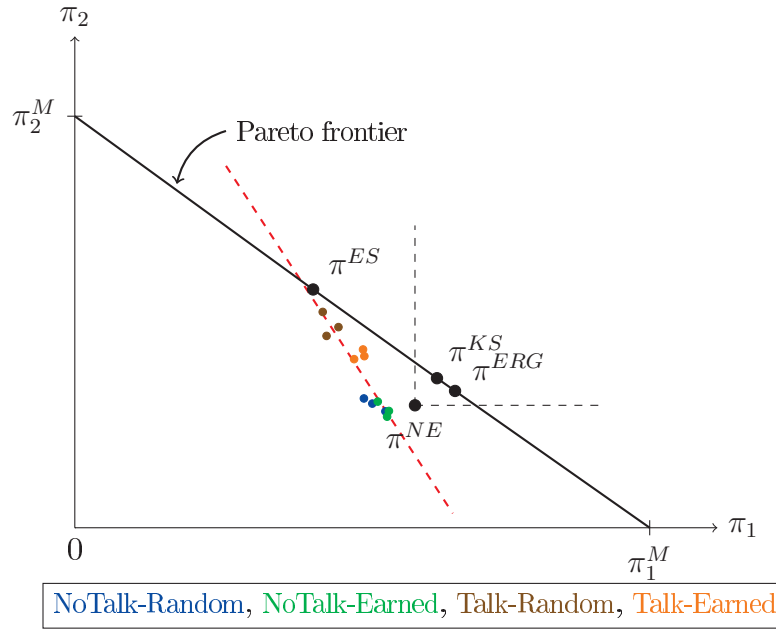


Figure 5.5: Session means in the π_1 - π_2 space and benchmarks. The dashed red line shows a linear regression of session averages, suggesting $\pi_2 = 1568 - 1.53\pi_1$ with an R^2 of .86.

Session averages are remarkably well organized by an (ad hoc) linear regression line – the red dashed line in Figure 5.5. Starting from the green NoTalk-Earned dots, session averages in the other treatments “move” along this line toward ES where the brown Talk-Random dots are close to the ES point. The slope of the regression line is -1.53 , indicating that joint profits increase when one compares NoTalk and Talk treatments. All Talk session averages, however, lie to the north-west of the green NoTalk-Earned dots, indicating that only firm 2 gains from talking. In any event, the good fit of the regression line with the data suggest that sessions settle between the polar points NE and ES.

5.5.5 How do firms collude in the Talk treatments?

Our second main research question is, provided firms manage to cooperate, how exactly do they collude? We have already seen that firm 1 largely bears the burden of collusion. This section deepens the analysis by examining the nature of the agreements the firms coordinated on. Since there are no tacit agreements in NoTalk, we focus here on Talk-Random and Talk-Earned.

To identify agreements in the chat of our 330 market pairs in Talk, we recruited two coders and provided them with incentives to code the chat data (Houser and Xiao 2011). We asked them to independently judge every chat dialog for whether a mutual and clearly specified agreement was present. Both were paid 10 cents for every evaluated market pair if and only if both drew the same conclusion regarding the agreement status. Their evaluations matched for 92% of market pairs. A measure of inter-coder reliability, Cohen's kappa, is $\kappa = 0.83$ which can be considered high.

The majority of market pairs in our Talk treatments decided to agree on a clearly specified joint production plan. In the Talk-Random sessions 74.3% of the markets came to an agreement whereas in the Talk-Earned sessions this rate was with 62.5% significantly lower (two-sided Fisher exact, $p = 0.083$). Only around 1% of market pairs did not make any use of the possibility to exchange messages at all.

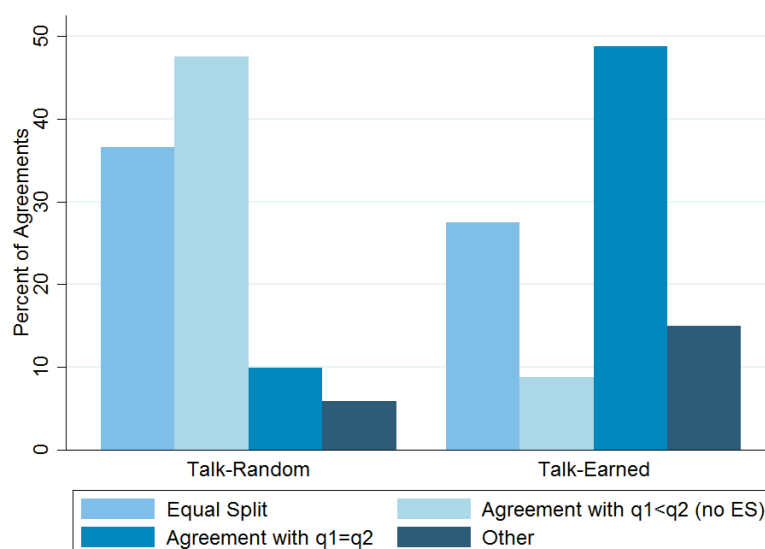


Figure 5.6: Distribution of agreement types, conditional on reaching an agreement in the first place.

Figure 5.6 gives an impression of the different and most prominent types of agreements in the two Talk treatments. We see clear treatment effects.¹² In Talk-Random, we note that a substantial share of firms (84.2%) agreed on either the equal split or other allocations in which the more efficient firm 1 produces a smaller amount than firm 2.¹³ Those agreements with $q_1 < q_2$ that do not fully imply the equal split nevertheless have a strong similarity to it; outputs in these cases were, on average, $(q_1, q_2) = (15.2, 21.6)$. Most prominent among the agreements not covered by the previous bins are agreements on equal quantities ($q_1 = q_2$) with about 9.9% of the agreements in that treatment.

In Talk-Earned, agreements at or close to the equal split still amount to about 27.5% but are clearly much more rare. With 48.8%, the most frequent type of agreement has both firms producing the same output. The vast majority of these agreements (74.4%) involve quantities of 19 or 20 and all other observations have slightly lower quantities. None of the agreements that we mainly observe in our experiments can be implemented as equilibria of the repeated game.




Figure 5.7: Deviation frequencies from agreements

The equal-output agreements are somewhat surprising at first sight and they do not appear in our list of benchmarks or, in fact, anywhere in the literature. One possible reason for these agreements is that equal quantities are strategically

¹²For the data categorization chosen for the plots, a two-sided 4×2 Fisher exact test shows that the distributions of agreement types differ significantly between the Talk treatments ($p < 0.001$).

¹³We included the following agreements into the bin for the equal split: $(q_1, q_2) \in \{(14, 20), (14, 21), (15, 20), (15, 21)\}$.

particularly simple and can serve as an easy and transparent target for agreement. In our data, equal quantities give firm 1 a payoff advantage when compared to the equal split outcomes. In fact, most agreements of this type involve $q_1 = q_2 = 19$ or $q_1 = q_2 = 20$ whereas the Pareto-efficient agreements (conditional on $q_1 = q_2$) range from $q_1 = q_2 = 16$ to $q_1 = q_2 = 20$. In other words, among the plausible $q_1 = q_2$ agreements, those favoring firm 1 are chosen almost throughout. To a certain extent, the efficient firm seems to make use of its cost advantage when roles in the Talk treatments are Earned.

Dependent variable	Deviation
Earned dummy	0.523 (0.340)
Firm 1 dummy	0.395*** (0.00394)
Initial period dummy	-0.917*** (0.323)
Earned \times Firm 1	-0.142** (0.0560)
Earned \times Initial period	-0.123 (0.408)
Initial Period \times Firm 1	0.325 (0.300)
Initial \times Firm 1 \times Earned	-0.144 (0.454)
Constant	-1.116*** (0.138)
Observations	1,206
Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1	

Table 5.5: Probit regression with dummy *deviation* as the dependent variable, clustered at the session level.

In the next step we analyze how stable the agreements were over the course of the supergame. Figure 5.7 shows by treatment, point in time, and firm type the share of periods where a (unilateral or bilateral) deviation from an agreement occurred in a market. Several interesting patterns emerge. First, significantly fewer deviations from agreements occur in the first period of the supergames when compared with the later periods. This relationship emerges in both treatments and for both firm types. Second, deviations tend to occur more often in the Talk-Earned than in the Talk-Random treatment, particularly for the non-initial periods of the game. Third,

the efficient firms deviate substantially more often than the inefficient firms. The probit regressions in Table 5.5 confirm these effects are statistically significant.

To understand these patterns in the deviations, it is instructive to look at the best responses (and therefore optimal deviations) given the other firm chooses its agreed upon quantity. Taking the two most popular agreement types – equal split and equal quantities – and comparing the size of the optimal deviations for firms 1 and 2 indicates that firm 1 should make a much larger deviation from its agreed quantity (+3 units when $(q_1 = 15, q_2 = 21)$ and +6 units when $q_1 = q_2 = 19$). Comparing the magnitude of the actual average deviations between firms confirms this pattern. The average deviation of firm 1 from its agreement quantity is +1.97, while for firm 2 it is with +0.22 significantly lower ($p = .011$). There is not much difference between the size of the average deviations in the two Talk treatments.

5.5.6 Text-mining analysis

Concluding the results section, we want to analyze in more detail the use of language in the chat treatments. What kind of language is useful to support collusion? Do efficient and inefficient firms differ in the content of their communication? And how do the Random vs. the Earned treatments affect chat?

Our method of analyzing the text will be *text mining*. To our knowledge, Möllers, Normann and Snyder (2016) were the first to use this method in experimental economics. Text-mining methods extract keywords from a body of text, referred to as a corpus. We will compare the most frequently used keywords for two corpora in order to find out how the corpora (the chats) differ. To be more precise, we will use Huerta's (2008) *relative rank difference* which tells us which keywords are comparatively more frequently used in corpus c relative to c' . Formally, we measure the keyness of word w in corpus c relative to c' by generating ranks $r_c(w)$ for all words w in corpus c according to frequency (and in descending order). The difference in the rank of w in corpus c relative to corpus c' is defined as

$$rd_c^{c'} = \frac{r_{c'}(w) - r_c(w)}{r_c(w)}$$

We restrict ourselves to keywords that are among the top 50 most common in corpus c , avoiding keywords with a high $rd_c^{c'}$ that are nevertheless rarely used. We omit conjunctions, prepositions, and articles and report keywords with $rd_c^{c'} > 1$.

We are interested in how treatments Random vs. Earned are reflected in the communication and how collusive supergames differ from non-collusive ones. We further expect different firm types to use different language.

The top panel of Table 5.6 reveals some interesting insights into the differences in chat in Random vs. Earned. In Random, the output pairs that yield the equal split ("15" and "21") occurs in the list, as does the neighboring "14" and the word "equal" (payoffs). There are some differences between firm types: inefficient firms

point out the resulting payoff of “630” that both players would earn and use the wink emoji whereas efficient firms suggests higher quantities or merely confirms (“ok”). Chat in the Earned treatments differs substantially. Keywords used by the efficient firms literally result in the equal-output strategy: “let us both [produce] 19 (20) in each round.” The words “very” and “only” appear in different contexts. When inefficient firms want to move away from the equal-output strategy, they have to suggest different outputs for different firms and hence use “firm 1.”

Language used in (un)successfully colluding duopolies can be obtained from the bottom panel of Table 5.6. Both types of successfully colluding firms use the confirming “yes” and the equal-payoff output combination (“14”, “19”). Successful collusion requires strategies that draw a distinction between “firm 1” and “firm 2.” Inefficient firms add “thanks,” indicating that they are aware that the efficient firm could earn more by playing non-cooperatively. The non-collusive groups mention higher output targets among their keywords. Words like “but” and “not” hardly indicate successful collusion, whereas “me” hardly indicates the mutuality inherent to agreements.

5.6 Conclusion

Asymmetric Cournot markets are intricate because there are no focal points and such oligopolies are accompanied by bargaining problems, inefficiencies (Bishop 1960, Schmalensee 1987 and Tirole 1988) and tight incentive constraints (Ivaldi et al. 2003). The main research questions of this paper are simple: do firms manage to collude at all in asymmetric Cournot markets and, if so, how?

The first set of findings indicates that our duopolies with express communication have both firms produce a positive amount in each period. The alternative technology, alternating monopoly, is occasionally discussed in the chats but are largely dismissed and are not frequent in the data. As for the set of results concerning the bargaining outcome, we find that almost all the gains from explicit communication go to the inefficient firm. Only when the role of the efficient firm is earned in a contest do the low-cost firms perform somewhat better. Unexpectedly, subjects revert to an equal-output strategy in this case.

We also confirm existing experiments (Mason et al. 1992, Mason and Philips 1997, Fonseca et al. 2005, and Normann et al. 2014) in that duopolies are close to the static Nash equilibrium without communication. We add to this literature by showing that allocating the role of the efficient firm through a contest is suitable for removing some discrepancies observed in the previous literature (and in our baseline treatment).

Our analysis suggests that players often make use of some symmetry criterion (equal profits, equal outputs) to reach an agreement. But since firms are asymmetric, this introduces inefficiencies which are an impediment to successful collusion. We

Efficient firm				Inefficient firm			
Random		Earned		Random		Earned	
word	rd	word	rd	word	rd	word	rd
14	7.94	very	9.78	14	13.11	units	1.46
21	7.25	only	7.41	;-)	2.93	firm 1	1.44
equal	5.20	round	6.06	equal	2.83	us	1.21
23	4.77	both	4.25	21	2.00	20	1.17
17	4.53	us	3.93	630	1.24		
ok	1.33	let	2.36				
15	1.20	quantity	1.98				
		each	1.58				
		19	1.40				
		20	1.33				
		best	1.19				

Efficient firm				Inefficient firm			
Collusive		Non-collusive		Collusive		Non-collusive	
word	rd	word	rd	word	rd	word	rd
firm 2	8.11	25	12.22	14	2.67	30	6.82
firm 1	2.24	me	9.81	thanks	2.66	25	2.82
14	2.10	produce	2.43	19	1.93	have	1.46
19	2.09	20	2.00	too	1.84	17	1.38
yes	1.50	30	1.33	yes	1.75	20	1.17
make	1.14	but	1.31	has	1.26		
	1.14	output	1.26	have	1.05		
	1.06	not	1.09				

Table 5.6: Text-mining analysis. We report words with absolute rank $r_c \leq 50$ and relative rank differential $rd \geq 1$.

thus strongly confirm that *asymmetries hinder collusion* (Ivaldi et al. 2003). Even with explicit communication firms do not manage to increase profits substantially. Our data imply that, given the choice of whether or not to (costly) communicate (that is, choose whether to participate in a cartel), the efficient firms would decline and would rather not communicate.

We conclude with a more general note. Experimental economists usually analyze bargaining and cooperation as separate phenomena. Bargaining experiments typically take a certain pie as given and abstract from problems that might arise when the pie has to be generated through cooperation in the first place. Likewise, cooperation experiments are bland in the bargaining dimension of the problem. In the standard prisoner's dilemma, for example, mutual cooperation implies a symmetric payoff division and, moreover, maximum joint payoffs, so in terms of bargaining there is not much to disagree on. The key point of this paper is that players often have to resolve bargaining and cooperation problems *uno actu*, which may be troublesome in the presence of asymmetries. Cartels need to simultaneously resolve bargaining frictions and ensure reliable cooperation, using one action. In symmetric settings, again, this is not exactly a major problem because players have the same preferences about, say, the cartel price, and the joint-profit maximum can be implemented. In asymmetric games including cartel, public-good, and common-pool games, the preferred bargaining outcomes might not be supported by stable cooperation, and outcomes that are incentive-compatible may not be desirable from a bargaining perspective. Our paper illustrates the relevance of the bargaining and the cooperation dimension in a Cournot setup. We believe that asymmetric cooperation games, especially with explicit communication, deserve more attention.

5.7 References

- BISHOP, R. L. (1960): "Duopoly: Collusion or Warfare?," *The American Economic Review*, 50(5), 933–961.
- CHARNESS, G., D. MASCLET, AND M. C. VILLEVAL (2013): "The Dark Side of Competition for Status," *Management Science*, 60(1), 38–55.
- CONLEY, J. P., AND S. WILKIE (1991): "The Bargaining Problem Without Convexity: Extending the Egalitarian and Kalai-Smorodinsky Solutions," *Economics Letters*, 36(4), 365–369.
- (1996): "An Extension of the Nash Bargaining Solution to Nonconvex Problems," *Games and Economic Behavior*, 13(1), 26–38.
- COOPER, D. J., AND K.-U. KÜHN (2014): "Communication, Renegotiation, and the Scope for Collusion," *American Economic Journal: Microeconomics*, 6(2), 247–78.

- DAL BÓ, P. (2005): “Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games,” *The American Economic Review*, 95(5), 1591–1604.
- FAHR, R., AND B. IRLBUSCH (2000): “Fairness as a Constraint on Trust in Reciprocity: Earned Property Rights in a Reciprocal Exchange Experiment,” *Economics Letters*, 66(3), 275–282.
- FISCHBACHER, U. (2007): “z-Tree: Zurich Toolbox for Ready-Made Economic Experiments,” *Experimental Economics*, 10(2), 171–178.
- FONSECA, M. A., S. HUCK, AND H.-T. NORMANN (2005): “Playing Cournot Although They Shouldn’t,” *Economic Theory*, 25(3), 669–677.
- FONSECA, M. A., AND H.-T. NORMANN (2008): “Mergers, Asymmetries and Collusion: Experimental Evidence,” *The Economic Journal*, 118(527), 387–400.
- (2012): “Explicit vs. Tacit Collusion—The Impact of Communication in Oligopoly Experiments,” *European Economic Review*, 56(8), 1759–1772.
- GÄCHTER, S., AND A. RIEDL (2005): “Moral Property Rights in Bargaining with Infeasible Claims,” *Management Science*, 51(2), 249–263.
- GREINER, B. (2015): “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE,” *Journal of the Economic Science Association*, 1(1), 114–125.
- HARRINGTON JR, J. E., R. H. GONZALEZ, AND P. KUJAL (2016): “The Relative Efficacy of Price Announcements and Express Communication for Collusion: Experimental Findings,” *Journal of Economic Behavior & Organization*, 128, 251–264.
- HOUSER, D., AND E. XIAO (2011): “Classification of Natural Language Messages Using A Coordination Game,” *Experimental Economics*, 14, 1–14.
- HUERTA, J. M. (2008): “Relative Rank Statistics for Dialog Analysis,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 965–972. Association for Computational Linguistics.
- IVALDI, M., B. JULLIEN, P. REY, P. SEABRIGHT, AND J. TIROLE (2003): “The Economics of Tacit Collusion,” *Final Report for DG Competition, European Commission*.
- KALAI, E., AND M. SMORODINSKY (1975): “Other Solutions to Nash’s Bargaining Problem,” *Econometrica*, 43(3), 513–518.
- KONOW, J. (2000): “Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions,” *The American Economic Review*, 90(4), 1072–1091.

- MASON, C. F., AND O. R. PHILLIPS (1997): “Information and Cost Asymmetry in Experimental Duopoly Markets,” *Review of Economics and Statistics*, 79(2), 290–299.
- MASON, C. F., O. R. PHILLIPS, AND C. NOWELL (1992): “Duopoly Behavior in Asymmetric Markets: An Experimental Evaluation,” *The Review of Economics and Statistics*, 74(4), 662–670.
- MÖLLERS, C., H.-T. NORMANN, AND C. M. SNYDER (2016): “Communication in Vertical Markets: Experimental Evidence,” *International Journal of Industrial Organization*, forthcoming, also: DICE Discussion Paper No 226.
- NASH, J. F. (1950): “The Bargaining Problem,” *Econometrica: Journal of the Econometric Society*, 18(2), 155–162.
- NORMANN, H.-T., T. REQUATE, AND I. WAICHMAN (2014): “Do Short-Term Laboratory Experiments Provide Valid Descriptions of Long-Term Economic Interactions? A Study of Cournot Markets,” *Experimental Economics*, 17(3), 371–390.
- OXOBY, R. J., AND J. SPRAGGON (2008): “Mine and Yours: Property Rights in Dictator Games,” *Journal of Economic Behavior & Organization*, 65(3), 703–713.
- ROTH, A. E. (1979): *Axiomatic Models of Bargaining*. Springer.
- SCHMALENSEE, R. (1987): “Competitive Advantage and Collusive Optima,” *International Journal of Industrial Organization*, 5(4), 351–367.
- SELTEN, R., M. MITZKEWITZ, AND G. R. UHLICH (1997): “Duopoly Strategies Programmed by Experienced Players,” *Econometrica*, 65(3), 517–555.
- TIROLE, J. (1988): *The Theory of Industrial Organization*. MIT Press.

5.A Appendix

Instructions (Talk-Earned treatment)

Part 1 (only used in the Earned treatments)

Welcome to our experiment. Please read these instructions carefully. Please do not speak with your neighbor and remain silent during the entire experiment. If you have any questions please raise your hand. We will come to your seat and answer your questions in person.

In this experiment you will repeatedly make decisions and thereby earn money. How much you earn depends on your decisions and the decisions of the other participants. You will receive a show up fee of €5 for participating in the experiment. This amount will increase by the earnings that you make during the course of the experiment. At the end of the experiment the Taler you earned will be exchanged at an exchange rate of 2200 Taler = €1. All participants receive (and currently read) the same set of instructions. You will remain anonymous to us and all other participants of the experiment. We do not save any data relating to your name.

In this experiment you will repeatedly make decisions for a firm in a market. There are two firms in the market: firm 1 and firm 2. The experiment consists of two parts. In the first part of the experiment your decisions and the decisions of all other participants will determine whether, in the second part of the experiment, you will take the role of firm 1 or firm 2. In part 2, one participant in the role of firm 1 will always interact with one participant in the role of firm 2. Both firms produce and sell the same product in a market, however firm 1 has lower production costs than firm 2. The payoff of a firm depends on its production costs.

Your task in part 1 is to translate letters into numbers during a period of five minutes. Your screen will show a table with two columns, where the first column shows letters and the second column shows the corresponding numbers. The computer will provide you with a letter and you have to enter the corresponding number into the box on your screen. Subsequently you click on “OK”.

When you confirm your answer you will be informed immediately whether your answer was right or wrong. In case the answer is wrong you will have to re-enter a number until your answer is correct. A new letter will only be shown once the current letter has been correctly translated into a number.

As soon as you have confirmed a correct answer the translation table will be recompiled with new letters and numbers and a new letter will be displayed for translation. You can translate an arbitrary amount of letters during the given time of five minutes.

In order to acquaint yourself with the program a test period will take place before the process starts.

For participating in part 1 of the experiment you will receive 4000 Taler, independently of the amount of translated letters.

Your role for part 2 of the experiment will be assigned as follows:

- After the five minutes have passed all participants will be assigned to one of two groups depending on the amount of letters that they translated. You will be assigned to group 1 if you conducted more correct translations than at least half of all participants. Otherwise you will be assigned to group 2. Thus, all participants in group 1 have translated more letters than all participants in group 2. In case several participants have translated the same amount of letters the computer will order them at random in order to guarantee an assignment to a group.
- If you were assigned to group 1 you will take the role of firm 1 in the following (firm 1 has lower production costs than firm 2)
- If you were assigned to group 2 you will take the role of firm 2 in the following (firm 2 has higher production costs than firm 1)

Immediately after the completion of the translation task you will be informed whether you have been assigned the role of firm 1 or firm 2. You will keep this role until the end of the experiment.

Please notice that in part 1 of the experiment participation in the translation task is not mandatory. Alternatively, you can e.g. read something or surf the Internet. However, please do not speak with your neighbor and please keep quiet.

Part 2 (only shown to subjects after completion of Part 1)

In this part of the experiment you repeatedly make decisions for a firm in a market. There are two firms in the market: firm 1 and firm 2. You have been assigned to the role of either firm 1 or firm 2 in the first part of the experiment. You will keep this role until the end of the experiment.

Both firms produce and sell the same product on a market. Every game of the experiment consists of several periods and in every period both firms must simultaneously make the same decision: Which quantity of the good do you want to produce? In every period, both firms can each produce a maximum quantity of 45 units of the good. Thus, the quantity you choose must be in between 0 and 45.

Information about the market and the firms:

- There is a uniform market price for both firms. The price (in Taler) which you receive for every unit of your product is calculated as follows:

$$\text{market price} = 91 - \text{production quantity firm 1} - \text{production quantity firm 2}$$

With every unit that your firm produces the market price for both firms is reduced by one Taler. Please keep in mind that the decision of the other firm has the same effect on the price: Every additional unit produced by the other firm also reduces the market price by one Taler.

- Both firms have different production costs: Firm 1 has per unit production costs of 13 Taler. Firm 2 has per unit production costs of 25 Taler.
- Your payoff per unit sold is the difference between the market price and your per unit production costs

$$\text{Payoff per unit} = \text{market price} - \text{production costs per unit}$$

Please note that you make a per unit loss if the market price is below your per unit production costs.

- Your per period payoff is equal to your per unit payoff multiplied by the number of sold units:

$$\text{Payoff per period} = \text{Payoff per unit} * \text{Number of sold units}$$

You can assume that all the units you produce can also be sold.

- In order that you can see which quantities lead to which payoffs we provide you with a payoff calculator. With it you can calculate the profits which result from different quantity combinations on your screen before you make your actual decision. Before the beginning of the first game you will have the opportunity to acquaint yourself with the profit calculator.
- At the end of every period you will receive feedback about the quantity decisions of both firms, the realized market price and your payoff. Additionally, the computer will show the total payoffs that you obtained so far.

Example: Suppose that you are firm 1. Thus, your production costs are 13 Taler per produced unit. You decide to produce 30 units of the good. Subsequent to your decision you receive the information that firm 1 decided to produce 18 units of the good itself. Hence, the resulting market price is 43 Taler and your per unit payoff in this period is 43 Taler - 13 Taler = 30 Taler. Thus, your payoff in this period is 30 Taler * 30 = 900 Taler.

Course of action: Every game consists of one or several periods. After every period, chance decides whether another period takes place: The computer randomly draws a number between 1 and 4. If a “1”, “2”, or “3” is drawn then another period takes place; with a “4” no further period is conducted and the current game ends. Hence, it can happen that a game is over already after a single period. Equally it can happen that a game continues for many rounds. As soon as a game ends a new game will be started. The experiment consists of a total of 5 games for which the following holds:

- All games of the experiment have the same structure. This means that in every period of every game the above-described production decision has to be made.

- If a game ends you will be assigned to a new partner at the beginning of the next period. You will not meet any previous partner in any future game.

Communication: At the beginning of every game the two firms have the possibility to communicate with each other. For this purpose a text box will appear on your screen. You and the other firm can exchange typed messages in it in which you can talk about anything. The only restriction is that you must not identify each other (e.g. do not write your name). There will be no further possibility to communicate during or after the periods.

The following Figure 5.8 gives a schematic summary of part 2 of the experiment.

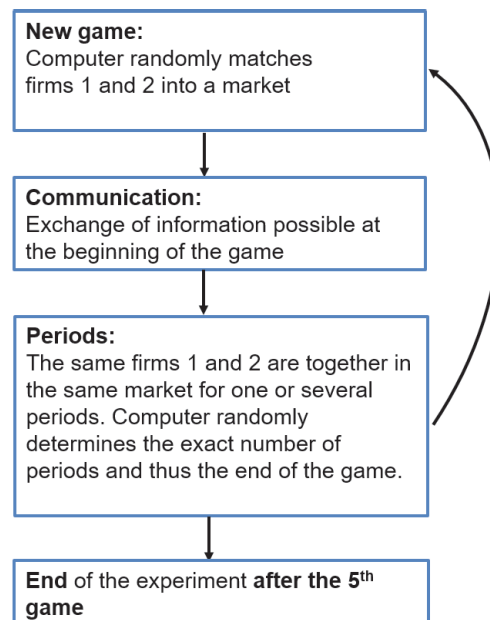


Figure 5.8: Schematic summary

Chapter 6

Conclusion

In this thesis, four articles on international trade and industrial organization have been presented. Three of them are concerned with the organization of firms in the multinational context. The forth article contributes to the literature on industrial organization by studying horizontal agreements between firms experimentally.

Chapter 2 investigates long-term relationships between headquarter corporations and their foreign suppliers. In an incomplete contracts environment we examine how informal relational contracts can be used to overcome productive inefficiencies in the presence of search frictions. Our major contribution is to highlight how the possibility to switch suppliers crucially affects the contractual nature of these buyer-supplier relationships. The feasibility of an efficiency-enhancing relational contract depends essentially on two components. First, the participants must be on average patient enough such that their incentive constraints can be met. Second, a relational contract may only become possible once the headquarter stops searching for a preferable partner and thus launches a long-term collaboration with her current match. We apply Chinese firm-product-level customs data and find evidence that confirms many of the predictions of our theory.

We employ the same data set to address the empirical relevance of the theory developed in the third chapter, which complements the analysis of Chapter 2 in several dimensions. This essay focuses on examining the firm's search and matching problem: We generalize the former theory to a continuous distribution of supplier types and elicit the role of uncertainty when the firm searches for a new partner.

In Chapter 4, I examine export relationship between firms under incomplete contracts. It is the case for many export transactions that shipping goods from one country to an other costs time. Since the payment morale of transaction partners can ex-ante often not be properly assessed but only be inferred from previous interactions, exporters are forced to care about the selection of appropriate payment contracts. The main finding of the article is that by choosing the payment contract for a transaction in an export relationship, the exporter is challenged with a trade-off between maximizing the profits from his current transaction and making use of screening opportunities that differ between format types. The article also studies the benefits of trade credit insurances that can help to foster export activities in early stages of export relationships.

Finally, the fifth chapter investigates firm relationships through a market experiment. We study the role of explicit communication for collusion in asymmetric Cournot markets where, from a theoretical perspective, the quality of coordination is highly unclear. Our experiment shows that when explicit communication is available total payoffs increase above the competitive level, however, game-theoretic concepts do not predict the outcome well: Nearly all the gains from talking go to the inefficient firm. When the role of the efficient firm is earned in a contest, the efficient firm earns higher profits and firms often collude by producing equal amounts. Without communication, firms fail to collude and essentially play the static Nash equilibrium.