Collaboration in Supply Chains: Design and Effects of Non-Contractual Mechanisms

by

Ana Ruth Beer

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2015

Doctoral Committee:

Professor Hyun-Soo Ahn, Co-Chair
Assistant Professor Stephen G. Leider, Co-Chair
Assistant Professor Erin L. Krupka
Assistant Professor Scott I. Rick

ACKNOWLEDGEMENTS

I am deeply grateful to my advisors, Hyun-Soo Ahn and Stephen Leider, for all their guidance and support throughout the years. They spent long hours working with me and teaching me, they shared my happiness when I was making progress, and they had patience and words of encouragement when things were hard. I feel extremely lucky for having them as advisors, and I cannot imagine transitioning this path without them by my side. I would also like to thank my dissertation committee members, Erin Krupka and Scott Rick for their helpful comments and contributions to my dissertation. Additionally, I would like to thank Roman Kapuscinski, Damian Beil, Wally Hopp, William Lovejoy, and Ravi Anupindi, who contributed to my learning in our departmental seminars and with valuable research feedback.

I am grateful to my friends – cohort mates, office mates, and my senior and junior colleagues in the PhD program – and to my Ann Arbor family, Yao Cui, Anyan Qi, Hakjin Chung, Evgeny Kagan, Remy Elbez, and Tiffany Vu; who were a wonderful companionship in this journey. And finally, my infinite gratitude to my parents, Noemi and Kike, my siblings, Deby and Ale, and my aunt Dorthy; for their unconditional love and support, and for always believing in me.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

<u>**Table**</u>

LIST OF APPENDICES

Appendix

ABSTRACT

Collaboration in Supply Chains: Design and Effects of Non-Contractual Mechanisms

by

Ana Ruth Beer

Co-Chairs: Hyun-Soo Ahn and Stephen Leider

As many companies and organizations gain global presence, buyer-supplier relationships become a very important topic in Operations Management. From both buyers and suppliers perspectives, the success of a supply chain relies on how well these relationships are managed. Contracts and mechanisms which are based on non-cooperative game theoretic models (e.g., zero-sum games), often result in poor outcomes such as poor quality and non-conformance, and hurt buyers and suppliers instead of helping them.

Building on game-theoretic frameworks, earlier work in this area has focused mostly on designing contracts that can achieve coordination of the supply chain. In practice, however, not all important aspects of a relationship can be contemplated in a contract. For example, desired quality or service level may be hard to specify. The suppliers expected reaction in the case of an unforeseen event, like a natural disaster, may also be hard to predetermine in advance. It is particularly in these cases when the nature and continuity of a relationship matter the most. My research focuses on non-contractual aspects of buyer-supplier relationships. I develop behavioral models to analyze industry practices that enhance collaboration in a supply chain, and then

test the theoretical models with laboratory experiments.

My dissertation explores actions that can be taken by buyers and suppliers to improve the relationship and promote a more efficient supply chain. The next three chapters answer questions that are important in understanding and designing successful buyer-supplier relationships: How can a buyer identify trustworthy suppliers? How should the buyer reward good suppliers? In which cases and how should a company invest in developing a long-term relationship with its suppliers? I show that higher profits and efficiency can be obtained when 1) suppliers make an upfront buyer-specific investment to signal that they are trustworthy, 2) buyers reward good suppliers with private symbolic awards, and 3) buyers allocate decision rights to long-run focused employees incentivizing suppliers to share with them their innovations.

CHAPTER I

Can Trustworthiness in a Supply Chain be Signaled?

1.1 Introduction

The relationship between a buyer and its suppliers is vital in almost every aspect of business. The operations management literature has explored in depth the problem of designing the optimal contracts for buyer-supplier relations. In many business contexts, however, it is not possible to describe every important aspect of the transaction in a contract. For example, desired quality or service level may be hard to specify (*Kaya and Özer* 2009). The supplier's responses to disruptions from unforeseen events, such as a natural disaster, may also be hard to determine. When a supplier fails to fulfill its obligations, the buying firm can suffer greatly. For example, Toyota's accelerator pedal quality problems in 2010 (due in part to supplier misbehavior) cost the company nearly two billion dollars and a significant decline in market share. To prevent such outcomes, many firms invest in identifying and maintaining good relationships with their business counterparts (e.g., suppliers, buyers).

In a supply chain setting, a relationship with a trustworthy supplier often results in significant benefit for a buyer. *Morgan and Hunt* (1994) find that when both commitment and trust are present in the buyer-supplier relationship it leads to increased efficiency, productivity and effectiveness. *Piboonrungorj and Disney* (2012) studied supplier relationships in the tourism industry and found that higher levels of inter-firm trust lead to better logistics performance. *Doney and Cannon* (1997) empirically found a positive correlation between the buying firm's trust in a supplier and the supplier's willingness to make relation-specific investments. A recent initiative by General Motors (GM) to establish strategic supplier relationships that the authors were involved in led to an improvement in the relationship with a key supplier of fascia, ultimately leading to the supplier building a new dedicated production facility. Often, supplier trustworthiness is demonstrated by the behavior of suppliers in areas not covered by the contract. Many buyers explicitly attempt to encourage this "above and beyond" behavior. Many companies including Delphi, Verizon, and AT&T have established outstanding supplier awards for the suppliers that go above and beyond their performance objectives. They reward their suppliers' efforts in terms of creative cost-reduction solutions, teamwork, customer service, response to natural disaster, sustainability, and social responsibility. For instance, a major store chain, Costco, states in its official Supplier's Code of Conduct that it encourages its suppliers to work to achieve *above and beyond goals* in excess of legal workplace requirements.

A standard argument for the emergence of a collaborative relationship under an incomplete contract is that long-term relationships between buyers and suppliers and concerns about reputation will limit opportunistic behavior. That is, if the expected long-term benefits of good behavior outweigh the immediate gratification of engaging in opportunistic behavior, then self-interested suppliers will perform collaboratively even in areas where the contract is silent. While relational and reputational incentives are certainly important, there are many cases where the incentives they provide are absent or insufficient to fully explain behavior. For example, many transactions are difficult for outsiders to monitor so that reputational incentives can steer supplier's behavior. The transactions may also be inherently one-time exchanges that fail to induce relational incentives. In these cases it is important for a buyer to identify suppliers that are trustworthy before signing a contract.

Trust can be generally defined as "a psychological state comprising the intention to accept vulnerability based upon positive expectations of the intentions or behavior of another" (*Rousseau et al.* 1998) [1]. In the context under study, a trusting buyer offers a generous price to a supplier when quality is non-contractible. A trustworthy supplier provides high quality when he was trusted with a high price. Thus, in our setting, the concept of trustworthiness is closely related to reciprocity: if a supplier is treated generously by the buyers (e.g., if they are offered a high price), it will reciprocate the gesture by, for example, providing high quality products. While trustworthiness has been observed in a number of settings among individuals, it is also a relevant characteristic in describing transactions between firms. As *Morgan and Hunt* (1994) and *Piboonrungorj and Disney* (2012) show, firms differ significantly in this dimension and transactions between firms with higher trustworthiness benefit both parties - suggesting that trustworthiness is an intrinsic attribute that can separate firms. A natural question is what leads suppliers to be inherently trustworthy. Why do firms differ in trustworthiness? One possible explanation could be that firms, regardless of their size or industry, execute their interactions with other firms through individuals. This is the case at GM, a large company with over two hundred thousand employees, where the relationship with each supplier is managed by a handful of people interacting with a small number of individuals from the supplier. Thus, these individuals' preferences and behavior will influence the relationship that develops between the firms. Another possible explanation is that a firm's trustworthiness is a deliberate business strategy supported by the firm's corporate culture. A firm's culture often reflects how the leadership would like employees to interact with cus-

[1] This definition is used by *Özer et al.* (2011) in their study of trust in forecast information sharing. In their context, a trusting supplier relies on the forecast information provided by the manufacturer to make a capacity decision and a trustworthy manufacturer is that who reports forecast information truthfully.

tomers and business partners. Some firms who are known to be "excellent" suppliers may try to cultivate cultures where "above and beyond" actions are rewarded. In such cases, reliability and trustworthiness will be more common in their relationships with other firms.

As suppliers may inherently differ in trustworthiness, an important question for buyers is how to identify these "good" suppliers before contracting. In this paper we propose that observing pre-contracting behavior by the supplier is one way to discern good suppliers from selfish ones. In particular, we argue that early buyer-specific investments by a supplier may signal that firm's trustworthiness. Relationship-specific investments are costly for suppliers, as investing in one particular buyer will tend to weaken the supplier's outside option and make the supplier vulnerable in negotiations with the buyer. However, these kind of relationship-specific investments are not rare. *Ganesan* (1994), in a study of buyer-supplier relationship in regional department store chains, found that transaction-specific marketing investments were quite common, including training the retailer's salesforce, developing product displays, providing dedicated electronic linkups for inventory control and offering information on new products. We found other examples through our own professional interaction with firms. Cosmax, an original design manufacturer which serves several of the world's largest cosmetics companies, invested in a buyer-specific equipment which, at the time, was only recommended by L'Oreal without even having a contract from L'Oreal. An Argentinean clothing manufacturer, Kayene, hired a dedicated quality assurance team to serve a specific retailer prior to having an agreement or a written contract with that firm.[2] In both cases it was feasible for the firms to make these investments after securing the contract. Why then would these firms make the costly investments in advance? We hypothesize that these kind of investments can be used as signals that trustworthy suppliers send to help the buying firm to discern trust-

[2]The authors worked with Cosmax and Kayene for several years.

worthy suppliers from selfish suppliers. While Cosmax anticipated that building trust with L'Oreal would lead to a long term relationship, in the case of Kayene the buying firm's objectives were focused mainly on the short run due the high volatility of the Argentinean economy, which makes future interactions highly unpredictable.

These motivations suggest several research questions that we address in this paper: What are the benefits of developing a more collaborative supplier relationship? Is it possible to identify trustworthy suppliers before contracting with them? If so, under what circumstances is it possible? Do these benefits persist in long term relationships? We hypothesize that the buyer can distinguish between trustworthy and untrustworthy suppliers based on the suppliers pre-contract investments. Specifically, we expect that suppliers who make a buyer-specific investment will be more likely to deliver higher non-contractible quality, leading to higher profits for both firms. These results should be further accentuated when firms have expectations of establishing a long term relationship.

To formalize this intuition, we develop a model in which a trustworthy supplier can make a relationship-specific investment (instead of a general investment) to signal his type to the buyer. The buyer then offers a price, and the supplier makes a non-contractible effort that determines product quality. We identify cases where a trustworthy supplier chooses the buyer-specific investment while the selfish supplier chooses the general investment in a separating equilibrium. Under this equilibrium trustworthy suppliers receive higher prices, and exert higher effort.

We test these predictions using an experimental supply chain game. Our results show that the specific investment leads to significantly higher prices and quality, and increases the profits for both the buyer and the supplier. Furthermore, the investment choice reflects persistent individual differences, with different subjects showing a preference for one investment over another. We show that there is a positive correlation between the suppliers' preference for the specific investment and their level of reciprocity. Hence, the investment choice represents an accurate signal of the underlying type of the supplier. We demonstrate that the signaling mechanism is essential in generating the benefits of the specific investment. In an additional treatment where the signaling mechanism is eliminated (by randomly assigning investments), the buyer specific investment no longer leads to higher quality or increased profits.

Finally, we analyze the case where firms interact repeatedly through several transactions after the supplier's investment decision, representing a (finite) long-term relationship between the buyer and the supplier. The existing literature shows that repeated interactions have a positive impact on trust and trustworthiness. *Özer et al.* (2011) find that repeated interactions further promote cooperation in forecast information sharing. Empirical research by *Doney and Cannon* (1997) find that concerns about reputation reinforce trust and that developing trusting relationships represents an investment for the long run. This suggests that repeated interactions could enhance the benefits of buyer-specific investments in presence of reciprocal suppliers. In the absence of reciprocal suppliers, equilibria with collaborative outcomes can never be supported with finitely repeated interactions [3]. However, in the presence of reciprocal suppliers, we characterize two different collaborative equilibria. First, as with the one-shot interaction, an equilibrium exists where a reciprocal supplier chooses the specific investment and a selfish supplier chooses the general investment. In another collaborative equilibrium, the selfish supplier mimics the reciprocal suppliers by

[3] In infinitely repeated interactions it is a well-know game-theoretical result that, if players care enough about the future, collaborative outcomes can be sustained in equilibrium. On the other hand, in finitely repeated interactions these equilibria usually fall apart. Through backwards induction, players know that their counterpart will defect in the last period, so this breaks apart collaborations in previous periods. However, previous research has shown that outcomes that are not equilibria of the single shot game can be equilibria of a finitely repeated game in certain circumstances, as in the case of incomplete information. For example, *Kreps and Wilson* (1982) show that reputation building can be an equilibrium in a finitely repeated version of Selten's finitely repeated chain-store game. Similarly, *Kreps et al.* (1982) show that reputation effects due to informational asymmetries can generate cooperative behavior in finitely repeated Prisoner's Dilemma, where "finking" at each stage is the only Nash equilibrium of the finitely repeated game. We show that the existence of reciprocal suppliers (and asymmetric information about the suppliers' type) allows for collaborative outcomes to arise in a finitely repeated game.

choosing the specific investment and offering high quality in all transactions except the last one. In both equilibria the specific investment generates higher effort and a greater surplus that the general investment, and, compared to the one-transaction game, the benefits of the specific investment over the general investment are magnified by the repeated interactions. To test these results, we conducted a new treatment with one investment decision and three subsequent trading periods. We find that three trading periods are enough to significantly increase the efficiency of the specific investment. Prices and effort under the specific investment are significantly higher than in the single interaction case. As a result, the profit premiums of the specific investment are significantly more prominent with repeated interactions for both buyers and suppliers.

1.2 Literature Survey

Improving buyer-supplier relations can lead to important performance gains, including enhanced supply-chain responsiveness (as a result of reduced cycle times) and higher profits (*Handfield and Bechtel*, 2002). Additionally, these relationships can benefit from the parties' willingness to make relationship-specific investments (*Dyer and Singh*, 1998). For example, *Asanuma* (1989) and *Dyer* (1996a) have shown that firms can derive improved performance and competitive advantages when relationship-specific investments were made. However, relationship specific investments also present problems. Hold-up problems arise from the fact that, once a party has made a specific investment, the other party has an incentive to be opportunistic. In many cases it is difficult to prevent such opportunism contractually, which may necessitate vertical integration to promote efficiency (see *Williamson* 1971, *Williamson* 1975, *Klein et al.* 1978, *Williamson* 1979, and *Grossman and Hart* 1986 for theoretical work; see *Monteverde and Teece* 1982, *Masten* 1984, and *Joskow* 1985 for empirical work).

Alternative mechanisms that can limit the scope for opportunism are long run relationships and the importance of firm reputation (*Larson* 1992, *Baker et al.* 2002, *Gibbons* 2005). There is experimental evidence of this in the Operations Management literature. *Özer et al.* (2011) find that trust and cooperation can be reinforced by reputation concerns in the context of forecast information sharing. *Heinrich and Brosig-Koch* (2011) find that when buyers can consider the reputation of bidders in procurement auctions, bidders supply higher quality leading to higher market efficiencies [4]. We consider first a setting where complete contracting, integration and relational incentives are not present to provide clear and direct evidence for the importance of trustworthiness. Our results also apply to settings where these factors may be present, but insufficient to incentivize proper behavior by the supplier. Then, we consider the case where firms interact repeatedly. This allows us to examine the role of reputation concerns on trust in the context of relationship-specific investments with hold up problems.

The importance of trust and trustworthiness has been demonstrated in a variety of settings. *Berg et al.* (1995) provide early experimental evidence on the importance of trust and trustworthiness in investment decisions. *Glaeser et al.* (2000) demonstrate that trust and trustworthiness reflects both past actions and beliefs about others. *Kosfeld et al.* (2005) show that trust has a biological basis. Trust varies between countries (*Bohnet et al.* 2008), often depending on culture and institutions (*Bohnet et al.* 2001, *Johnson et al.* 2002). The level of trust in a country has significant effects on the rate of economic growth (*Zak and Knack* 2001), as many economic transactions require trusting the other party. *Özer et al.* (2011) find evidence that trust and trustworthiness allow for cooperative forecast information sharing in a supply chain context. In cross-country supply chains, trust, trustworthiness, and strategic

[4] More recent research by Haruvey, Katok, Ma and Sethi (2014) also focuses on the effects of reputation on the provision of quality. They conduct an experimental study of the role of reputation building when a seller makes non-contractible effort towards the production of a good.

information sharing is affected by the country of origin of the supply chain members (*Özer et al.* 2014).

Trustworthiness is often modeled as a preference for equity or reciprocity (*Rabin* 1993, *Fehr and Schmidt* 1999, *Bolton and Ockenfels* 2000, *Dufwenberg and Kirchsteiger* 2004), an approach that we follow. *King-Casas et al.* (2005) demonstrate that reciprocal actions lead to future trust and trustworthiness. Reciprocal motives have been demonstrated experimentally in a labor market setting where higher wages lead to higher effort (*Fehr and Falk* 1999), and a buyer-seller transaction setting where higher prices lead to higher quality goods (*Fehr et al.* 1993). Mutual reciprocity/trustworthiness (as in our setting) is a particularly powerful way of addressing problems of contractual incompleteness (*Fehr et al.* 1997). Since reciprocal counterparties are more profitable, it is a natural question how one might identify a reciprocal individual from a selfish one. A few studies have looked at signals from outside the transaction, such as the image of the other person's face (*Scharlemann et al.* 2001) or information about past charitable giving (*Fehrler* 2010). This paper, however, focuses on the role of investments within the context of the transaction as a potential signal of trustworthiness.

The field of Operations Management has produced a vast literature on buyer-supplier relationships. Most papers in this category focus on designing optimal contracts or comparing contracts in different settings. This attention towards contracting problems stems from the challenge of coordinating each firm's objective with that of the supply chain, particularly due to double marginalization (*Cachon* 2003, *Spengler* 1950). The most usual setting for these problems is the newsvendor model (*Silver et al.* 1998) for which different types of contracts have been explored, including wholesale price (*Lariviere and Porteus* 2001; *Bresnahan and Reiss* 1985), buy-back (*Pasternack* 1985) and revenue sharing contracts (*Cachon and Lariviere* 2005). Rather than investigating the quantity decision, we focus on non-contractible aspects of buyer-supplier relations such as effort and quality. While other have studied incentive problems relating to non-contractible capacity investments (*Tomlin* 2003) or product quality (*Kaya and Özer* 2009) we are unaware of other papers that examine investment as a signaling mechanism in this context.

In the behavioral operations literature, contracting theories in buyer-supplier interactions have been tested experimentally and revised to account for social preferences or decision biases, beginning with *Schweitzer and Cachon* (2000).[5] Several papers have identified concerns for fairness as an important influence on supply chain performance (*Cui et al.* 2007, *Pavlov and Katok* 2009). *Özer et al.* (2011) studies the importance of trust and trustworthiness in sharing forecast information within a supply chain. *Loch and Wu* (2008) find that forming a relationship prior to a transaction leads both parties to take more collaborative actions. *Brinkhoff et al.* (2014) show that trust is a strong predictor of supply chain project success when mediated by project-level factors. *Cui and Mallucci* (2010) study how investment decisions are affected when the retailer can have distributive fairness concerns with respect to the manufacturer. We identify a specific action that buyers can take beforehand which can lead to more collaborative relationships. We model a situation where the supplier can signal its type by making a relationship-specific investment before the buyer offers a contract and propose that this signal allows the buyer to screen for reciprocal suppliers, which turns out into a more collaborative and profitable relationship between the parties.

1.3 Theoretical Model

We consider a three-stage game in which a buyer trades with a single supplier for a non-divisible good. In the first stage the supplier makes a pre-contractual investment.

[5]See also *Bolton and Katok* (2008), *Becker-Peth et al.* (2011), *Katok and Wu* (2009) and *Ho and Zhang* (2008).

The buyer observes this and offers a take-it-or-leave-it price offer in the second stage. Then, the supplier decides if he should accept the offer, and, if so, how much effort he will exert towards generating quality. The buyer's value of the good depends on the good's quality, which depends on the supplier's non-contractible effort and his investment choice.

At the beginning of the game, the supplier needs to choose between two different investment options: a *general investment* (denoted by g) and a *buyer-specific investment* (denoted by b). We assume that the firm has the resources available to make one investment, and that either investment would be a better use of capital than the alternatives - hence choosing one investment is the optimal decision [6]. Both options have equal financial cost, however they benefit the supplier in different ways. The *general investment* directly increases the supplier's outside option value (i.e., the reservation utility) which is the monetary value the supplier receives when both parties cannot strike a deal. Since the buyer has to compensate at least the outside option value in order to close a deal, the general investment benefits the supplier by improving the supplier's bargaining power. Examples of the general investment include industry standard certification (e.g., ISO 9000), building a multi-purpose automated production line, and increasing the capability and man-power in B2B marketing. On the other hand, the *buyer-specific investment* will increase the value of the good for the buyer for a given effort choice of supplier. This investment will benefit the supplier only if the buyer shares the increased value created by the supplier's investment and

[6] We considered an alternative model where the supplier makes a decision i from two options: to invest ($i = I$) or not to invest ($i = NI$). If he invests, the quality coefficient is α_I and if he does not invest it is α_{NI}, with $\alpha_I > \alpha_{NI} > 0$. Under both decisions the outside option remains $\bar{u}_I = \bar{u}_{NI} = \bar{u} > 0$ and making the investment has a fixed cost K. We assume in this case, that a reciprocal supplier considers an offer to be generous if the price not only compensates him for his outside option but also for his investment cost, K. We find that there is no set of parameters under which a Separating Equilibrium can arise in this model. In particular, the interesting Separating Equilibrium in which the selfish supplier chooses not to invest and the reciprocal supplier chooses to make the specific investment cannot happen. This is because it is never incentive compatible for the selfish supplier not to invest for two reasons: first, because the price offered to suppliers who invest is too high since it needs to compensate for K, and second, not investing does not raise the supplier's outside option.

effort through the take-it-or-leave-it price. Examples of the buyer-specific investment include purchasing a buyer-specific machine or fixture, adopting a higher quality standard that is only requested by a particular buyer, or hiring a team for a specific buyer. Note that we are considering the case where the supplier is already about to make an investment since making either one of the investments is profitable (better than not investing at all). Thus, the supplier has already incurred in the initial cost of investing, which will be considered sunk cost. However, since the supplier can only choose one investment, the buyer-specific investment has an opportunity cost - the supplier must forgo the chance to increase his outside option.

To formally capture this, we assume that, if the supplier chooses investment i ($i = g$, or b) and exerts effort, e, the value of the good the buyer receives is $\alpha_i e$ where $\alpha_b > \alpha_g > 0$. In other words, for given effort level, e, the supplier who chose the buyer-specific investment provides a higher quality, and hence a higher value to the buyer, ($\alpha_b e$) than the supplier who chose the general investment ($\alpha_g e$). We will refer to α_i as the *quality coefficient* from now onwards. We assume that the outside option value from the general investment ($\bar{u}_g$) is higher than that from the specific investment ($\bar{u}_b$): $\bar{u}_g > \bar{u}_b \geq 0$.

After observing the supplier's investment, the buyer offers a take-it-or-leave-it price offer, p, to a supplier. In the final stage of the game, the supplier evaluates the contract and determines whether to accept the buyer's offer or not. If the supplier rejects the buyer's offer, the buyer receives zero payoff and the supplier receives the outside option value ($\bar{u}_g$ or $\bar{u}_b$), depending on the supplier's pre-contract investment. If the supplier accepts the contract, the supplier then chooses an effort level, e, which incurs cost $c(e)$, which we assume to be strictly increasing and strictly convex in e.

We assume that there are two types of suppliers–*selfish* and *trustworthy*– in the market place. The selfish supplier cares about his own monetary payoff exclusively. Thus, he only aims to maximize its own profit. If the selfish supplier with investment type i accepts the buyer's price offer (p) and chooses effort level e, his utility is simply his monetary payoff and is defined as follows

$$U^s(e|i,p) = p - c(e). \tag{1.1}$$

When offered a contract, the selfish supplier will compare the maximum utility he can receive from accepting the buyer's offer to his outside option ($\bar{u}_i$), and will choose the option that yields a higher monetary payoff.

On the other hand, the trustworthy supplier differs from the selfish supplier by having reciprocal preferences. If the buyer's offer is sufficiently generous, the trustworthy supplier's utility depends on both total supply chain profits as well as his own monetary payoff. To capture this, let $\gamma > 0$ be the minimum premium that the trustworthy supplier needs to receive in order to perceive that the buyer's offer is generous. If the buyer's offer to the supplier with investment type i is not generous, that is, $p < \gamma + \bar{u}_i$, then the trustworthy supplier will act selfishly and will maximize his monetary payoff, $p - c(e)$. On the other hand, if the buyer's offer, p, is generous, then the supplier with investment i who accepts the contract will maximize a utility function that accounts for both his monetary payoff and the total surplus of the supply chain: $(1-\phi)[p-c(e)]+\phi[\alpha_i e - c(e)]$ for some $\phi \in [0,1]$. We define ϕ to be the coefficient of reciprocity, which represents the degree of the supplier's reciprocity toward the buyer. Note than when $\phi = 0$, then this payoff is identical to that of the selfish supplier. On the other hand, when $\phi = 1$, the supplier becomes totally altruistic and interested in maximizing the total surplus. Thus, the higher ϕ is, the more reciprocal the supplier is. This notion of reciprocity is similar in spirit to *perceived kindness* used in *Rabin* (1993) in simultaneous move games, and *Dufwenberg and Kirchsteiger* (2004) and *Falk and Fischbacher* (2006) in sequential games, or *inequity aversion* used in *Fehr and Schmidt* (1999). However, our model is a stylized simplification of other models of reciprocity in two respects. First, our reciprocal supplier cares about total surplus, rather than the buyer's profit. Surplus maximization is more intuitive and prevents inefficient over-provision of quality, which in the context of buyer-supplier relations would be unrealistic. [7] Second, reciprocity is binary, depending on whether the offer is sufficiently generous. [8]

Combining these two cases, the utility that the reciprocal supplier with investment i gains when he accepts the buyer's price offer, p and exerts an effort level, e is

$$U^r(e|i,p) = \begin{cases} p - c(e) & \text{if } p - \bar{u}_i < \gamma \\ (1-\phi)(p - c(e)) + \phi(\alpha_i e - c(e)) & \text{if } p - \bar{u}_i \geq \gamma \end{cases} \tag{1.2}$$

The supplier compares the maximum utility that he can receive from accepting the offer and the outside option ($\bar{u}_i$), and chooses the option with a higher value.

The buyer's utility from offering a price, p to the supplier with investment type i ($i = b$ or g) is [9] [10]

[7] If the surplus is replaced by buyer's profit (with the adjustment that the reciprocity coefficient, ϕ, needs to range between $[0, \frac{1}{2}]$) behavior does not change.

[8] This simplification provides modeling tractability, however none of our main results depend on this assumption. The assumption is similar to that in *Englmaier and Leider* (2012) where, in a principal-agent context, a "generous" contract is one that provides the agent with an expected monetary utility in excess of his outside option. We consider a binary version of that model and introduce the additional individual-specific parameter γ, which reflects how generous the offer needs to be.

[9] We consider the simpler case where the buyer is modeled as selfish, which is sufficient to derive separating equilibrium results. Because the supplier moves last, if the supplier is reciprocal then even a selfish buyer has strategic reasons to offer a high price. A reciprocal buyer would have an even greater incentive to offer high prices to suppliers choosing the specific investment, strengthening our results. This setting is similar to *Englmaier and Leider* (2012) where, in a principal-agent context, the agent is modeled as reciprocal and the principal as selfish when solving for the optimal contract.

[10] We consider a setting where under investment i the buyer has an outside option w_i if the deal does not close. This setting favors the buyer in terms of the allocation of surplus however all our other main results hold with minor changes. In the full information case, there is an upward shift in the threshold for conditions C_b and C_g. When the buyer has an outside option, the buyer requires a higher quality coefficient in order to offer a reciprocal contract to a reciprocal supplier. Otherwise, offering a reciprocal contract is not worthwhile. Similarly, in the asymmetric information case, a new condition is necessary for the separating equilibrium to arise. We need the buyer's outside option not to be too high so that it is incentive compatible for the buyer to offer a reciprocal contract under the specific investment. Specifically, we need $w_b \leq \frac{c(c'^{-1}(e_b^{r*}))}{\phi}$.

$$U^B(p|e,i) = \begin{cases} \alpha_i e - p & \text{if the supplier accepts the offer and exert an effort level, } e \\ 0 & \text{if the supplier rejects the offer.} \end{cases} \tag{1.3}$$

We first study the full information case, in which the supplier's type is common knowledge and we then study the case where the supplier's type is private information.

1.3.1 Full Information Case

We begin by analyzing the case where the buyer has full information about the supplier type – *trustworthy* or *selfish*– as a benchmark. We first characterize the supplier's action in the third stage: whether the supplier should accept the buyer's offer and, if so, how much effort he should exert. We then apply backward induction and analyze the buyer's offer problem (2nd stage) and the supplier's choice of pre-contractual investment (1st stage).

In the third stage, a supplier decides between accepting the buyer's offer and rejecting the offer for an outside option. If the supplier accepts the offer, he must decide how much effort he exerts. We first consider a selfish supplier who chose type-i investment in the first stage and received the buyer's offer, p. If he rejects the offer, then he would receive the utility from his outside option, $\overline{u}_i$. If he accepts the offer, from (1.1), it is easy to observe that the optimal effort for the selfish supplier is always zero regardless of the price, p.

Now, consider a trustworthy supplier with type-i investment. As in the selfish supplier case, if the reciprocal supplier rejects the offer, he earns his outside option, $\overline{u}_i$. On the other hand, if he accepts, his optimal effort depends on whether he perceives the buyer's contract to be generous. If $p - \overline{u} < \gamma$, then the offer is not considered to be generous. Thus, the supplier will act selfish and will maximize the utility function, $p - c(e)$ by exerting zero effort. If $p - \overline{u} \geq \gamma$, then the supplier finds the offer generous. Then, his best effort is derived from the following optimization problem:

$$\max_{e \geq 0} \quad (1-\phi)(p - c(e)) + \phi(\alpha_i e - c(e)) \quad \text{ s.t. } p - \bar{u}_i \geq \gamma.$$

The solution to this problem is $c'(e) = \phi\alpha_i$. Note that, because c(e) is a strictly increasing convex function of e, $c'(e)$ is always positive, increasing in e, and invertible. Additionally, since $c'(e)$ is strictly increasing in e, $c'^{-1}(\phi\alpha_i)$ is also increasing. As a result, the solution to the above problem can also be written as $e^{t*}(p,i) = c'^{-1}(\phi\alpha_i)$. After combining both cases, it can be shown that the trustworthy supplier's optimal effort, denoted by $e^{t*}(p, \hat{e}, i)$, is

$$e^{t*}(p,i) = \begin{cases} 0 & \text{if } p < \bar{u}_i + \gamma \\ c'^{-1}(\alpha_i\phi) & \text{otherwise.} \end{cases} \tag{1.4}$$

We then compare the two options– accepting and rejecting– and characterize the supplier's optimal action in the following lemma.

Lemma I.1. *Consider a supplier who chose type-i investment and faces the buyer's price offer, p.*

(i) If $p > \bar{u}_i$*, the selfish supplier accepts the offer and exerts zero effort:* $e^{s*}(p,i) = 0$*. If* $p \leq \bar{u}_i$*, he rejects the offer and earns* $\bar{u}_i$*.*

(ii) If $p \geq \gamma + \bar{u}_i$*, the trustworthy supplier accepts the offer and exerts* $e^{t*}(p,i) = c'^{-1}(\alpha_i\phi)$*. If* $\bar{u}_i < p \leq \gamma + \bar{u}_i$*, he accepts the offer and exerts zero effort,* $e^{t*}(p,i) = 0$*. If* $p \leq \bar{u}_i$*, he rejects the buyer's offer and earns* $\overline{u}_i$*.*

Lemma I.1.(i) implies that the selfish supplier will never choose strictly positive effort. Since the buyer will never earn positive profit from a selfish supplier, it is optimal for the buyer to offer $p = 0$, and induce the supplier to reject. [11] On the

[11] Although any price $p < \bar{u}_i$ can be an equilibrium, we focus on the case of $p = 0$ for expositional

other hand, facing the trustworthy supplier, the buyer must compare the two options – offering a generous contract that makes the supplier exert strictly positive effort and offering a very low offer so that the supplier rejects the contract. In order to characterize the optimal offer, we denote an offer, $p = 0$, as a *null contract.* Similarly, if the buyer offers $p = (\bar{u}_i + \gamma)$ to the supplier with type-i investment, we call this a *trusting contract* and denote by T_i, $i = b, g$.

In preparation for our following Lemma, let us define $\underline{\alpha_i} = min\{\alpha_i | \alpha_i(c'^{-1}(\alpha_i \phi)) - \bar{u}_i - \gamma \geq 0\}$ as the minimum value of α_i that satisfies $\alpha_i(c'^{-1}(\alpha_i \phi)) - \bar{u}_i - \gamma \geq 0$. Because $c'^{-1}()$ is strictly increasing in α, there exists some threshold $\underline{\alpha_i} > 0$ above which the buyer finds it profitable to offer a reciprocal contract.

Lemma I.2. *Suppose the supplier chose type-i investment in the first stage. Then, offering the null contract is optimal when the buyer faces either the selfish supplier or the trustworthy supplier with low quality coefficient: $\alpha_i \leq \underline{\alpha_i}$. Offering the trusting contract, T_i, is optimal if the buyer faces a trustworthy supplier with high quality coefficient, $\alpha_i > \underline{\alpha_i}$.*

Lemma I.2 implies that the buyer offers a trusting contract to the supplier when the supplier can provide sufficiently high value when type-i investment is made: $\alpha_i(c'^{-1}(\alpha_i \phi)) - \bar{u}_i - \gamma \geq 0$. Rewriting the condition for both types of investment, the condition in Lemma I.2 can be expressed as follows:

$$\text{Condition } C_b: \quad \alpha_b \geq \underline{\alpha_b} \quad \text{and} \quad \text{Condition } C_g: \quad \alpha_g \geq \underline{\alpha_g} \tag{1.5}$$

Now consider the supplier's investment in the first stage. From Lemma I.2, the selfish supplier will receive the null contract no matter what he chose in the first stage. Since the supplier will always reject the null contract, it is optimal for the selfish supplier to choose the general investment to raise his outside option value to

purpose.

$\overline{u}_g$. On the other hand, the optimal action for the trustworthy supplier depends on which of the two conditions – C_b and C_g is met. Since $\alpha_b \geq \alpha_g$ and $\overline{u}_b < \overline{u}_g$, it suffices to consider the following three cases (the fourth case, condition C_g is met and C_b is not, cannot occur). The next result characterizes the equilibrium under the full information.

Theorem I.3. *In equilibrium, the following statements hold.*

a) The selfish supplier chooses the general investment, the buyer offers the null contract, and the supplier then rejects the offer.

(Parts b) to d) apply to the trustworthy supplier:)

b) Suppose that both C_b and C_g hold. If $(1-\phi)(\bar{u}_b+\gamma)-c(c'^{-1}(\alpha_b\phi))+\phi\alpha_b c'^{-1}(\alpha_b\phi) \geq (1-\phi)(\bar{u}_g+\gamma)-c(c'^{-1}(\alpha_g\phi))+\phi\alpha_g c'^{-1}(\alpha_g\phi)$, then the supplier chooses the buyer-specific investment, the buyer offers the trusting contract, T_b, and the supplier chooses the effort level: $e_b^{t*} = c'^{-1}(\alpha_b\phi)$. Otherwise, the supplier chooses the general investment, the buyer offers T_g, and the supplier chooses $e_g^{t*} = \alpha_g\phi$.

c) Suppose that only condition C_b holds. If $(1-\phi)(\bar{u}_b+\gamma)-c(c'^{-1}(\alpha_b\phi))+\phi\alpha_b c'^{-1}(\alpha_b\phi) \geq \bar{u}_g$, then the supplier chooses the buyer-specific investment, the buyer offers T_b, and the supplier chooses the effort level: e_b^{t*}. Otherwise, the supplier chooses the general investment, the buyer offers the null contract, and the supplier rejects the buyer's offer.

d) Suppose that neither C_b nor C_g holds. Then, the supplier chooses the general investment, the buyer offers the null contract, and the supplier rejects the buyer's offer.

1.3.2 Asymmetric Information Case

We now analyze the case where the supplier's type is private information. As in *Spence* (1973), we use the perfect Bayesian equilibrium (PBE) as our solution concept, imposing the restriction that the buyer's belief is consistent with the buyer's knowledge of the supplier's behavior in equilibrium. In particular, we characterize a separating equilibrium under which the supplier's investment acts as a signal. We also derive pooling equilibria in which neither supplier reveals his type.[12]

1.3.2.1 Separating Equilibrium

We first claim that the selfish supplier chooses the general investment and the trustworthy supplier chooses the buyer-specific investment in a separating equilibrium. To see why this must be the case, suppose that there exists a separating equilibrium in which the selfish supplier chooses the buyer-specific investment and the trustworthy supplier chooses the general investment. From Lemma I.2, the buyer will offer the null contract to the selfish supplier, who rejects the offer and earns the outside option payoff $\bar{u}_b$. Since $\bar{u}_g \geq \bar{u}_b$, the selfish supplier is better off by deviating and making a general investment, and this contradicts the equilibrium. We also note that, in a separating equilibrium, the buyer should offer the null contract to the selfish supplier and contract T_b to the trustworthy supplier. Consequently, the selfish supplier rejects the offer and the trustworthy supplier accepts the offer and exerts effort $e_b^{t*} = c'^{-1}(\alpha_b \phi)$.

We characterize a sufficient condition under which the separating equilibrium exists in the next lemma. In preparation, let θ be the real fraction of reciprocal suppliers in the marketplace, $\theta_j \in [0, 1]$ be the buyer's prior belief that the supplier's type is j, $j \in \{t = \text{trustworthy}, s = \text{selfish}\}$, and $\theta(j|i)$ be the buyer's updated belief about the supplier's type when the supplier chooses investment i, $i \in \{b, g\}$.

Theorem I.4. *There exists a separating equilibrium in which the selfish supplier chooses the general investment and the trustworthy supplier chooses the buyer-specific investment, the buyer offers the null contract to the selfish supplier and contract T_b to*

[12] Under certain conditions, semi-pooling equilibria may arise in which one type of supplier chooses a pure strategy and the other uses a mixed strategy when choosing the investment type. We focus on the separating and pooling equilibria as they are most relevant to our experimental results.

the trustworthy supplier, and the selfish supplier rejects the offer and the trustworthy supplier accepts the offer and exerts effort $e_b^{t*} = c'^{-1}(\alpha_b \phi)$*, resulting in* $\theta(t|b) = 1$ *and* $\theta(t|g) = 0$ *if and only if the following condition holds:*

i) $\bar{u}_g \geq \bar{u}_b + \gamma$,

ii) $(1-\phi)(\bar{u}_b + \gamma) - c(e_b^{t*}) + \alpha_b \phi e_b^{t*} \geq \bar{u}_g$.

The first condition guarantees that the selfish supplier's outside option is greater than what he would get by choosing the buyer-specific investment and exerting zero effort. The second condition guarantees that the trustworthy supplier's utility with the buyer-specific investment is greater than his profit when he mimicks to be selfish. We note that i) and ii) together imply $(1-\phi)(\bar{u}_b+\gamma)-c(e_b^{t*})+\alpha_b\phi e_b^{t*} \geq \bar{u}_b+\gamma$, which means that condition C_b holds. The result implies that pre-contractual investment can be a signal when $\overline{u}_g$ is high enough so that the selfish supplier is incentivized to choose the general investment, and, at the same time, α_b is high enough that fulfilling the buyer's contract is more attractive to the trustworthy supplier.

Under the buyer-specific investment, the supplier exerts effort e_b^{t*} and the buyer pays price $\bar{u}_b + \gamma$, so the buyer's profit is $\alpha_b e_b^{t*} - \bar{u}_b - \gamma$, which is greater than zero by condition C_b. Under the general investment the buyer earns zero profits. Under the buyer-specific investment, trustworthy suppliers earn a monetary profit of $\bar{u}_b+\gamma-c(e_b^{t*})$, and derive utility $(1-\phi)(\bar{u}_b+\gamma)-c(e_b^{t*})+\alpha_b\phi e_b^{t*}$ (note that, if everything is held constant, the trustworthy suppliers' utility increases more than their monetary profits as effort increases). Under the general investment, selfish suppliers earn $\bar{u}_g$. Condition i) in Theorem I.4 means that suppliers' monetary profits are higher under the general investment. Finally, total profits are $\alpha_b e_b^{t*} - c(e_b^{t*})$ under the specific investment and $\bar{u}_g$ under the general investment. Because of the convexity of $c(e)$, $\alpha_b e_b^{t*} - c(e_b^{t*})$ exceeds $\bar{u}_g$ if e_b^{t*} is large enough.

1.3.2.2 Pooling Equilibrium

In a pooling equilibrium, both suppliers will choose the same investment, thus the buyer is unable to discern the supplier type. In our setting, two pooling equilibria can exist– both types choosing the general investment and both types choosing the specific investment. To avoid a potentially large number of equilibria, we refine multiple equilibria with the *intuitive criterion* (*Cho and Kreps* 1987). The intuitive criterion states that for any belief the uninformed player may have after seeing a deviation, if one type of player receives a worse payoff by deviating than his equilibrium payoff and the other type does not, then the deviation should not be attributed to the player whose payoff decreases.

In the next result, we characterize three pooling equilibria that survive the intuitive criterion. In preparation, define a threshold $\tilde{\theta}_i = \frac{\bar{u}_i+\gamma}{\alpha_i c'^{-1}(\alpha_i\phi)}$ for $i = b$ and g. Since $\bar{u}_b \leq \bar{u}_g$ and $c'^{-1}()$ is increasing in α and $\alpha_b > \alpha_g$, then $\tilde{\theta}_g \geq \tilde{\theta}_b$.

Theorem I.5. *There are three pooling equilibria that survive the intuitive criterion.*
a) If $\theta \geq \tilde{\theta}_b$, a pooling equilibrium under which both suppliers choose the buyer-specific investment arises. In this equilibrium, the buyer offers a trusting contract T_b, and both suppliers accept the offer, the selfish supplier exerts zero effort and the trustworthy supplier exerts effort $e_b^{t*} = c'^{-1}(\alpha_b\phi)$.
b) If $\theta \geq \tilde{\theta}_g$, a pooling equilibrium under which both suppliers choose the general investment arises. In this equilibrium, the buyer offers a trusting contract T_g, both suppliers accept the offer, the selfish supplier exerts zero effort and the trustworthy supplier exerts effort $e_g^{t*} = c'^{-1}(\alpha_g\phi)$.
c) If $\theta < \tilde{\theta}_g$, a pooling equilibrium under which both suppliers choose the general investment arises. In this equilibrium, the buyer offers the null contract and both suppliers reject the buyer's offer.

Intuitively, if the buyer believes that the supplier is likely to be trustworthy after observing the supplier's investment, the buyer offers the corresponding trusting contract: T_b for the buyer-specific investment and T_g for the general investment. Otherwise, the buyer offers the null contract. From an earlier result, the selfish supplier always exerts zero effort. However, the trustworthy supplier exerts positive effort in response to the trusting contract, and zero effort in response to the null contract. The detailed condition under which each of the three equilibrium exists is relegated to the appendix.

In all three pooling equilibria described in Theorem I.5 buyers' expected profits depend on the probability the buyer is facing a trustworthy supplier, θ. In the pooling equilibrium described in parts a) and b), buyers' expected profits are $\theta(\alpha_i e_i^{t*}) - \bar{u}_i - \gamma$, selfish suppliers earn profit $\bar{u}_i + \gamma$ and trustworthy suppliers earn profit $\bar{u}_i + \gamma - c(e_i^{t*})$ and get utility $(1-\phi)(\bar{u}_i + \gamma) - c(e_i^{t*}) + \alpha_i \phi e_i^{t*}$. Thus, expected total surplus in the pooling equilibria described in a) and b) is $\theta(\alpha_i e_i^{t*}) - c(e_i^{t*})$. In the pooling equilibrium described in part c) the buyer earns zero profits and both types of suppliers earn $\bar{u}_g$, so total surplus is $\bar{u}_g$.

1.3.2.3 Repeated Interaction

We extended the previous model to the case where, after the supplier chooses an investment, the buyer and the supplier engage in a finite number of repeated transactions ("periods"). We first characterize a separating equilibrium, analogous to the one described in the one-period model. Under this equilibrium, the buyer offers a trusting contract under the specific investment and a null contract under the general investment in each transaction. The trustworthy supplier chooses the specific investment, then the buyer offers a trusting contract to which the supplier reciprocates by exerting effort e_b^{t*}. Likewise, the selfish supplier chooses the general investment and is offered a null contract, which the supplier rejects. The sufficient condition for the separating equilibrium in a one-transaction game also guarantees a separating equilibrium in the finitely repeated game. This result is summarized in Theorem I.6.

Theorem I.6. *In a finitely repeated interaction game; there exists a separating equilibrium that is the same as in Theorem* (I.4)*.*

We also characterize another equilibrium which leads to collaborative outcomes denoted "semi-separating". In this equilibrium, both types of supplier choose the buyer-specific investment, and are offered a trusting contract in each transaction. Both suppliers exert effort e_b^{t*} for the first $N-1$ transactions. In the last transaction, the trustworthy supplier exerts effort e_b^{t*} and the selfish supplier exerts zero effort. We summarize this result in Theorem I.7:

Theorem I.7. *In a finitely repeated game, there exists a semi-separating equilibrium under which both suppliers choose the specific investment, and the buyer offers a trusting contract T_b in every period. Upon receiving the contract, the trustworthy supplier exerts effort $e_b^{t*} = c'^{-1}(\alpha_b \phi)$ for all periods, and the selfish supplier exerts the same effort except in period n in which he exerts zero effort.*

Since in repeated interactions firms may be concerned about how their actions in the current period affect their profits in future period, this new equilibrium arises in which selfish suppliers mimic trustworthy suppliers and exert high effort for a number of periods. We show in the proof that this can be part of a semi-separating equilibrium if the fraction of trustworthy suppliers is high enough. The detailed description of the model and the proofs are relegated to the Appendix.

1.4 Experimental Design

Our experiment consisted of ten rounds of the supply chain game and, after the supply chain ended, one round of each of two additional tasks: an investment game

(*Berg et al.*, 1995) to measure trust and trustworthiness and a lottery task (*Dohmen and Falk*, 2011) to measure risk attitudes.

1.4.1 The Supply Chain Game

Subjects were randomly assigned to the role of supplier or buyer, which they kept for all ten periods. In each period subjects were randomly and anonymously matched. This setup rules out reputation or repeated game effects. The supply chain game proceeded as described in our theoretical model: the supplier chooses between the buyer-specific or general investment, the buyer makes a price offer, and finally the supplier accepts or rejects the offer and makes an effort choice. For the buyer-specific investment, we set $\alpha_b = 12$, $\bar{u}_b = 0$. For the general investment, we set $\alpha_g = 3$, $\bar{u}_g = 15$. We also assume that the supplier incurs costs for his effort according to the canonical form, $c(e) = \frac{1}{2}e^2$. In order to simplify the subjects' task, they were presented with the following table:

Table 1.1: Cost of Effort Function

e	0	1	2	3	4	5	6	7	8	9	10
$c(e)$	0	0.5	2	4.5	8	12.5	18	24.5	32	40.5	50

In order to rule out negative payoffs, we added 60 points to the payoff of suppliers and 100 points to the payoff of buyers. Hence, the suppliers' payoff was $\pi^S = 60 + p - c(e)$ if he accepted the offer or $\pi^S = 60 + \bar{u}$ if he rejected, while the buyers' payoff was $\pi^B = 100 + \alpha e - p$ if the supplier accepted the offer or $\pi^B = 100$ if the supplier rejected. At the end of each round, subjects were informed their own payoff and the other subject's payoff.

1.4.2 Two Additional Tasks

The investment game has two roles: senders and receivers. Both senders and receivers are initially endowed with twenty points. The sender can transfer a portion of his endowment to the receiver, with any amount transfer being tripled. The receiver can then make a return transfer (without tripling) to the sender. We use the strategy method, with each subject choosing how much to send if they are the sender, and how much to return for each possible transfer amount if they are the receiver [13]. Subjects were then randomly assigned a role and matched to another subject for payment.

The lottery task gave subjects fifteen choices between a fixed payoff, which ranges from 2.5 to 37.5 in increments of 2.5, or a 50-50 lottery between a payoff of 40 points and a payoff of zero points. One decision was randomly selected for payment. The number of choices of the fixed payoff provides a measure of risk aversion.

1.4.3 Additional Treatments

First, we conduct a "random" treatment which is equivalent to the main treatment but suppliers are randomly assigned to an investment, both with equal probability. By assigning investments exogenously, we eliminate the signaling mechanism. Thus, this treatment is key to isolate the signaling power of the investment decision.

Secondly, we conduct a repeated interactions treatment in which subjects play six rounds of a repeated version of the supply chain game. In each round, the supplier makes one investment decision which is followed by three transactions. In each transaction, the buyer offers a price and the supplier decides whether he accepts the offer and, if so, a quality level. The values of the parameters are the same as those of the main treatment; $\alpha_b = 12$, $\bar{u}_b = 0$ for the buyer-specific investment, and $\alpha_g = 3$,

[13] Using the strategy method means that the receiver, instead of being asked how much he would like to send back given the amount he received, was asked how much to return for each possible transfer amount. In this way, we are able to elicit his complete strategy rather than his action in one particular case.

$\bar{u}_g = 15$ for the general investment. Suppliers start the each game with 60 points and buyers with 100 points. The supplier's and buyer's payoffs from the game are given by their initial endowments plus the sum of their payoffs from all three transactions. At the end of the experiment, one of the six rounds of the repeated supply chain game is randomly selected for payment.

Finally, we conducted two additional treatments of the supply chain game as robustness checks. In the "low benefit" treatment we reduce α_b from 12 to 6, making the specific investment less attractive and therefore reducing the range of individual parameters, ϕ and γ, within which the separating equilibrium arises. We also conducted an additional treatment where we increased α_b (18, versus 12 in the main treatment). While some of the results of these treatments are similar to those in the main treatment, these two additional treatments provide some interesting insights about the robustness of our results to changes in the values of the parameters. The analysis of these two cases has been relegated to the Appendix except for minor comments included in the main body of the paper.

1.4.4 Hypotheses

Our model predicts that a separating equilibrium can exist when the following two conditions are met. First, the quality improvement under the buyer-specific investment must be large enough so that the trustworthy supplier has enough incentive to exert effort on behalf of the buyer. Second, the outside option payoff under the general investment should be large enough so that the selfish supplier is incentivized to choose the general investment in order to improve his outside option value when the transaction does not close, but not too large so that the trustworthy supplier is not tempted to choose the general investment. If at least one of these two conditions is not met, the separating equilibrium breaks down [14].

[14] For example, if α_b is very low, both suppliers will choose the general investment. On the other hand, if $\bar{u}_g$ is very low, both suppliers prefer to choose the buyer-specific investment. We test this

In the experiment, we conduct a main treatment where the parameters of each investment are such that the separating equilibrium is likely to occur for reasonable values of individual-specific parameters (ϕ and γ) [15]. To derive hypotheses for the main treatment, we examine the comparative statics of the separating equilibrium from the model. We also derive hypotheses from the underlying cognitive mechanisms of reciprocity that generates those equilibrium results.

Recall that in the separating equilibrium trustworthy suppliers choose the buyer-specific investment and are offered a positive price, which they accept and exert positive effort, and selfish suppliers choose the general investment and are offered a null contract, which they reject. Thus, we expect to see higher price, higher acceptance and higher effort under the specific investment. Buyers' profits under the buyer-specific investment are $\alpha_b e_b^{t*} - \bar{u}_b - \gamma$ and are zero under the general investment and suppliers earn $\bar{u}_b + \gamma - c(e_b^{t*})$ under the buyer-specific investment and $\bar{u}_g$ under the general investment. Total profits should be higher under the buyer-specific investment for values of effort above a certain level. For the values of the parameters adopted in our experiment, this should be true for any effort greater than 1.33. Thus, the buyer-specific investment should also lead to higher buyer profits and higher total surplus than the general investment.

We can also observe the separation mechanism based on reciprocity in how suppliers respond to different price offers. Note that while the equilibrium makes specific point predictions for prices, the experimental data is likely to have a range of price offers. Lemma I.1 describes how we should expect subjects to respond to different price offers. Because subjects who choose the general investment are predicted to be selfish, they will provide the same (low) effort for any price offer. However, subjects who choose the specific investment reciprocate high price offers. Hence we would expect

prediction with the additional treatment with a low value of α_b that is presented in the Appendix.

[15] We present in the Appendix a figure indicating the range of parameters for which the separating holds.

that under the specific investment low prices will receive low effort, while high prices will receive high effort.[16] Therefore, there should be a strong positive correlation between price and effort in the specific investment, and a weak or zero correlation under the general investment. We will measure the strength of the effort-price relationship with the slope coefficient of the regression of effort on price.

Most importantly, our model assumes that investment decisions are driven by suppliers separating based on intrinsic characteristics. That is, intrinsically trustworthy suppliers choose the specific investment and intrinsically selfish suppliers choose the general investment. As a result, we expect that at the individual level subjects will differ in their propensity to choose the buyer-specific investment. In particular, we expect that there will be a positive correlation between subjects choosing the specific investment more often and subjects demonstrating a more "trustworthy behavior" in our experiment. We will identify "trustworthy behavior" in two ways. First, we measure subjects' price-effort correlation when they choose the specific investment and use the slope of the effort-price regression as a measure of trustworthiness. This is common in the experimental reciprocity literature. For example *Fehr et al.* (1993) use the wage-effort relationship in the gift exchange game, while *Berg et al.* (1995) use the ratio of amount sent to amount returned in an investment game. Second, we create a measure of reciprocity based on the subject's return transfer decisions in the investment game. Therefore, our theory predicts:

HYPOTHESIS 1. [Main Treatment] *We expect the relationship between supplier and buyer to be more collaborative under the buyer-specific investment than under the general investment. In addition, the frequency of choosing the specific investment is positively correlated with suppliers being more trustworthy.*

1.a - under the buyer-specific investment, buyers will offer higher prices, suppliers

[16] Additionally, we have simplified things theoretically by assuming that all trustworthy suppliers have the same ϕ and γ. This additional heterogeneity will further enhance and smooth out the price-effort correlation we describe.

will accept offers more often and exert higher effort, the price-effort relationship will be stronger, and buyers' profits and total profits will increase relative to the general investment, and

1.b - the frequency of choosing the specific investment is positively correlated with suppliers being more trustworthy as measured by a higher effort-price correlation in the supply chain game and more trustworthy behavior in the additional investment game.

In order to test whether the underlying mechanism driving the results in the main treatment is that in a separating equilibrium suppliers choose the buyer-specific investment to signal trustworthiness, we conduct an additional *random treatment*. In this treatment, suppliers are randomly assigned to an investment, which severs the connection between investment choice and the supplier's underlying preferences; the fraction of reciprocal suppliers is the same under both investments and therefore, the differences across investments can only be attributed to the specific investment having a higher quality coefficient. Thus, comparing the differences across investments in the random treatment relative to the main treatment allows to measure the impact of the separating mechanism.

To represent the random treatment we adapt our theoretical model by adding an initial move by nature that randomly assigns an investment to the supplier. Building on previous results, we identified three possible cases: If the buyer believes that the probability that the supplier is trustworthy is high enough, then he offers a trusting contract regardless of the investment. If the belief is moderate, he only offers a trusting contract under the buyer-specific investment and, if the belief is low, he does not offer a trusting contract in any case. The model is described in the appendix.

Assuming either case 2 or 3 arises, that is assuming there is a moderate/low number of trustworthy suppliers [17], the expected overall effort and total profits should be lower in the random treatment than in the main treatment. The expected overall effort is θe_b^{t*} in the main treatment and $\frac{1}{2}\theta e_b^{t*}$ if case 2 arises or zero if case 3 arises in the random treatment. The expected total profit is $\theta[\alpha_b e_b^{t*} - c(e_b^{t*})] + (1-\theta)\bar{u}_g$ in the main treatment and $\frac{1}{2}\theta[\alpha_b e_b^{t*} - c(e_b^{t*})] + \frac{1}{2}\bar{u}_g$ and $\frac{1}{2}\bar{u}_b + \frac{1}{2}\bar{u}_g$ in cases 2 and 3 of the random treatment respectively. Thus, if $\alpha_b e_b^{t*} - c(e_b^{t*}) > \bar{u}_g$, as predicted by Hypothesis 1, then expected total profits should be lower in the random treatment.

We also predict that the difference in expected effort and expected total profit across investments should be smaller in the random treatment relative to the main treatment. To see why, note that while under the main treatment all the suppliers who choose the specific investment exert effort e_b^{t*} and there is no provision of effort under the general investment, in the random treatment we expect the reciprocal effort e_b^{t*} to be exerted by at most a fraction θ of the suppliers with the specific investment and no provision of effort under the general investment if either case 2 or 3 arise. By the same argument, in the random treatment the expected total profit under the specific investment is $\theta[\alpha_b e_b^{t*} - c(e_b^{t*})]$ if case 2 arises and $\bar{u}_b$ if case 3 arises (both lower than $\alpha_b e_b^{t*} - c(e_b^{t*})$ under the main treatment) and under the general investment it is $\bar{u}_g$ in both treatments.

Finally, the random treatment allows to test the signaling role of the investment choice. Since in the random treatment the investment is no longer related to the the supplier's inherent type, we do not expect the specific investment to be positively correlated with the subject's trustworthiness. These predictions are presented in our last hypothesis.

[17] Case 1 is very unlikely to arise with the values of parameters used in the experiment, as it would require an extremely high fraction of trustworthy suppliers. For example, in the most extreme case where trustworthy subjects are perfectly reciprocal (i.e. surplus maximizers) case 1 would require 41% of suppliers to be trustworthy. Previous literature shows evidence of much lower rates: in modified dictator games, *Andreoni and Vesterlund* (2001) find that 21% are pure surplus maximizers (27% of men and 9% of women) and *Andreoni and Miller* (2002) find it to be 17% (or at most 22% including inexact classifications).

HYPOTHESIS 2. [Random Treatment] *Under the random treatment, expected effort and expected total profits are lower than in the main treatment, and the differences between the two investments is smaller than under the main treatment. The specific investment is not positively correlated with the subjects' trustworthiness.*
2.a - The overall effort provided by suppliers and total profits are lower than in the main treatment.
2.b - The difference in expected effort and expected total profits between the two investments is smaller than in the main treatment.
2.c - There will be no relationship between trustworthiness (measured by the effort-price correlation and the behavior in the additional investment game) and the specific investment.

We found that two different collaborative equilibria may arise in the repeated interactions setting. First, a separating equilibrium where a trustworthy supplier chooses the specific investment and a selfish supplier chooses the general investment, analogous to the one described in the one-shot interaction. Second, a semi-separating equilibrium where the selfish supplier mimics the trustworthy supplier by choosing the specific investment and offering high quality in all transactions except the last one.

Based on these theoretical findings, we expect the relationships in the repeated interactions case to be more collaborative under the specific investment. To see why, note that while the separating equilibrium predicts that the outcome of each transaction will be as in the single-transaction case, the semi-separating equilibrium predicts even more collaborative relationships under the specific investment (since selfish suppliers also choose the specific investment and provide high effort in at least some of the transactions). Our theoretical findings also suggest that in presence of trustworthy suppliers, repeated interactions should accentuate the social surplus of a specific investment relative to the single-transaction case. Both equilibria predict that overall effort and total surplus[18] should be at least as high as in the main treatment. The separating equilibrium predicts in each transaction the same overall effort and total surplus as in the separating equilibrium of the single-transaction game (overall effort: θe_b^{t*}, total surplus: $\theta[\alpha_b e_b^{t*} - c(e_b^{t*})] + (1-\theta)\bar{u}_g$). The semi-separating equilibrium predicts effort e_b^{t*} in periods 1 through $N-1$ and θe_b^{t*} in period N. Additionally, it predicts a total surplus of $\alpha_b e_b^{t*} - c(e_b^{t*})$ in periods 1 though $N-1$ and $\theta[\alpha_b e_b^{t*} - c(e_b^{t*})]$ in period N. Note that if Hypothesis 1 is true, and $\alpha_b e_b^{t*} - c(e_b^{t*}) > \bar{u}_g$, then the total surplus in every transaction period 1 through $N-1$ in repeated interactions should be higher than in the main treatment. Only in the last period, the semi-pooling predicts a lower total surplus than in the main treatment. We summarize these predictions in Hypothesis 5:

HYPOTHESIS 3. [Repeated Interactions Treatment]

The relationship between supplier and buyer is more collaborative under the buyer-specific investment.

3.a - Buyers offer higher prices, suppliers accept offers more often and exert higher effort under the specific investment than under the general investment. Buyers' profits and total profits are higher under the specific investment than under the general investment.

3.b - In each transaction, expected overall effort is at least as high as in the main treatment. Expected total surplus is at least as high as in the main treatment in every transaction except in the last one, where it can be higher or lower than in the main treatment.

[18] We define "surplus in transaction i" as the net profit a subject gets from that particular transaction (it does not include the initial endowment).

1.4.5 Experimental procedure

The experiment was conducted in z-Tree (*Fischbacher*, 2007) at the University of Michigan between June and July of 2011 and September and November of 2014. Subjects were paid for one randomly selected round of the Supply Chain game, for the investment game, and for one randomly selected choice from the lottery task. Subjects received $0.05 per point earned plus a $5 show up fee. Average payoffs were $12 (including the show up fee) and each session lasted approximately one hour.

1.5 Experimental Results

We conducted thirteen sessions of the main treatment of the experiment with between eight to fourteen subjects each time, who each played ten rounds of the supply chain game and one round of each additional task [19]. Overall, we had a total of 134 participants for the main treatment, 67 of which played as suppliers and 67 as buyers. In addition, we conducted five sessions of the random treatment with 54 participants in total, 27 playing as suppliers and 27 as buyers [20].

1.5.1 Differences between investments

Table 1.2 reports the fraction of times each investment was chosen and the average price, proportion of acceptance, average effort and average effort in accepted offers under the two investments. Since the supplier had the option of rejecting the buyer's offer and getting his outside option payoff, we distinguish the cases when the supplier

[19] The analysis presented in the Experimental Results section includes the data of the ten rounds. The results are qualitatively the same if we consider only the last five periods of play.

[20] The subjects were students at the University of Michigan. No subject participated in more than one session of the experiment. These treatments correspond to the first set of data, collected in 2011. Average age was 21.4 years, 43% were female and 57% were male. When asked about ethnicity, 49% of the subjects identified themselves as white, 36% as Asian or Pacific Islander, 9% as Black/African American, and the remainder as Hispanic, Multiracial, or Other. Students were from a range of different majors: 21% from Social Sciences, 20% from Sciences, 20% from Engineering, 10% from Economics, 8% from Business Administration, 5% from Arts and humanities, 5% from Medicine, and 11% from other fields.

accepts the offer and exerts zero effort from those when he rejects the offer. This allows us to observe in isolation the cases where the transaction did occur. Additionally, we observe what happens to the proportion of acceptance and to effort when prices are greater than 15. Note that, because under the general investment the supplier has an outside option of 15, for prices greater or equal to 15 the two investments have the same earnings potential, that is for every effort level suppliers make the same profit under both investments.

Table 1.2: Investment Comparison - Additional Treatments

Treatment	Investment	% Chosen	Average Price	% Accept	Average Effort	Average Effort (Accepted)	Average Effort (Accepted) Price $\geq$ 15
Main	Specific	67.46%	24.64	87%	1.78	2.03	2.74
	General	32.54%	12.62	39%	0.80	2.04	2.02
	p-value*		< 0.001	< 0.001	< 0.001	0.90	0.023
Random	Specific		21.52	84%	1.040	1.236	1.63
	General		14.00	40%	0.479	1.213	1.33
	p-value*		< 0.01	< 0.001	< 0.001	0.815	0.64

* Wilcoxon Rank-sum test of difference in average price, acceptance, effort and total profit between general and specific investments.

Our model predicts that in the main treatment buyers will offer higher prices to suppliers who chose the specific investment. In line with this prediction, we observe that the average price offered by the buyers when the suppliers choose the specific investment is nearly double than the price offered when the suppliers choose the general investment (Wilcoxon rank-sum test: $p < 0.001$). We also verify this result by regressing price on a dummy variable for the specific investment. We find that choosing the specific investment increases the price the buyer offers by 10.52 points ($p < 0.01$), supporting Hypothesis 1.a. At an individual level, we find that 72% of the subjects offer a higher price under the specific investment than under the general investment (Wilcoxon signed-rank test: $p < 0.001$).

In the random treatment, we observe that the difference in average price under the two investments is smaller than in the main treatment (7.52 versus 12.02) [21]. Since both treatments have the same quality coefficients, the higher price premium under the specific investment in the main treatment is attributed to the separating mechanism, which is not present in the random treatment.

In line with Hypothesis 1.a, the specific investment led to higher acceptance rates by suppliers in the main treatment (see Table 1.2). To control for price, we regress acceptance on price and investment type and present the results in Table **??**. We find that, even after controlling for price, choosing the specific investment increases the probability of acceptance by 26.38 percentage points ($\beta = 1.884$, $p < 0.001$, marginal effects = 0.2638). To correctly control for the difference in the outside option value, we repeat the same regression restricting the sample to offers with a price of 15 or higher, and find that there is still a significant positive correlation between choosing the specific investment and the probability of acceptance ($\beta = 1.323$, $p = 0.003$, marginal effects = 0.071). Together these results provide support to Hypothesis 1.a.

1.5.1.1 Effort

Table 1.2 shows that, in the main treatment, the overall provision of effort under the specific investment more than doubled that of the general investment and that the difference is significant (Wilcoxon rank-sum test: $p < 0.001$). Additionally, individual-level data shows that 67% of subjects exert higher or equal effort under the specific investment than under the general investment (Wilcoxon signed-rank test: $p = 0.016$). If we consider accepted offers only, we find no difference across investments. Our reciprocity model predicts not only higher effort levels under the specific investment but specifically effort that depends on price. While the general investment has an outside option of 15, the specific investment has an outside option of

[21] A regression of price on a specific investment dummy for each treatment shows that the price premium for the specific investment in the random treatment is directionally smaller but not statistically significant ($\beta = 10.52$ for the main treatment and $\beta = 8.96$ for the random treatment, $p = 0.65$).

zero. Hence, rejections are more likely to happen under the general investment than under the specific investment for prices lower than 15 (28% rejection under specific versus 87% rejection under general, $p < 0.001$), and hence the set of accepted offers for the specific investment includes more offers with a low price. Therefore, we are interested in testing whether there is a difference in effort across investments for accepted offers with price greater than 15 (rejections are 0.4% under the specific investment and 9.5% under the general investment). We find that, for prices greater or equal than 15, average effort in accepted offers under the specific investment is 2.74 and under the general investment it is 2.02 and the difference is statistically significant ($p = 0.023$). This result provides further support for Hypothesis 1.a. In the random treatment, overall effort is significantly lower than in the main treatment if we consider all offers (0.79 vs. 1.47, $p = 0.0001$) and considering accepted offers only (1.23 vs. 2.03, $p = 0.0002$), as predicted by Hypothesis 2.a. Furthermore, a regression of total effort on treatment dummies shows that the random treatment has a negative effect on overall effort relative to the main treatment ($\beta = -0.673$, $p < 0.001$). Hypothesis 2.b predicts that the difference in effort across investments is smaller in the random treatment than in the main treatment. The results support this hypothesis. The difference in average effort levels was 0.561 in the random treatment, versus 0.98 in the main treatment. While the difference in average effort across treatments is not significant under the general investment, (0.807 for the main treatment and 0.479 for the random treatment, $p = 0.223$), the difference is significant under the specific investment (1.783 and 1.040 respectively, $p < 0.001$). Additionally, if we consider only accepted offers with price greater or equal to 15, we find that in the random treatment there is no difference across investments, while in the main treatment the difference is significant.

Table 1.3 presents the results of regressing effort on price (Tobit regression for effort censored at zero) under the two investments type using price as a continuous independent variable and using a price dummy, which takes the value of one if price is greater or equal to 15 and zero otherwise. In the main treatment, we find a higher effort coefficient for price under the specific investment. In particular, the difference in the price dummy coefficients between the specific and general investments is statistically significant ($p = 0.019$). In the random treatment, we find that the price coefficient under the specific investment is much smaller than in the main treatment. Similarly, the difference between investments is smaller than in the main treatment and the difference in coefficients across investments is not significant ($p = 0.754$ and $p = 962$ for continuous and discrete price variables respectively).

Because we expect that the response to price will not necessarily be linear, and that the differences between the specific and general investments will be greatest at the higher quantiles of the effort distribution, Panel A in Table 1.4 estimates the effects of price on effort for the two investments separately at the $20th$, $40th$, $50th$, $60th$, and $80th$ quantiles in the main treatment. In Panel B, we test at each price level whether the investment type has a significant impact on the predicted values of effort for the different quantile regressions. This allows to estimate the effort distributions under the two investments separately for different price levels.

In Panel A we observe that the specific investment presents higher price sensitivity than the general investment. In particular, the difference becomes significant for the $60th$ and $80th$ quantiles. In addition, for the specific investment there is an increasing trend in price sensitivity as we go from the lowest quantiles to the highest quantiles. This implies that suppliers are significantly more price sensitive at the right tale of the effort distribution. On the other hand, under the general investment, suppliers are not very price sensitive at any part of the effort distribution. This means that those suppliers who choose high effort under the general investment are not being reciprocal in response to a generous offer, but rather being altruistic.

In Panel B we observe the effort predicted by the quantile regressions for different price levels under the two investments. We find that at a moderate price (20), the predicted level of effort is very similar under the two investments at all quantiles, indicating that both investments have a similar effort distribution. For prices greater than 20, the high quantile regressions have significantly higher predicted effort under the specific investment than under the general investment. This indicates that the effort distribution for the specific investment shifts towards the right and has a longer right tale than the distribution of effort under the general investment. Thus, for high price offers $(40, 60)$ the specific investment is most beneficial for the buyer: it makes him better off at the right tale of the effort distribution and not worse off at the left tale of the effort distribution. For prices smaller than 20, we find that the effort distributions have shorter tales under both investments but now the distribution under the general investment is the one that is slightly shifted to the right. The reason for this is that under the general investment some suppliers who accept low price offers (though this does not happen very often since suppliers have an outside option of 15) are not price sensitive and choose high effort. On the other hand, under the specific investment suppliers are price sensitive so if they accept a low price, they choose low effort. Thus, for low price offers $(2, 10)$, buyers are better off under the general investment where there are some altruistic suppliers who offer high prices even when the price offered was low. This also explains the results on Table 1.2 where we observe that the premium in average effort under the specific investment only becomes significant for price offers greater or equal to 15.

1.5.1.2 Profit

We next examine the impact of the specific investment on buyers' profits and total profits in the supply chain.

Hypothesis 1.*a* predicts that in the main treatment buyers' profit is higher under the specific investment. Our data supports this - buyers' average profit was 96.88 un-

Table 1.3: Price-Effort Relationship for Specific vs. General - Main and Random Treatments

Coefficients	Effort (Accepted)	Effort (Accepted)
Price x Specific x Main Treatment	0.085***	3.408***
	(0.006)	(0.367)
Price x General x Main Treatment	0.061***	1.290
	(0.017)	(0.821)
Price x Specific x Random Treatment	0.049***	2.167***
	(0.011)	(0.640)
Price x General x Random Treatment	0.042*	2.271
	(0.021)	(2.065)
Specific x Main Treatment	-1.681***	-2.092***
	(0.591)	(0.790)
Specific x Random Treatment	0.020	0.739
	(0.841)	(2.089)
Random Treatment	-1.400	-2.925
	(1.051)	(2.195)
Constant	0.094	0.562
	(0.599)	(0.771)
Observations	657	657
Nr. of Subjects	94	94

Test Price x Specific = Price x General	$p - value$	$p - value$
Main Treatment	0.206	0.019
Random Treatment	0.754	0.962

Tobit regressions with subject random effects. Standard errors reported in parentheses. Column 1 has price as independent variable and column 2 has a price dummy (which takes value one if the price is greater or equal to 15 and zero otherwise) as independent variable. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 1.4: Price-Effort Quantile Regressions - Main Treatment

Panel A

Quantile	$Q_{0.20}$	$Q_{0.40}$	$Q_{0.50}$	$Q_{0.60}$	$Q_{0.80}$
Coefficients			Effort		
Price x Specific	0.017**	0.053***	0.076***	0.086***	0.091***
	(0.008)	(0.009)	(0.008)	(0.006)	(0.004)
Price x General	0	0.040	0.057***	0.045***	0.030
	(0.009)	(0.025)	(0.017)	(0.016)	(0.020)
Specific	-0.167	-0.267	-0.076	-0.866	-1.727**
	(0.208)	(0.482)	(0.423)	(0.547)	(0.817)
Constant	0	0	0	0.866	2.455***
	(0.170)	(0.477)	(0.413)	(0.540)	(0.794)
Observations	483	483	483	483	483

Difference in Price Slopes (Test Specific = General)					
$(p - value)$	0.171	0.618	0.323	0.017	0.003

Panel B

Total Investment Effects	At price	$Q_{0.20}$	$Q_{0.40}$	$Q_{0.50}$	$Q_{0.60}$	$Q_{0.80}$
Specific	2	-0.133	-0.160***	0.076	0.171**	0.909***
General	2	0	0.080	0.114	0.955*	2.515***
Specific - General	2	-0.133	-0.240	-0.038	-0.784	-1.606
$(p - value)$	2	0.472	0.587	0.922	0.132	0.041
Specific	10	0	0.267***	0.684***	0.857***	1.636***
General	10	0	0.400	0.571**	1.313***	2.758***
Specific - General	10	0	-0.133	0.112	-0.456	-1.121
$(p - value)$	10	1	0.682	0.708	0.276	0.086
Specific	20	0.167**	0.800***	1.443***	1.714***	2.545***
General	20	0	0.800**	1.143***	1.761***	3.061***
Specific - General	20	0.167	-5.55e-17	0.300	-0.047	-0.515
$(p - value)$	20	0.026	1.000	0.267	0.887	0.316
Specific	40	0.500**	1.867***	2.962***	3.429***	4.364***
General	40	0	1.600**	2.286***	2.657***	3.667***
Specific - General	40	0.500	0.267	0.676	0.772	0.697
$(p - value)$	40	0.088	0.728	0.183	0.049	0.115
Specific	60	0.833**	2.933***	4.481***	5.143***	6.182***
General	60	0	2.400**	3.429***	3.552***	4.273***
Specific - General	60	0.833	0.533	1.052	1.591	1.909
$(p - value)$	60	0.119	0.676	0.219	0.016	0.005

Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Total Investment Effects estimates the total difference between investments given prices at the corresponding quantiles of the price distribution.

der the specific investment and 92.94 under the general investment ($p = 0.02$). When we regress buyers' profit on the specific investment, we find a positive and significant effect ($\beta = 4.725$, $p < 0.01$, presented in Table A.3 in the Appendix). Additionally, a within-subject comparison of average profit under the two investments shows that, for an individual buyer, profits were higher when he was paired with a supplier who chose the specific investment than when he was paired with a supplier who chose the general investment (Wilcoxon signed-rank test: $p = 0.016$, 61% of the subjects)[22]. While the data shows that, as the theory predicts, buyers make significantly higher profits under the specific investment than under the general investment, we also note that the effect is not big enough to allow buyers to earn strictly positive net profits in the main treatment. Buyers do better in the High α treatment, where the specific investment allows them to break even (average buyers' profit is 101.968 and not statistically different from the initial endowment of 100, see the Appendix for more details on this treatment). This result is consistent with previous experimental studies which show that the return to trust seldom pays back (*Camerer* 2003, p.86) [23]. As we report in Section 5.3, buyers do earn strictly positive average profits under the specific investment in the repeated interactions treatment.

Hypothesis 1.a also predicts that total profits are higher under the specific investment. A regression of total profit on a dummy variable for choosing the specific investment shows that the specific investment has a significant positive effect on total profits, consistently with Hypothesis 1.a ($\beta = 5.352$, $p < 0.01$, presented in Table A.3 in the appendix).

Hypothesis 2.a predicts that total profits are lower in the random treatment rel-

[22]Buyers offered strictly positive prices in 383 out of the 452 transactions where suppliers chose the specific investment. Out of these 383 transactions, 143 of them resulted in the buyers earning strictly positive profits (37%). On the other hand, suppliers chose the general investment in 218 transactions and out of these, 158 had strictly positive prices. Only 2 out of these 158 transactions resulted in strictly positive profits for the buyers (1%).

[23]In the investment game reported in *Berg et al.* (1995) the average return is 90% of the amount sent and in the additional investment game subjects play in our experiment it is 85%.

ative to the main treatment. Average total profits in the random treatment were 170.10, compared to the main treatment 174.63 (Rank-sum test: $p = 0.005$). When we regress total profit on treatment dummies (presented in Table A.4 in the Appendix), we observe that the random treatment has significant negative effects on total profit relative to the main treatment ($\beta = -4.53$ with $p = 0.003$), consistent with our hypotheses. Hypothesis 2.b predicts that the difference in total profits across investments is lower in the random treatment than in the main treatment. The difference in average total profit across investments is 0.69 in the random treatment and 7.63 in the main treatment. Additionally, we regress profits on separate indicator variables for the specific investment for each treatment, as well as a treatment dummy. We find that the effect of the specific investment on total profit in the main treatment is greater than in the random treatment ($\beta = 5.352, p = 0.001$ in the main treatment, $\beta = 1.132, p = 0.559$ in the random treatment, difference $p = 0.100$). The result is presented on Table A.3 in the Appendix.

Similar results hold for buyers' profits. Average buyer profit is 91.17 in the random treatment, significantly lower than the profit of 95.6 in the main treatment ($p = 0.005$). The difference in average buyer's profit across the two investments is smaller in the random treatment than in the main treatment (0.67 in the random treatment vs. 3.94 in the main treatment). We also find that, at individual level, in the random treatment a buyer does not make a significantly higher profit under the specific investment than under the general investment (Wilcoxon signed-rank test: $p > 0.20$), as was the case in the main treatment. In addition, a regression with the interaction effects of the specific investment and treatment on profits shows that in the random treatment the specific investment has no significant effect on profits (Table A.3 in the Appendix). The strong connection between the main treatment and the higher benefits of the specific investment, support the argument that it is the separating equilibrium which drives the results of the main treatment.

1.5.2 Individual Differences

We next examine whether the aggregate results discussed in the previous section are caused by all subjects behaving differently under the specific investment, or whether individuals who choose the specific investment are inherently different from those who choose the general one.

1.5.2.1 Heterogeneity in Investment Choice

We predict that in the main treatment trustworthy suppliers signal their type by choosing the buyer-specific investment. If investment choice is driven by sorting based on an underlying preference type, some subjects should persistently choose either the specific or the general investment. We conduct several tests which confirm that subjects exhibit significant heterogeneity in investment choices. First, we find that subjects who choose the specific investment more often in the initial five periods are significantly more likely to continue to choose it in the later five periods ($\rho = 0.588$, $p < 0.05$). We find similar results using a non-parametric test for trends ($p < 0.01$)[24]. Lastly, a permutation test indicates that significantly more subjects choose the specific investment at least 8 times than would be expected if all subjects chose between investments with a common probability in each period ($p < 0.01$) [25]. Figure A.2 in the Appendix displays the fraction of subjects who chose the specific investment a given number of times.

[24] We used a "Wilcoxon-like test for trends" introduced by *Cuzick* (1985). The test conducts a non-parametric test for trend across ordered groups, which is an extension of the Wilcoxon rank-sum test. It works by computing the average ranks for one group and then correlating the average ranks with the values in the other group. It tests for a trend of (increasing) values in the ranks of one group across the values of the other group.

[25] For each period, we shuffled the investment decisions across subjects and compared the number of times subjects got the specific investment with the frequencies observed in the experimental results. We conducted 100.000 iterations of the shuffling.

1.5.2.2 Heterogeneity in Effort Choice

We next examine whether this difference in investment choice between subjects corresponds with different effort choices. We expect that subjects who choose the specific investment frequently will be inherently more trustworthy. Hence, as predicted by Hypothesis 1.b, we should find that these subjects have a larger price-effort relationship. To test for this, in Table 1.5 we regress effort on price, conditional on having the specific investment in this period, separately for subjects who frequently (infrequently) choose the specific investment across all periods[26]. We find that the effect of price on effort is approximately twice as large for subjects who frequently chose the specific investment and the difference in effort-price correlation between subjects who choose the specific investment with high and low frequency is statistically significant ($p < 0.001$). This shows that suppliers who choose the specific investment more often were also more trustworthy, providing support for Hypothesis 1.b. Choosing the specific investment does accurately signal that the supplier is inherently more reciprocal, and will choose a higher effort if offered a high price.

When we examine the subject-level behavior in the random treatment we find that the relationship between the specific investment and trustworthiness is no longer present. As predicted by Hypothesis 2.c, when we compare the price-effort relationships of subjects that were assigned the specific investment six times or more to those assigned it five or fewer times, we find that high-frequency subjects actually have a lower price coefficient than low-frequency subjects with only marginal significant difference ($p = 0.097$ presented in Table 1.5) [27].

These results explain why the specific investment is no longer more profitable in the random treatment. The fact that the specific investment is no longer positively

[26]The cutoff point of choosing the specific investment at least eight was chosen based on the results of the permutation test reported above.

[27]The cutoff point was set at those subjects who were assigned the specific investment six times or more so that the fraction of subjects above and below the cutoff point was close to that in the main treatment.

Table 1.5: Individual Specific Price-Effort Relationship Under Specific Investment - Main and Random Treatments

Coefficients	Effort (Accepted)
Price x Specific Often x Main Treatment	0.100***
	(0.007)
Price x General Often x Main Treatment	0.048***
	(0.009)
Price x Specific Often x Random Treatment	0.034***
	(0.013)
Price x General Often x Random Treatment	0.067***
	(0.017)
Specific Often x Main Treatment	-0.685
	(0.640)
Specific Often x Random Treatment	-0.404
	(1.048)
Random Treatment	-0.112
	(0.899)
Constant	-0.872*
	(0.451)
Observations	524
Nr. of Subjects	92

Test Price x Specific Often = Price x General Often	$p - value$
Main Treatment	< 0.001
Random Treatment	0.119

Tobit regressions with subject random effects. Standard errors reported in parentheses. Specific Often is a dummy variable which takes value 1 if the supplier's frequency of the specific investment is greater or equal to 8 for the main treatment and 6 for the random treatment. General Often is a dummy variable which takes value 1 if the subject chose the specific investment 7 times or less in the main treatment and 5 times or less in the random treatment. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

correlated with the suppliers' type shows that the investment type is no longer a good predictor of trustworthiness. These results provide support to the hypothesis that the separating equilibrium is the main driver behind our results.

1.5.2.3 Heterogeneity in Trust and Risk Attitudes

We further examined two additional questions: whether other behavioral factors influence sorting, and whether other measures of trustworthiness correlate with the measure obtained from the effort-price correlation in the supply chain game. After playing the supply chain game, subjects played one round of the investment game and completed a risk aversion task [28]. We measured risk aversion as the fraction of times the subject chose the fixed payoff over the $50 - 50$ chance lottery, yielding a distribution of subjects' risk aversion between 0 and 1. In the investment game subjects could choose to send between 0 and 20 points in increments of two points to some other subject they were randomly and anonymously paired with. Any amount sent was tripled. Subjects were then asked how much they would like to send back for different amounts they could have received, up to the total amount received. We used the amount sent as a measure of subjects' trust, therefore trust ranged between 0 and 20. We created a measure of subject's trustworthiness based on their answers to how much they would return by taking the difference between the maximum and the minimum amounts they wanted to return [29]. Given that in the vast majority of the

[28]One potential concern is that the course of play in the preceding supply chain game influenced subjects choices in the the additional tasks. While we cannot fully rule out this form of reverse-causality, we tried to minimize the connections by using contextualized instructions for the supply chain game and abstract instructions for the additional tasks. Additionally, if there were substantial spill-over effects from the supply chain game, one might expect that subjects who had been playing different roles would make different choices in the additional tasks. However, we do not find a significant difference between suppliers and buyers ($p > 0.20$ for risk aversion and trustworthiness, $p > 0.10$ for trust). Additionally, subjects who were randomly assigned the specific investment more or less often in the random treatment (described below) do not make significantly different choices in the additional tasks.

[29]We considered two other measures of trustworthiness, one was the difference between the minimum and maximum amount returned as a fraction of the amount received and the other one was the sum of all the net returns. All the main results remained the same regardless of which measure was used.

cases the amount returned was (at least weakly) increasing in the amount sent, this measure of trustworthiness captured how different the subjects' response was when the sender was kind from when the sender was unkind. Trustworthiness could then range between 0 and 60. Table A.5 in the Appendix summarizes risk aversion, trust and reciprocity observed in the two additional tasks.

Our results show that trustworthiness in the investment game is a good predictor of sorting in the investment choice of the supply chain game. Specifically, having an investment game measure of trustworthiness above the median is correlated with a higher likelihood of choosing the specific investment ($\beta = 0.630$, $p = 0.045$, marginal effects = 0.243)). This provides support for Hypothesis 1.b. Additionally, suppliers who are more trusting (i.e. sent more than the median in the trust game) are significantly more likely to choose the specific investment eight times or more ($\beta = 1.298$, $p < 0.001$, marginal effects = 0.452). This suggests that suppliers sort not only on their willingness to repay high prices within the supply chain game and their reciprocity measure in the investment game, but also on their propensity to trust in others. Similarly, we find that trust also predicts buyers' willingness to offer high prices ($\beta = 9.920$, $p = 0.009$). We also find that higher levels of trust and trustworthiness in the investment game are associated with a higher effort-price correlation in the supply chain game (when we test the difference in coefficients for high versus low trust, trustworthiness, and risk aversion we get $p = 0.013, 0.061$, and 0.228 respectively). In the random treatment we find that, as predicted by Hypothesis 2.c, neither trust nor trustworthiness in the trust game are correlated with a high frequency of the specific investment (see Table A.7 in the Appendix).

1.5.3 Repeated Interactions

In the previous sections, we have considered the effects of trustworthiness and up-front buyer specific investments on supply-chain efficiency when firms interact only once. In this section we study the impact of longer relationships on the sorting mechanism with a finitely repeated version of the supply chain game. We conducted five sessions of the repeated interactions treatment, with a total of 50 subjects [30]. Total payoffs from the experiment include the payoff from the Supply Chain Game, the payoffs from the two additional tasks and a $7 participation fee. The average payoff was $15 and each session lasted approximately 90 minutes.

Table 1.6 shows that, as predicted by Hypothesis 5.a, average price, acceptance rate, and effort are higher under the specific investment than under the general investment in all three periods, and the differences are significant. Table 1.7 shows that the effort-price correlation is also higher under the specific investment in every transaction. Buyers make higher profits[31] when the buyer-specific investment is chosen than when the general investment is chosen (110.961 vs. 63.75, Rank-sum test p-value: < 0.0001). This is the result of a higher surplus in every transaction (see Table 1.6). Total profits are also significantly higher under the specific investment, as predicted by Hypothesis 5.a (247.51 for specific vs. 189.26 for general, $p < 0.001$). This result is confirmed by the regression of buyers' profits and total profits on a specific investment indicator variable presented in Table A.9 in the Appendix (buyer's profit: $\beta = 45.640$, $p < 0.001$, total profits: $\beta = 44.931$, $p < 0.001$).

Hypothesis 5.b predicts that in each transaction, expected overall effort is at least as high as in the main treatment and that expected total surplus is at least as high as in the main treatment in every transaction except in the last one, where it can be higher or lower than in the main treatment. Total effort is 1.466 in the main treatment and 3.2 and 2.733 in transactions 1 and 2 of the repeated interactions

[30]This treatment corresponds to the second set of data, collected in 2014. Average age was 21.8 years, 61% were female and 39% were male. When asked about ethnicity, 37% of the subjects identified themselves as white, 45% as Asian or Pacific Islander, 12% as Black/African American, and the remainder as Hispanic, Multiracial, or Other. The demographics by major were: 17% from Social Sciences, 29% from Sciences, 17% from Engineering, 17% from Economics, 4% from Business Administration, and the reminder were from Arts and humanities, Medicine, and other fields.

[31]Buyers' and suppliers' profits are defined as their initial endowment plus the sum of their surplus in all three trading periods.

treatment respectively (both differences are significant, $p < 0.001$). In transaction 3, the difference with the main treatment is not significant (1.593, $p = 0.182$). Similarly, total surplus is 14.628 in the main treatment and 27.813 and 24.48 in transaction 1 and 2 of the repeated interactions treatment respectively (both differences have $p < 0.001$). In transaction 3, total surplus is 16.577, which is not significantly different from the total surplus in the main treatment ($p = 0.857$).

Buyers' surplus under the specific investment in transactions 1 and 2 are significantly higher than in the main treatment (6.127 and 9.608 respectively versus -3.12, both with $p < 0.001$) and not different in transaction 3 (-4.77 versus -3.12, $p = 0.214$). While buyers' profits under the specific investment increase relative to the main treatment, buyers' profits under the general investment significantly decrease relative to the main treatment. As a result, the difference in buyers' profits across the two investments increases in repeated interactions relative to the main treatment (47.211 versus 4.44). Similarly, the difference in total profits across investments is significantly larger under the repeated interactions treatment than in the main treatment (58.25 in repeated interactions vs. 7.63 in the main treatment). These results support the argument that the benefits of the specific investment are enhanced in repeated interactions.

We find two additional important results. First, there is strong evidence that the buyers' profit under the buyer-specific investment is greater than the initial endowment of 100. That is, the increase in buyers' surplus in transactions 1 and 2 more than compensates for the drop in transaction 3. We conduct a pair-wise comparison of each retailer's profit observation with a variable that has all its values equal to 100, and reject the hypothesis that the two variables are equivalent ($p = 0.015$). Additionally, we regress buyers' profit on a dummy variable for the buyer-specific investment (the results are presented on Table A.9 in the Appendix). A 95% confidence interval for buyers' profit under the specific investment is $[100.97; 120.22]$. Therefore at $\alpha = 0.05$, it confirms that buyers' profits are greater than 100. This result stresses the impact of buyer-specific investments on buyers' profits in repeated interactions. Buyers can expect to make positive profits even in a context of non-contractible quality. Secondly, we observe that both price and effort remain relatively high under the specific investment in the last transaction (average price is 31.569 and average effort is 3.139). Similarly, Table 1.7 shows that the effort-price correlation is higher under the specific investment than under the general investment and the difference is even higher than in previous periods. These results provide support for the existence of a fairly high number of trustworthy suppliers, which is a requirement for the semi-separating equilibrium to arise.

Table 1.6: Average Price, Acceptance, Effort, and Surplus

	Investment	Transaction 1	Transaction 2	Transaction 3
Price	General	20.708	18.167	10.958
(all offers)	Specific	42.863***	33.020***	31.569***
Price	General	38.541	35.818	36.546
(accepted offers)	Specific	45.021	37.798	44.431
Acceptance	General	0.5	0.458	0.229
	Specific	0.951***	0.873***	0.706***
Effort	General	2.667	2.182	1.182
(accepted offers)	Specific	4.289***	4.067***	3.139***
Supplier's Surplus	General	24.083	22.125	19.302
	Specific	30.627**	22.294	23.627
Buyer's Surplus	General	−15.271	−13.427	−7.562
	Specific	6.127***	9.608***	−4.774*

Note: Surplus refers to the net profit from a transactions and does not include the initial endowment. Significant differences across investments is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

1.6 Discussion

The experimental results largely confirm our hypotheses based on a signaling model. First, we show that the upfront choice of a specific investment results in a

Table 1.7: Price - Effort Relationship - Repeated Interactions

Coefficients	Transaction 1		Transaction 2		Transaction 3	
	Effort Specific	Effort General	Effort Specific	Effort General	Effort Specific	Effort General
Price	0.091***	0.051**	0.108***	0.078**	0.080***	-0.051
	(0.008)	(0.025)	(0.006)	(0.036)	(0.014)	(0.071)
Constant	-0.360	0.540	-0.223	-1.619	-1.546	1.243
	(0.566)	(1.110)	(0.322)	(1.701)	(0.969)	(2.671)
Observations	97	24	89	22	72	11
Nr. of Subjects	24	15	22	12	21	10

Tobit regressions (accepted offers only) with subject random effects, except in column 6 which has 11 observations and 10 subjects. Standard errors reported in parentheses. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

more collaborative relationship between buyers and suppliers. Buyers offer higher prices, suppliers accept offers more often, the overall provision of effort is higher, and this results in higher total profits. We show that this is possible because there exists a positive correlation between choosing the specific investment and the supplier being trustworthy. Therefore, the investment choice can help buyers identify trustworthy suppliers before contracting. We also find that when suppliers are randomly assigned an investment, the efficiency premium of the specific investment disappears. This supports the hypothesis that it is the signaling effect of the investment choice what drives the specific investments efficiency enhancement. Finally, we show that repeated interactions magnify the benefits of the specific investment, leading to even more collaborative relationships. This could be either attributed to a scenario where trustworthy suppliers choose the specific investment and selfish suppliers choose the general investment, or a scenario where both types of suppliers choose the specific investment and selfish suppliers mimic trustworthy suppliers for some number of periods.

One result not anticipated by our theoretical model is the importance of trust for both buyers and suppliers. In our model, we account for the supplier's trustworthiness and this is sufficient for the separating equilibrium to arise. The experimental data (presented in Table A.6) shows that, suppliers must trust that buyers actually play the equilibrium and respond to their investment choice with high prices. Similarly, buyers must trust that suppliers will reciprocate high prices with high effort. By contrast, the supplier's effort choice, which has no subsequent buyer action, does not depend on trust since this choice does not make the supplier vulnerable. Since trust plays a role in subjects' decision making, a related consideration is whether subjects' decisions are influenced by "betrayal aversion". Previous research on the trust game suggests that people may be averse to being betrayed (*Bohnet and Zeckhauser* 2004). In order to analyze this, we identify as a "betrayal" the case where a buyer offered a strictly positive price and received effort less or equal to one, conditional on acceptance. We find that the price offers in periods 6 to 10 are correlated with the number of betrayals experienced in the first five periods ($\beta = -4.440$, $p < 0.001$). In addition, we find that a betrayal in the period immediate previous to the current one has a significant negative effect on buyers' price offers ($\beta = -6.641$, $p < 0.001$). This further suggest that in practice buyers have some wariness about whether suppliers will in fact be trustworthy.

Another surprising result is the increase in supplier's monetary profits under the specific investment. While our theoretical model predicts that trustworthy suppliers will have higher *utility* under the specific investment, this is predicted to be entirely due to the non-monetary reciprocal utility. Under the separating equilibrium the suppliers monetary profits should be higher under the general investment to prevent selfish suppliers from switching their investment choice. In our data, however, suppliers' average profit under the specific investment were not significantly lower than that under the general investment (80.23 under specific vs. 76.54 under general, Wilcoxon rank-sum test: $p = 0.97$). Similarly, the within-subject pairwise comparison of average profit under the two investments shows that a supplier's profit was not significantly lower under the specific investment than under the general investment (Wilcoxon signed-rank test: $p = 0.1935$). A regression shows that the specific investment increases suppliers' profit by 3.216 points and the effect is significant a the 5% significance level. One possible explanation for what keeps selfish suppliers from switching to the specific investment despite the monetary benefits is strategic uncertainty: since there is a wide range of prices buyers offer under the specific investment, selfish suppliers who are risk averse may want to guarantee for themselves the general investment's higher outside option.

Finally, subjects' decisions in the game could be affected by other interpersonal concerns such as inequality aversion. If they perceive that their previous payoff was too low (or too high) compared to their partners' payoff, they may want to adjust their behavior in the following period to make payoffs more equitable. This could happen even if subjects play with different subjects in every round. To analyze this, we tested whether firms' profit premium (over their partner's profit) earned in the previous period affected their decision in the current period. For buyers, we find that price offers were not affected by profit inequality in the previous period. For suppliers, we find that the choice of effort is negatively correlated with their profit premium in the previous period. This suggests the opposite of inequality aversion: the more suppliers have earned over their partners in the previous period, the lower the effort they provide in the current period (which will further increase inequality). These results suggest that inequality aversion does not play an important role in this context.

1.7 Conclusion

We investigate how firms can benefit from identifying trustworthy suppliers when non-contractible factors such as quality are important. We suggest that upfront relationship-specific investments can signal that a supplier is trustworthy and will provide high quality if awarded a high price contract. This provides an explanation to why certain suppliers want to make a buyer-specific investment before contracting. We identify theoretical conditions where this signaling mechanism can generate a separating equilibrium with selfish suppliers choosing a general investment (that improves the supplier's outside option), while trustworthy suppliers choose the specific investment (that increases the efficiency of the supplier's quality-generating effort).

We test our model using a laboratory experiment. The results of our supply chain game confirm the effect of signaling on supply chain performance. Subjects who consistently choose the specific investment are significantly more trustworthy as suppliers. As a result, contracting with suppliers who made the relationship-specific investment leads to higher buyer profits and supply chain profits. Offering a price premium to suppliers who chose the specific investment leads to higher quality, as well as higher profits for the supply chain and both individual firms. Thus, buying firms facing a supplier who made an up-front specific investment should consider offering generous contracts even when quality is non-contractible. Our model determines that for the signaling mechanism to arise, the buyer-specific investment needs to be efficient enough and the general investment must provide sufficient monetary incentives to the selfish supplier. The experimental results show that, when these conditions are not present and the signaling mechanism is reduced or eliminated, the relation between the specific investment and trustworthiness is no longer present. As a result, the relationship-specific investment no longer leads to higher profits.

Finally, we show that the benefits of upfront buyer-specific investments in promoting collaboration and increasing profits are further strengthened with repeated interactions. We characterize two possible equilibria: one where trustworthy supplier fully separate from the selfish suppliers, and one where selfish suppliers mimic the investment choice and initial effort decisions of the reciprocal suppliers. In both cases repeated interactions only increase the transaction surplus when there are sufficiently many trustworthy firms. Our experimental results confirm this intuition: supply chain profits are substantially increased with repeated transactions, and buyers benefit heavily from working with trustworthy suppliers.

Taken together our results show that there is great value in a buyer being able to identify a trustworthy supplier, and suggest one potential avenue for trustworthy suppliers to distinguish themselves. Future research can explore other ways that trustworthy suppliers can identify themselves, and other supply chain settings where this signaling is important.

CHAPTER II

The Signaling and Incentive Effects of Supplier Awards

2.1 Introduction

One of the most important aspects of managing a good supply chain is to build and retain relationships with suppliers, as what suppliers do affects a firm and its customers. The field of operations management has produced a vast literature in buyer-supplier relationships. Traditionally, the most obvious objective that has been studied is efficiency or minimizing fulfillment cost. The vast majority of the literature focuses on designing optimal contracts and/or comparing different contracts (e.g., *Cachon* 2003; *Pasternack* 2008; *Cachon and Lariviere* 2005). Although a contract can be an effective tool to agree on obvious metrics such as cost and order quantities, there are other aspects of the relationship where what suppliers do affects the firm's bottom line, many of which are not contractible. For example, random disruptions and catastrophic events are difficult to predict, therefore, the parties have seldom agreed on subsequent actions (response, compensation, etc.) in a standard fulfillment contract. In case of a complex component, it may not be feasible to specify numerous details. As a result, the supplier has a discretion to choose the quality level for these dimensions. For instance, if the noise level of a hard disc spindle is not specified in a contract, the supplier is not required (although noise reduction is certainly desired by the buyer) to reduce the noise level of the product (*Kaya and Özer* 2009). In a similar vein, the quality of a multi-faceted service is not fully contractible and enforceable.

In these cases, non-contractual instruments such as behavioral and social preferences greatly impact the nature of relationship and the supply chain performance; however, these instruments are largely unstudied in the OM literature. One of such instruments is the role of an award in a supply chain relationship. Many firms reward the top-performing suppliers with an "Outstanding Supplier" or "Supplier of the Year" award for the supplier's "above and beyond" efforts and commitment. The awards are typically given to suppliers who excel in creative cost-reduction solutions, teamwork, customer service, response to supply chain disruption, or sustainability. For example, in 2011, Verizon recognized Ciena with an "Outstanding Performance Award" for the outstanding performance of Ciena's packet-optical switches during Japanese earthquake. [1] Similarly, the power solutions division of Johnson Controls received the GM's 2009 Supplier of the Year Award for consistently exceeding GM's expectations. [2] The award-winning suppliers typically publicize the awards through a press release.

In addition to free press coverage, there exist a number of reasons why some suppliers may care about symbolic awards. In some cases, a supplier's corporate culture or incentives strategically encourage its managers to seek awards from business counterparts. For instance, if a manager's performance and incentives (i.e., promotion, bonuses, incentives) are tied to the buying firm's recognition or feedback, awards will affect his/her efforts and response to the buying firm. Another plausible explanation is that even in a supplier-buyer relationship, decision makers of both firms who are involved in forming, retaining, and dissolving the relationship are, after all, individ-

[1] Source:http://www.ciena.com/connect/blog/Verizon-recognizes-Ciena-with-Supplier-of-the-Year-award.html

[2] Source:http://www.prnewswire.com/news-releases/johnson-controls-receives-2009-supplier-of-the-year-award-from-general-motors-corp-89136407.html

uals. Even in large corporations such as GM, the relationship with a key supplier is managed by a handful of individuals whose social preferences and behavior will influence the relationship that develops between the firms. Thus, if symbolic awards induce intrinsic motivation of individuals working in the supplier as proven through a number of experiments in behavioral economics (*Lacetera and Macis* 2010b, *Kosfeld and Neckermann* 2011, *Ashraf et al.* 2012, *Bradler et al.* 2013), their actions affect the relationship between the two firms at a corporate level.

From a buying firm's perspective, it looks as if these symbolic awards do no harm at all. They cost little to the buyer, yet a positive feedback recognizes the supplier's effort and improves the nature of relationship between the two firms. In case the supplier cares about the award, awards incentivize suppliers to put above and beyond efforts for the buyer. Thus, a symbolic award seems to be an effective tool to help the buyer "identify" and "retain" a good supplier. However, when there are multiple buyers and suppliers in the market place, the effect of a symbolic award is more complicated than in a setting with one principal and one agent or a setting with one principal and multiple agents (*Kosfeld and Neckermann* 2011, *Eriksson and Villeval* 2012, *Bradler et al.* 2013). For instance, if there are multiple buyers (all of them want to locate and build relationships with good suppliers), it is not clear whether these "symbolic" awards are indeed cost-free. A publicized award informs other firms about which supplier is good. If the gain from working with a good supplier is significant and/or there is scarcity in good suppliers (e.g., a good supplier is hard to come by and/or the capacity of a supplier is limited), the presence of other firms make it costly for the firm to build and retain the relationship with a good supplier. In a market with multiple buyers, a supplier may use an award as a tool to increase profit. Thus, it is possible that a "selfish" supplier could initially exert high efforts to get a favorable contract and then reap off benefit by shirking.

The main question of this paper is whether giving a symbolic award to a supplier enhances or disrupts a collaborative relationship in a competitive supply chain with multiple buyers and suppliers. Two opposite effects are central to our paper. The first is a motivational effect by which symbolic awards improve the supplier's effort. The other force is a signaling effect whereby awards intensify competition among buyers. Our main research questions are as follows: (1) Do awards have the motivational effect? (2) In a market with multiple buyers, is the competition effect strong enough to undercut the benefits of the motivational effects? (3) Considering both effects combined, when is giving out an award beneficial for the buyer in a market with multiple buyers? To answer these questions, we develop a model where we analyze four different settings: a benchmark case with no awards, a setting where the awards are private, a setting where awards are public, and a setting where the awards are private but the supplier's choice of quality is public information.

In the next sections, we analyze a stylized model and derive provable hypotheses. We then test the validity of our theoretical hypotheses with a series of lab experiments which reproduce the settings analyzed in the theoretical model. Our results show that symbolic awards have the incentive effect and lead to higher provision of quality. At the same time, the competition effect results in buyers paying higher prices to transact with good suppliers when awards are made public. We also discuss the implication of a symbolic award to the firm's profit and show that, while the symbolic award increases total profits, it does not fully restore all efficiency loss.

2.2 Literature Review

As supply chain structure becomes more complex and decentralized, one of the key issues is how to design and manage an efficient supply chain. Earlier papers in this area focus on reducing supply chain inefficiency using optimal mechanisms and/or creative contracts (see *Cachon* 2003 for an extensive review). However, both academics and practitioners well recognize that a good relationship goes beyond price and quantity negotiation (*Liker and Choi* 2004). The relationship between a supplier and a buyer is often multi-faceted, and many aspects that define a relationship cannot be explicitly written in a contract. One way to improve supply chain efficiency in such case is to recognize and utilize social attributes and preferences such as norms, fairness, and trust. *Haitao Cui et al.* (2007), *Katok and Pavlov* (2013), *Cui and Mallucci* (2012) find that fairness plays an important role in supply chain performance. *Özer et al.* (2011) find that trust and trustworthiness impact how a buyer and a supplier share information for demand planning. *Loch and Wu* (2008) show that a good relationship promotes cooperation and efficiency. These results show that incorporating social preferences has become important in understanding how supply chains can be managed more effectively.

In our paper, we focus on how a symbolic award (which is a form of feedback) affects a buyer-supplier relationship. In the economics literature, a number of papers establish that feedback on status or relative comparison incentivizes agents: *Lazear and Rosen* (1979), *Green and Stokey* (1982), and *Nalebuff and Stiglitz* (1983) study principal-agent problems in a tournament or contest and show that linking monetary reward to rank orders or relative performance outperforms the payment based on absolute performance. A stream of literature in behavioral economics shows that negative or positive feedback has motivational effects in a situation where a moral hazard or free riding exists. *Fehr and Gächter* (2000) shows that a costly (monetary) sanction reduces free-riding and increases contribution in a public goods game. *Noussair and Tucker* (2005) and *Masclet et al.* (2003) further show that even informal sanctions such as expression of a disapproval can increase contributions. Similarly, *Ellingsen and Johannesson* (2008) allow unrestricted verbal feedback in a dictator game and show that allowing for feedback reduces extreme split (e.g., 100-0) and increases the likelihood of equal splits. In a similar vein, the experiments by *Gächter and Fehr* (1999) show that verbal approval or praise induces cooperative actions. In

Bolton et al. (2004) feedback improves transaction efficiency in an online market.

The motivational effects of non-pecuniary rewards have been studied in principal-agent settings as well. Using a field experiment, *Kube et al.* (2012) show that gifts that are costly to an employer are effective in inducing reciprocal actions by employees. In particular, they show that non-monetary gifts have larger impact than monetary gifts of similar value. *Baron and Kreps* (1999) propose that workers endow a higher value to a gift if the gift is costly to the employer in labor setting. The existing work show that rewards do not have to have motivational value to motivate agents. Theoretical work by *Besley and Ghatak* (2008) shows that status rewards (such as a job title or a medal) can work as incentive devices. In a laboratory experiment, *Peeters and Vorsatz* (2013) find that sending approval via a smiley face increases contributions in the context of a public goods game. Through a field experiment, *Lacetera and Macis* (2010b) show that the frequency of blood donation increases when symbolic awards (i.e., rewarding donors with symbolic "medals") are made public. *Ashraf et al.* (2012) show that non-financial awards (e.g., a star) are more effective in increasing workers' effort than financial incentives in a public health setting. *Bradler et al.* (2013) and *Kosfeld and Neckermann* (2011) show that congratulatory cards have motivational effects on individuals performing a data entry task. While these work establishes the effectiveness of symbolic awards, we should note that all of these papers consider the case where there is a single principal (employer). In contrast, we explicitly consider the setting when there are multiple buying firms (principals) and multiple suppliers, which is a common market condition in many supply chains. When firms compete, we posit that awards may have different implications than when there is only one buying firm. The aim of this paper is to study the incentive and competition effects of symbolic awards on the provision of non-contractible quality.

The study that is perhaps most relevant to our work is *Eriksson and Villeval* (2012). They study the effects of symbolic awards on the length of employment by conducting an experiment where an employer can issue a costly award to an employee. In their experimental design, the award is only visible to the recipient (private) and is interpreted as an expression of recognition or "respect". They show that a costly symbolic award works as a coordinating device to initiate a relational contract and is used strategically by the employer to prolong a profitable relationship. They analyze different settings varying by the labor market condition (excess supply, excess demand, balanced), and find that the impact and role of an award changes depending on the market condition. For instance, when there is excess supply of labor, employers strategically use awards to prolong profitable relationships. On the other hand, when there is excess demand of labor, suppliers have a stronger bargaining power, thus awards do not play a significant role. They show that the incentive role of a symbolic award is the most significant in a balanced market.

Our paper is different from theirs in several ways. First, unlike their setting in which awards are symbolic to agents, but costly to the principal, we consider truly symbolic awards as most of these awards incur no or very little cost to the buyer (they usually take the form of a plaque of a negligible monetary cost [3]). More importantly, although we also use multiple principals and agents, we reflect the industry practice that many of these awards are announced in public (through press releases and/or events) by explicitly comparing the private and public award settings. In addition, to examine how the nature of a relationship between a buyer and a supplier changes by a symbolic award, we consider a two-period game and examine how contractual terms and efforts change over time as we vary the award setting from private to public. We find that publicizing an award affects the intrinsic value that a supplier gives to the award. In the private award scenario, the incentive effect is present in both periods, suggesting that some suppliers are motivated by intrinsic values. In the public award scenario, however, we find that suppliers exert lower efforts in the second period,

[3]See for example: http://www.ciena.com/connect/blog/Verizon-recognizes-Ciena-with-Supplier-of-the-Year-award.html and http://about.usps.com/suppliers/performance-awards.htm

indicating that publicizing awards weakens motivational role of a symbolic award . The result is consistent with a number of papers that show monetizing or publicizing the awards decrease intrinsic motivation (*Lacetera and Macis* (2010a); *Ariely et al.* (2009); *Gneezy and Rustichini* (2000)). Furthermore, we find that the transaction price for a supplier with an award increases, which partly cancels out the buyer's gain from the award. By analyzing a number of different settings, our aim is to identify when giving an award to a supplier is most beneficial to a buyer.

Some previous studies have found evidence of awards crowding out motivation. For example, a few papers have shown that monetary awards can decrease intrinsic motivation. *Lacetera and Macis* (2010a) find a substantial drop in blood donation if awards are given in form of small cash. Similarly, *Ariely et al.* (2009) find that monetary incentives have no effects on efforts made in public but they do increase efforts made in private. *Gneezy and Rustichini* (2000) find that in high-school and university students performing a task, monetary incentives decreases intrinsic motivation if awards are not sufficiently large. In the case of symbolic awards, whether higher visibility (public awards) strengthens or crowds out intrinsic motivation (relative to private awards) is still unclear. *Charness et al.* (2013) conduct an experiment where participants perform a task and show that public symbolic awards (a medal for an out-performer or a donkey hat for an under-performer) crowd out the positive effect of ranking feedback on performance. On the other hand, *Lacetera and Macis* (2010b) and *Ashraf et al.* (2012) find that increased visibility of a symbolic award allows for peer comparison and makes the award more valuable. In our setting, a public award not only allows for social comparison, but also works as a signal which grants the recipient of the award better price offers in a subsequent period. Our experimental results show that public awards are no longer valuable in a second and final period (while private awards are valuable in both periods). This suggests that public awards may be perceived by suppliers merely as a "tool" to get better contracts and as a result, publicizing the award crowds out the intrinsic value of the award to suppliers.

2.3 Theoretical Model

We first consider a two-period model with two buyers and two suppliers (we denote them supplier i, $i = 1$ or 2, and buyer j, $j = 1$ or 2). In each transaction period t ($t = 1, 2$), each buyer initiates a transaction by making an offer to a supplier. The offer consists of two parts, a price that he will pay to the supplier, p_t and the minimum effort level that the supplier needs to exert to receive a symbolic award, $\hat{q}_t$. Upon receiving the offer, the supplier decides whether to accept the offer or not. If the supplier accepts, she must decide how much effort it will exert towards quality, which is denoted by q_t. We assume that the effort is costly and the cost of effort, $C(q)$, is strictly convex and increasing in effort level, q. Once a buyer sees the supplier efforts, he can give out a symbolic award to the supplier. We note that, while effort level is observable *ex post*, it is neither contractible nor enforceable. As the requested quality $\hat{q}_t$ is only tied to a symbolic award, the supplier is by no means obligated to comply for monetary reason. Once the buyer observes the effort, he can choose to give or not give an award to the supplier at his discretion.

In the first transaction period, each buyer is initially matched with one supplier and can make an offer only to the matched supplier. In the second transaction, both buyers are free to trade with any of the two suppliers. To avoid plethora of equilibria, we assume that the buyer chooses one of the three actions: (1) making a new offer to the other supplier, (2) making a new offer to the existing supplier, (3) keep the existing offer to the existing supplier (we call this a continuing offer) with an option to match any competing offer from the other buyer. In addition, if a buyer fails to secure a deal with a supplier (e.g., his offer is either out-bid or rejected), he is free to make a new offer to a different supplier.

We assume that there are two types of suppliers in the marketplace, different in how they value a symbolic award. A "low (motivational) type" supplier has no intrinsic value for a symbolic award, therefore, she chooses an action that maximizes his monetary profit over two periods (in each period, her profit is the price minus the cost of the effort). On the other hand, a "high (motivational) type" supplier draws intrinsic value from winning a symbolic award. Thus, the utility of a high-type supplier accounts for both monetary payoffs and transaction utility from winning an award. To formalize this, let ϕ be the transaction utility a high type supplier draws from an award. Then, the high-type supplier's utility is given by

$$U^h(A_t, q_t) = \sum_{t=1}^{2}[p_t + \phi A_t - C(q_t)] \tag{2.1}$$

where A_t is an indicator function that takes value 1 if the supplier gets the award in period t. The low-type supplier's utility, $U^l(\cdot)$ is just the monetary profit over two periods. We assume that a supplier is a high-type with probability π, and the type is private information to the supplier . We also assume that the type of supplier 1 is independent of the type of supplier 2.

On the other hand, the payoff that the buyer receives from supplier's effort, q_t, is αq_t, where α is an *efficiency coefficient*. Since the award is costless for the buyer, the buyer's utility is given as follows

$$U^B(p_t, \hat{q}_t | q_t) = \sum_{t=1}^{2}[\alpha q_t - p_t]. \tag{2.2}$$

To draw hypothesis, we analyze the equilibrium in four different scenarios. In the no-award case (denoted by NA), no award is allowed. In the private award case (PrA), a buyer can give a private award, thus the award is not known to the other buyer and supplier. In the public award case (PuA), a buyer can give an award, which will be announced in public. Finally, in the public quality case (PuQ), while the award is private, the supplier's effort level becomes public information.

We first analyze the no-award case and show that, in equilibrium, no supplier exert efforts. All proofs are relegated to the appendix.

Proposition II.1. *If awards are not available (scenario NA), all suppliers exert zero effort in both periods regardless of the offer they receive in equilibrium.*

To see why this is the case, note that effort is not enforceable, thus it is a dominant strategy for all suppliers to exert no effort in the last period. Anticipating this, no buyer will pay positive prices in equilibrium. Since there is no award, both types of suppliers act exactly the same. Establishing this as a benchmark, the next result shows that a high-type supplier will act differently from a low-type supplier if a symbolic award is given.

Proposition II.2. *In the private award scenario (PrA), there exists a separating equilibrium in which a low-type supplier exerts no effort and does not receive an award while a high-type supplier exerts strictly positive effort and receives an award in both periods.*

We show that, in the private award setting, a separating equilibrium in which a low-type supplier puts no efforts and does not get the award, and a high-type supplier gets the award in both periods by meeting the requested quality ($q_t = \hat{q}_t$) arises. In this equilibrium, the buyers have the first mover's advantage which allows them to set the requested quality level that extracts the surplus of a high-type supplier, which results in a high-type supplier fully working for the intrinsic value of an award. Because of this reason, all truth-telling equilibria have zero price in both periods. In contrast, the nature of an equilibrium changes when the award becomes public.

Proposition II.3. *The price in the public award scenario (PuA) is higher than the price in the private award case (PrA) on average.*

In particular, in the public award scenario we show the existence of an equilibrium where if one of the suppliers is of high type and the other supplier is of low type, price in period 2 is strictly positive. The equilibrium is fully characterized in the proof presented in the Appendix.

The difference between the private and public awards cases is that, when the award is private the buyers cannot update their beliefs about the other supplier's type after transaction period 1. Therefore, making an offer with positive price to the other supplier in period 2 would earn the buyer negative expected profits: with probability π the other supplier is of high type and any profitable offer would be matched by the other buyer, and with probability $(1 - \pi)$ the other supplier is of low type and always chooses zero quality in period 2. On the other hand, when the award is public, a buyer can update his belief about the other supplier's type after observing whether he received an award in period 1. In a truth-telling equilibrium, a buyer who is matched with a low type supplier knows the other supplier's type with certainty and can choose to make him an offer with strictly positive price (but low enough that it is still profitable) only when the supplier is of high type. In this case, the other buyer will match that offer to keep the high type supplier, which will result in equilibrium prices greater than zero.

Proposition II.4. *In both private and public award scenarios, on average, a buyer's profit is higher than in the no-award case. The buyer's profit in the public quality scenario is is the same as the price in the public award setting (PuA). Depending on the values of the parameters ϕ and π, a buyer's expected profit can be higher with private awards or with public awards.*

In the proof, we show that both the private and public award scenarios (PrA and PuA) result in higher expected buyer's profits than the no-award case (NA). To see why, first note that both the no-award and the private awards cases have prices equal to zero in all possible equilibria, but while the no-award case has zero quality in both periods, the private award scenario can result in equilibria with positive quality (Proposition II.2). Secondly, while in the no-award case buyers earn zero profits, we show that in the public awards scenario an equilibrium exists (the one characterized in Proposition II.3) where buyers' expected profits are strictly positive.

In addition, we show that in truth-telling equilibria, the scenario with public quality and private award (PuQ) is equivalent to the scenario with public award and private quality (PuA). The only difference is that rather than updating their beliefs about the suppliers' types based on whether they received an award in period 1 or not, buyers update their beliefs based on the suppliers' choice of quality in period 1. The formal proof is presented in the Appendix.

Based on these theoretical results, we derive hypotheses for the incentive and competition effects to be tested with a laboratory experiment.

2.4 Experimental Design

Each session consisted of six plays of the two-stage game. We refer to each play as a round. We used the z-Tree software package (*Fischbacher*, 2007), developed at the Institute for Empirical Research in Economics at the University of Zurich.

In each session of the experiment, subjects play six rounds of a computerized game. In each round subjects are matched in groups of four. The groups remain together for the entire round and are dissolved when the round is over. New groups are formed for the following round. At the beginning of each round, the members of a group are assigned a role (two subjects will play as buyers and two subjects will play as suppliers) and a label (Buyer 1, Buyer 2, Supplier 1, Supplier 2) which they will keep for the entire round.

Each round of the game consists of two transactions. For the first transaction, each buyer is randomly paired up with a supplier in the same group and can only make offers to that supplier. The first transaction consists of the following four stages: 1) The buyer makes an offer to the supplier. The offer consists of a price he is willing to pay for the good, which can range between 0 and 100 and a requested quality, which can range between 0 and 10. 2) The supplier decides whether he accepts the offer or opts out of the game for the first transaction. 3) If the supplier accepts the offer, he chooses a quality level for the good he provides. The buyer's expected quality level from the contract is non-binding so the supplier can choose quality freely. Quality can range between 0 and 10 and is costly for the supplier. We assume that the cost of quality takes the canonical form, $c(q) = \frac{1}{2}q^2$. In order to simplify the subjects' task, they were presented with table 2.1 on their screens. 4) If awards are available, the buyer sees the quality chosen by the supplier and decides if he gives him a symbolic award. Whether the award is only shown to the supplier or if it is shown to everyone varies by treatment. After the first transaction, suppliers and buyers are no longer matched, but the groups of four subjects remain together.

Table 2.1: Cost of Quality Function

q	0	1	2	3	4	5	6	7	8	9	10
$c(q)$	0	0.5	2	4.5	8	12.5	18	24.5	32	40.5	50

In the second transaction, buyers can make an offer to any of the two suppliers in their group. The stages are as follows: 1) The buyer chooses between the following options: a) Make a continuing offer to the same supplier, in which case the offer (price and requested quality) remains the same as in transaction 1. This option gives the buyer the right to match any offer the other buyer makes to his supplier. b) Make a new offer to the same supplier with different price and requested quality than in transaction 1. In this case, the buyer does not have the opportunity to match the other buyer's offer to his supplier. c) Make an offer, of any price and any requested quality, to the other supplier. 2) If a buyer made a continuing offer to his supplier and the other buyer also made an offer to this same supplier, the former can choose between matching the other buyer's offer and keeping his original continuing offer. In either case, the supplier receives both buyers' offers. 3) The supplier observes the offers he received and accepts the offer of one of the suppliers or opts out of the second period of the game. 4) If the supplier accepts an offer, then he chooses a quality level for the good he provides. 5) In the case where both buyers made an offer to the same supplier, the buyer who was not picked can make an offer to the other supplier. He makes a new offer (price and requested quality) and the supplier decides whether he accepts the offer and if so a quality level. 6) If awards are available, each buyer decides whether they give the supplier a symbolic award. When the second trading period is over, new groups are formed for the following round and subjects are relabeled.

The symbolic award is a picture of a ribbon that appears on the screen of the supplier who is granted the award. Giving an award has no cost for the buyer and does not grant extra profits for the suppliers. The award does not affect the monetary payoffs of the participants in any way. We conducted the following four treatments. The first treatment is a No-Award/Private Information case where awards are not available and the quality choices are only shown to the respective buyer and supplier. The second treatment is a Private Award/Private Information case in which only the corresponding buyer and supplier can observe the quality chosen and the award, if there is one. The third treatment is a Public Award/Private information treatment in which the award can be seen by everyone but the quality level can only be observed by the corresponding buyer and supplier. The last treatment is a Private Award/ Public information case in which only the supplier who gets the award and the buyer who gives it can see it, but the information about quality is public to everyone. The treatments are summarized in Table 2.2.[4]

The payoffs are expressed in tokens and converted into dollars (at a rate of \$0.05/token) for payment at the end of the experiment. One of the ten rounds played is randomly chosen for payment. In each round, the subjects' payoff is the sum of their payoff in both transactions. The supplier begins each round with 60 points,

[4]The stage in which suppliers may receive a symbolic award, is only present in the second, third and fourth treatments.

Table 2.2: Experimental Design - Treatments Summary

Treatment	Award	Information about Quality
No Award (NA)	N/A	Private
Private Award (PrA)	Private	Private
Public Award (PuA)	Public	Private
Public Quality (PuQ)	Private	Public

and the retailer with 100 points, to avoid negative payoffs. For each transaction, the payoffs are calculated as follows: the buyer's payoff is 100 plus ten times the quality chosen by the supplier minus the price he pays, $\pi^B = 100 + 10 * q - p$, or $\pi^B = 0$ if his offer was rejected. The supplier's payoff is 60 plus the price he gets for the good minus his cost of quality, $\pi^S = 60 + p - c(q)$, or $\pi^S = 0$ if he rejects all offers and opts out of the trading period. Subjects receive a $7 participation fee, plus their earnings from the experiment (which can range between $0 and $20).

2.5 Hypotheses

We derive hypotheses for the incentive and competition effects of symbolic awards based on the previous theoretical results. The first hypothesis predicts that there exists a criterion with which buyers give out symbolic awards. From Proposition II.2 we know that, when awards are available, a separating equilibrium arises in which high type suppliers choose a quality equal to the requested quality in order to get the award. In this equilibrium, it is a dominating strategy for the buyer to set the requested quality equal to the maximum quality high type suppliers are willing to offer in order to get the award. If he offers awards universally, high type suppliers have no incentive to offer high quality. In the experiment, we expect buyers' decision to give the award to be positively correlated with a higher provision of quality from the supplier.

HYPOTHESIS 1. [Award Criterion] *Buyers offer awards predominantly to suppliers who meet or exceed the requested quality.*

Based on Propositions II.1 and II.2, the second hypothesis predicts that suppliers will provide higher quality in the treatments where awards are available. When awards are available, a separating equilibrium arises in which buyers request some positive level of quality to give out the award. Thus, suppliers who care about the award, are willing to provide additional positive quality in order to receive one.

HYPOTHESIS 2. [Quality Effects of Awards] *The average quality when awards are available is higher than in the no-awards case.*

The third hypothesis predicts that, due to the competition effect, buyers pay higher prices to good suppliers when awards are public. This prediction is based on Proposition II.3, which states that in the public awards case expected price is at least as high as in the private award case and there always exists an equilibrium where it is strictly higher. The equilibria which arise in the public awards case where price is strictly higher than in the private awards case, have one high type supplier who receives an award in the first transaction. Because the award is observable to all players, in the second transaction both buyers compete to win the deal with the high type supplier. As a result, the winning buyer pays a higher price than if the awards are private or if awards are not available.

HYPOTHESIS 3. [Competition Effects] *The expected price paid to good suppliers in the second transaction when the award is public is higher than when awards are private or unavailable.*

Lastly, the fourth hypothesis makes predictions about buyers' profits based on Proposition II.4. According to the theoretical results, buyers' profits are higher when awards are available than when they are not. In the private awards case, the only effect which can potentially be present is the incentive effect. Thus, an increase in buyers' profit relative to the no-award case would capture how the buyer benefits from the incentive effects generated by the award. Additionally, the theory predicts that, with public awards average quality in the first transaction and average price in the second transaction both increase relative to the private awards case. As a result, buyers' profits can be higher or lower depending on the magnitude of these two differences. We represent this by alternate Hypotheses (4.a and 4.b). The public quality case is expected to be equivalent to the public award case, based on Proposition II.4.

HYPOTHESIS 4. [Profits] *Buyers' profits are higher when awards are available than when they are not. The public award and public quality cases have equal buyers' profits.*
4.a - Buyers' profits are higher when are awards are private than when they are public.
4.b - Buyers' profits are higher when are awards are public than when they are private.

2.6 Experimental Results

The experiment was conducted at the University of Michigan in the months of March and April of 2013. The subjects were undergraduate students from a variety of Departments at the University. We conducted sixteen sessions of the experiment, with eight or twelve subjects per session. We had 40 subjects for the no award treatment, 40 subjects for the private award treatment, 36 subjects for the public award treatment, and 36 subjects for the public quality treatment.

In each session, subjects played six rounds of the game. Sessions lasted on average an hour an a half and in the end, one of the six rounds was randomly selected for payment. Subjects made on average between 15 and 16 dollars.

2.6.1 General Results

Table 2.3 summarizes our results. We observe that in the three treatments where awards were available, buyers chose to give awards between twenty and fifty percent of the times. Average accepted price ranged between 30 and 40 points and it appears to be higher in the public award and public quality treatments, although we find this increase to be statistically significant only in the public quality treatment. Requested quality seems to be very stable across treatments. In the public quality treatment, requested quality is significantly higher than in the other three treatments. Average quality appears to be the lowest in the no-award treatment and the highest in the public quality treatment, both in period 1 and period 2. In the private and public award treatments it ranges between these two values. Similarly, the probability of suppliers choosing a quality level greater or equal to the buyer's requested quality level increases significantly in the three treatments where awards were available relative to the no-awards case.

Table 2.3: General Results - Period 1

	Treatment	Price offers	Acceptance (%)	Price	Quality	Award (%)
Period 1	No Award	29.1	80.3	31.4	2.198	
	Private Award	25.4	75	30.3	2.767	0.42
	Public Award	31.8	85.19	32.1	2.848	0.47
	Public Quality	38	79.63	40.4	3.872	0.49
Period 2	No Award	25.3	81.1	31.3	2.084	
	Private Award	23	80.8	28.5	2.486	0.33
	Public Award	27.6	77.8	35.5	2.262	0.41
	Public Quality	32.1	75	42.8	3.185	0.37

2.6.2 Incentive Effect

The analysis of the incentive effects has two components: first, we observe whether buyers give awards to those suppliers who choose higher quality, and then we test whether suppliers choose higher quality when awards are available. We study these two effects in the following subsections.

2.6.2.1 Award Criterion

As predicted by our first hypothesis, buyers give awards to those suppliers whose quality choice meets or exceeds their expectations. Figure 2.1 shows the probability of a buyer giving out an award as a function of the difference between the quality chosen by the supplier and the buyer's requested quality. In both periods we observe a significant climb at zero (that is, when quality equals requested quality), where the probability of receiving an award increases from around 20% to close to 100%. These results are presented formally in Table 2.4. We present the regression of the probability of receiving an award on the choice of a quality level that is greater or equal to the requested quality separately for each treatment. The coefficients are positive and significant in all cases where awards are available, both in transactions 1 and 2. This indicates that meeting or exceeding the buyer's quality expectations leads to higher probability of receiving an award in all treatments.

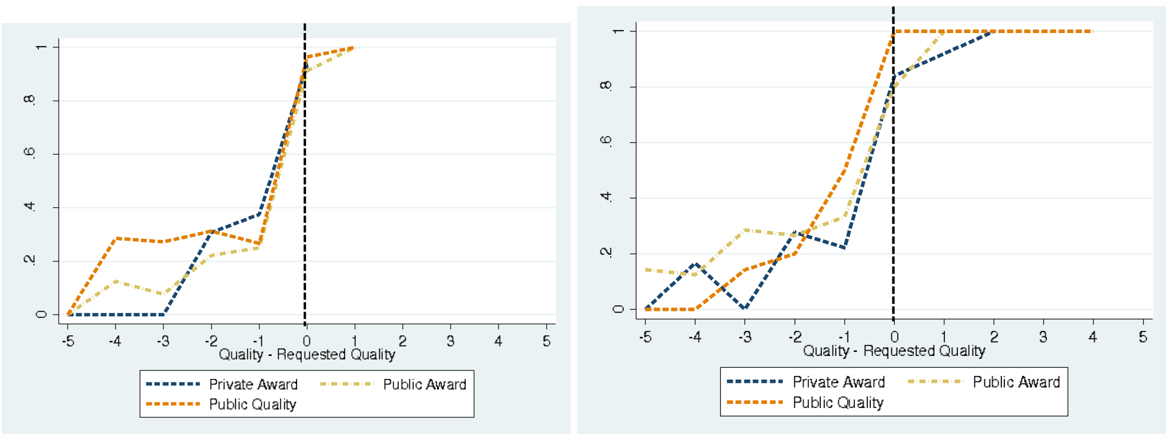

Figure 2.1: Probability of receiving award

Table 2.4: Interaction effects of quality and treatment on the award decision

Coefficients	Award Period 1	Award Period 2	Award Period 1	Award Period 2	Award Period 1	Award Period 2	Award Period 1	Award Period 2
(Q≥RQ)* x PrA	0.748***	0.717***					0.772***	0.719***
	(0.074)	(0.078)					(0.069)	(0.078)
(Q≥RQ)* x PuA			0.756***	0.624***			0.761***	0.632***
			(0.072)	(0.077)			(0.070)	(0.078)
(Q≥RQ)* x PuQ					0.752***	0.845***	0.719***	0.819***
					(0.070)	(0.060)	(0.073)	(0.059)
PuA							-0.005	0.031
							(0.068)	(0.082)
PuQ							0.073	0.019
							(0.078)	(0.087)
Price	-0.003	0.001	0.0001	0.005**	0.004**	-0.002	0.001	0.002
	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)
Constant	0.262***	0.142	0.158*	0.044	0.066	0.286***	0.141**	0.126**
	(0.091)	(0.087)	(0.090)	(0.089)	(0.070)	(0.091)	(0.064)	(0.061)
Observations	90	97	92	84	86	81	268	262
Nr. of Subjects	38	38	36	35	35	33	109	106

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Note: Probit estimates show similar results but present high standard errors in columns 6 and 8 due to lack of observations for the case "Quality Higher than Requested Quality and No Award" in period 2 of the public quality treatment. Marginal effects from Probit regressions are 0.774, 0.875, 0.756, 0.726, 0.840, 0.900 for columns 1 to 6 respectively. Note:(*) ($Q \geq RQ$) is a dummy variable which takes value one if Quality is greater than Requested Quality, and zero otherwise.

2.6.2.2 Quality Effects of Awards

Our second hypothesis predicts the existence of an incentive effect. In accordance with Hypothesis 2, we expect to see that when awards are available, average quality is greater than in the no-awards case. Table 2.5 presents statistical results that confirm this prediction. The first two columns correspond to the comparison between the no-award and private award treatments only. We regress quality on a dummy for the private award treatment controlling for price. We find that the private award treatment leads to higher quality in both transactions (regression coefficients equal to 0.708 and 737 in transactions 1 and 2 respectively). Since the only difference between these two treatments is the opportunity to give an award to the suppliers and since the award was private, we attribute this increase in quality exclusively to the supplier's value for the symbolic award. Columns three and four in Table 2.5 compare the no-award treatment with the public award treatment. Interestingly, the public award treatment leads to higher quality than the no award treatment in the first transaction but not in the second transaction. In transaction 1, the coefficient is 0.669 and it is statistically significant, while in transaction 2 it is 0.0075 and it is no longer significant. Similar results are found in columns five and six when comparing the no-awards and the public quality treatments (in transaction 1 the coefficient is 0.919 and significant, while in transaction 2 it is 0.261 and not significant). These results suggest that when the award is made public, suppliers perceive it as a tool to get better contracts in the second transaction, and the award looses its intrinsic value. Therefore, in the second transaction, suppliers are no longer willing to exert higher effort to get an award. This "crowding out" of the incentive effect when the awards are public is consistent with previous field experiment studies which found that public rewards (both monetary and non-monetary) can have detrimental effects on performance when completing a task (*Gneezy and Rustichini* 2000, *Harackiewicz* 1979, *Charness et al.* 2013) or in contributions for charity (*Lacetera and Macis* 2010a, *Ariely et al.* 2009). Two explanations have been used to explain this phenomena. First, individuals may not want to be perceived as naive and give the impression that they are willing to increase effort in exchange for some non-tangible reward. Second, extrinsic incentives make it harder for the supplier to show that he gives high quality because he is "good". That is, offering higher quality is a less clear signal that the supplier is "good" if they receive an award for it.

Table 2.5: Effects of Treatments on Quality

	PrA vs. NA		PuA vs. NA		PuQ vs. NA		All vs. NA	
Coeff	Quality Period 1	Quality Period 2	Quality Period 1	Quality Period 2	Quality Period 1	Quality Period 2	Quality Period 1	Quality Period 2
PrA	0.708*	0.737**					0.721*	0.697**
	(0.402)	(0.346)					(0.406)	(0.348)
PuA			0.669*	0.0075			0.663*	-0.013
			(0.361)	(0.402)			(0.36)	(0.404)
PuQ					0.919**	0.261	0.998***	0.392
					(0.404)	(0.43)	(0.386)	(0.423)
Price	0.0510***	0.0738***	0.0616***	0.0574***	0.0779***	0.0768***	0.0688***	0.0639***
	(0.011)	(0.009)	(0.009)	(0.01)	(0.008)	(0.008)	(0.006)	(0.007)
Cons	0.617*	-0.283	0.278	0.231	-0.245	-0.383	0.0463	0.0339
	(0.328)	(0.291)	(0.269)	(0.332)	(0.25)	(0.273)	(0.271)	(0.295)
Obser	196	204	198	191	192	188	374	369
# of S.	81	82	78	79	79	77	150	150

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

2.6.3 Competition Effect

The third hypothesis predicts that when either the award or the quality are public, buyers pay higher prices to the "good" suppliers in the second transaction. We define distinguish "good" suppliers in two ways: first, we consider those suppliers who received an award in transaction 1 and second, we consider those suppliers who chose in transaction 1 a level of quality that was above the median for their treatment. Figure 2.2 shows for each of the four treatments, a first column with the average price paid to all suppliers in the first transaction, a second column with the average price paid in transaction 2 to those suppliers who meet the definition of "good", and a third column with the average price paid in transaction 2 to those suppliers who do not meet the definition. We observe that in the treatments where either the award or quality were public, buyers paid higher prices to good suppliers in transaction 2. This result supports the third hypothesis which proposes that making the award public increases competition for the good suppliers and makes buyers have to pay more for them.

Further analysis on the competition effect is presented in Table 2.6. It shows the regression of price in transaction 2 on having received an award in the first transaction separately for each treatment. The first column presents all price offers and the second column presents accepted offers only. We observe that when awards are public, receiving an award in the first transaction has a stronger impact on the price offers received in the second transaction that when the awards are private (coefficients: 8.261 and 17.30 for the private and public award treatments respectively, both statistically significant). When considering only the accepted offers, we find that the transaction prices in period 2 are also higher with public awards than when the awards are private (coefficients: 5.44 and not significant in the private awards case, and 10.32 and significant with public awards). These results reinforce the support for the competition effect hypothesis.

2.6.4 Profits and Efficiency

Hypothesis 4 predicts that buyers' profits should be higher with awards (either private or public) than when awards are not available. We find that buyers' profits are only higher with private awards. When awards are public, there is no significant difference with the no-awards case. Average profits are presented in Table 2.7 and regressions with the effects of the different treatments on profits are presented in Table 2.8. Buyers' average profit is 195.1 with private awards and 184 in the no awards

Table 2.6: Interaction Effects of Award Treatments on Price

Coefficients	Price Offers in period 2	Transaction price in period 2
Award x Private Award	8.261**	5.44
	(4.04)	(4.36)
Award x Public Award	17.30***	10.32**
	(4.45)	(4.66)
Award x Public Quality	16.27***	7.883*
	(4.78)	(4.3)
Public Award	-2.01	4.023
	(4.49)	(5.1)
Public Quality	5.177	12.29**
	(5.15)	(5.02)
Constant	21.60***	25.98***
	(3.08)	(3.41)
Observations	268	221
Nr. of Subjects	106	100

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

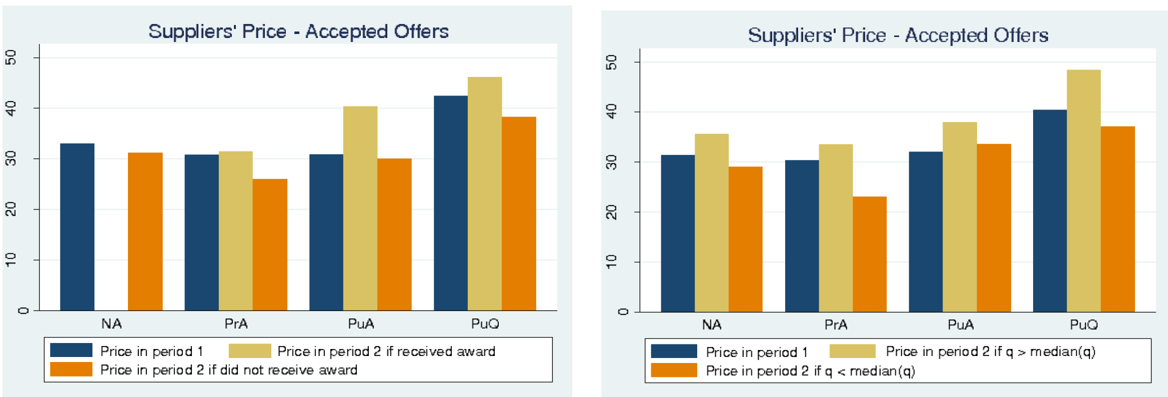

Figure 2.2: Transaction Price

case. In addition, the coefficient of a private award treatment dummy shows an effect of 11.6 points (p-value < 0.01) relative to the treatment with no awards. In the public award case, average buyers' profit is 186.5 and the coefficient for the treatment dummy shows no significant difference with the no award case. This suggests that the incentive effects in isolation have a positive impact on buyers' profits and that the additional effects which result from making the award public dilute this profit premium.

The theoretical results predict on average lower quality in transaction 1, lower price in transaction 2, and equal quality in transaction 2, in the private awards treatment relative to the public awards treatment. As a result, buyers' profits could be higher in either of these two treatments depending on the magnitude of these differences. We find that buyers' profits are significantly higher under private awards than with public awards, providing support for Hypothesis 4.*b* (further evidence is presented on Table B.1 in the Appendix). This is because, while in the public award treatment price is higher in transaction 2 (and in particular prices are higher for "good suppliers"), quality does not significantly increase in transaction 1. Furthermore, due to the crowding out effects, quality in transaction 2 is even lower than with private awards resulting in lower buyers' profits.

Finally, Hypotheses 4 predicts that the public award and public quality treatments are equivalent. The experimental results show that buyers' profit is not different with public quality than with public awards. In Table 2.8 we observe that the coefficients for the effects of the public award and public quality treatments on buyers' profits are 3.312 and 6.946 respectively and not significantly different (p-value = 0.435).

Tables 2.7 and 2.8 also present suppliers' monetary payoffs. We do not derive formal hypotheses for suppliers' profits since, unlike buyers' profits, these are an imperfect measure of total utility as they do not capture the utility generated by the awards. The experimental results show that, contrary to what happens with buyers' profits, suppliers' profits are lower when awards are private than when awards are public and when awards are not available (156 with private awards, and 164.8 and 162.8 with public and no awards respectively). This result indicate that the private awards treatments leads simultaneously to the highest buyers' profit and the lowest suppliers' profit. An interesting question then is how total profits change by treatment.

Table 2.7: Buyers and Suppliers Profit

	Average Profit - Buyers			Average Profit - Suppliers		
Treatment	Period 1	Period 2	Periods 1 + 2	Period 1	Period 2	Periods 1 + 2
No Award	92.439	91.55	184	81.44	81.34	162.8
Private Award	98	97.05	195.1	77.7	78.32	156
Public Award	96.944	89.97	186.9	81.49	83.3	164.8
Public Quality	98.6	91.8	190.4	83.53	85.5	169

Table 2.8: Effects of Treatments on Profit

	Buyers Profit			Suppliers Profit		
Coefficients	Period 1	Period 2	Periods 1 + 2	Period 1	Period 2	Periods 1 + 2
Private Award	5.853**	5.652**	11.597***	-4.079	-3.126	-7.433*
	(2.721)	(2.372)	(3.942)	(2.617)	(2.219)	(4.056)
Public Award	4.816**	-1.519	3.312	-0.074	1.79	1.479
	(2.398)	(3.174)	(4.581)	(2.648)	(2.658)	(4.136)
Public Quality	6.438***	0.416	6.946*	2.095	4.165	6.279
	(2.205)	(2.805)	(3.885)	(2.727)	(2.842)	(4.929)
Constant	92.086***	91.446***	183.488***	81.547***	81.393***	163.033***
	(1.626)	(1.858)	(2.685)	(1.937)	(1.389)	(2.986)
Observations	468	468	468	468	468	468
Nr. of Subjects	155	155	155	155	155	155

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 2.9 presents the results for total surplus. The first three columns show the buyers' dyads[5] and the fourth column shows the sum of the profits of all four members

[5]We calculate total surplus across the two periods keeping the buyer fixed. That is if a buyer was matched with one supplier in period one and with the other supplier in period two, we calculate total surplus as the sum of the buyer's profits in both periods plus the sum of the profits of each of the two suppliers in the respective periods in which they played with this supplier

Table 2.9: Total Surplus

Treatment	Total Surplus - Buyers' dyads Period 1	Period 2	Periods 1 + 2	Total Surplus per Group Periods 1 + 2
No Award	173.845	172.8	345.4	690.7
Private Award	175.492	175.4	350.9	701.7
Public Award	178.435	173.1	351.5	703
Public Quality	182.069	177.3	359.4	718.7

Table 2.10: Effects of Treatments on Total Surplus

Coefficients	Total Surplus - Buyers' dyads Period 1	Period 2	Periods 1 + 2	Total surplus per group Periods 1 + 2
Private Award	1.258	2.377	4.576	12.75
	(2.366)	(2.396)	(4.3)	(8.98)
Public Award	4.363**	-0.022	5.356	14.24*
	(2.176)	(2.388)	(4.2)	(7.73)
Public Quality	8.335***	4.647*	14.160***	31.94***
	(2.506)	(2.615)	(4.61)	(8.35)
Constant	173.934***	172.905***	345.742***	688.5***
	(1.433)	(1.54)	(2.8)	(5.83)
Observations	468	468	468	234
Nr. of Subjects	155	155	155	129

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

of a group. Average total surplus is only significantly higher in the public quality treatment (at a group level, average total profits were 718.7 in the public quality treatment versus 703 or less in all other treatments). Similar results are found in Table 2.10, which shows the regression of total profits on treatment dummies. The public quality treatment has an increase in total surplus of more than 30 points relative to the no-awards case and this coefficient is statistically higher than those of the private and public award treatments. In the public quality treatment, buyers can observe the actual suppliers' quality choice in the first transaction, as opposed to some imperfect signal (the award). This leads to higher prices and higher quality benefitting both buyers and suppliers [6].

Finally, we note that average total surplus is not significantly different in the no award and private awards treatments. This is because due to the incentive effects, in the private award treatment buyers' profits increase and suppliers' profits decrease relative to the no awards treatment. Similarly, total profits are not statistically different in the private and public awards treatments. This indicates that making the awards public results in a transfer of profits from buyers back to the suppliers, and it is not a loss in efficiency.

2.7 Discussion

Our results contribute to the discussion of where the value of symbolic awards stems from. *Frey and Neckermann* (2008) identify several reasons why awards may have motivational effects: because they make the recipient feel good about about himself, because the recipients value the opinion of the authority who gives out the award, because the awards generate social prestige and bring recognition within the

[6]This result is consistent with previous literature that shows that clearer signals of performance lead to higher efficiency. For example, *Bolton et al.* (2004) find experimentally that while online feedback systems have positive effects on transaction efficiency between strangers, they cannot fully capture the benefits of reputation building in markets where the same people interact with each other repeatedly.

peer group, or because some subjects enjoy joining the competition for the award. The first and second explanations are more salient in cases where the awards resemble feedback or praise, as it is usually the case in supplier awards. If workers in the supplier firm care about their hard work being recognized by the firm's clients, managers may seek to get these awards to motivate their workers to provide higher effort. The third and fourth explanations are usually associated with competition prizes. In line with this last group, some previous experiments find that awards that are scarce generate status and have higher motivational effects the more visible they are. In practice, we observe that supplier awards do not have the feature of being particularly scarce. For example, in 2013 USPS granted seven supplier awards (in different categories such as innovation, diversity, sustainability). In year 2014, Whole Foods Market gave twenty supplier awards including best partnership, best new product, product quality, innovation, etc. In our setting both suppliers may simultaneously get an award, thus there is no competition for the award between suppliers. Our experimental results show that making the awards public, rather than increasing the value of the award, crowds out their motivational value.

Another possibility is that suppliers care about receiving the award since it is an indicator of a higher probability of continuing relationship. This explanation seems plausible; our results show that when awards are available, receiving an award leads to a continuing offer in 52% of the cases compared to only 18% when suppliers do not receive an award. However, this argument is not sufficient to fully explain the outcomes we obtain from our experiment. In the private awards treatment, suppliers exert higher effort than in the no-award scenario both in transactions 1 and 2. Concerns about the continuation of the relationship fail to explain higher quality in the last transaction.

Introducing private awards in our experimental design allows us to disentangle the incentive effects, and furthermore enabled us to identify crowding out effects when awards are public. While useful in the laboratory, keeping an award private is usually hard to accomplish in most industry contexts. Awards are typically visible, either because they are announced at a public ceremony or because the award itself can be publicly displayed through websites and press releases. Therefore, it is common to observe firms offering informal recognitions which are harder for the recipient to publicize. For example, a great honor for L'Oreal's suppliers is to sit at the table of the CEO of L'Oreal at their biggest annual event. These types of informal rewards may still carry the incentive effects of symbolic awards, without the downside of the competition effects.

In our experimental design, the public quality case resembles full transparency and serves as a benchmark case for analyzing the awards as a signal of supplier's quality. In reality however, it may be harder for either the buyer or the supplier to control how visible the supplier's quality is. For example, while in some industries it is common to resort to "consumer reports" where third parties validate whether a certain service level is met, in many cases suppliers who provide components are not evaluated by third parties nor have a strong brand that allows others to track their performance.

2.8 Conclusions

We observed that, when available, buyers frequently choose to give out symbolic awards to those suppliers who meet or exceed their quality expectations. Furthermore, they persist to do so with the same frequency when the award is public, in spite of the fact that it may signal that the supplier is good and increase competition for the good supplier. These symbolic awards have shown to have value for the suppliers and induce them to exert higher effort to get an award, even in a second (and final) period. However, we also observed that in the case where awards are public, subjects are no longer willing to exert higher effort to get an award in the second transaction. This suggests that making the award public crowds out the intrinsic value of the award. As the award becomes a “tool” for the awarded suppliers to get better contractual terms in a following transaction, suppliers seem to loose the appreciation for the award itself.

We find evidence that symbolic awards also affect competition. When buyers give out public awards, they are signaling to other buyers that a supplier is good. By doing so, they trigger competition and they end up having to pay more for the good suppliers in the second transaction.

Buyers’ profits are higher with private awards than when awards are not available, which suggests that the incentive effect has a positive impact on buyers’ profits. This profit premium due to the incentive effect is not accompanied by an increase in efficiency. Suppliers’ monetary profits are (marginally) lower and, as a result, total profits are not significantly affected by the incentive effects.

The buyers’ profit premium in the private awards case, is no longer present when awards are instead public. This is because with public awards buyers pay higher prices in transaction 2 (in particular, they pay more to good suppliers) and do not receive higher quality in response. Average quality in transaction 1 does not increase and, due to the crowding out effect, quality in transaction 2 even decreases relative to the no-award case. While buyers’ profits decrease, suppliers’ monetary profits increase relative to the private awards case. As a result, the buyers’ profit loss from making the awards public, is not a loss in efficiency but a transfer to suppliers.

Finally, total surplus is significantly higher only when awards are private and quality is public. While a public award can signal the supplier’s type, this signal carries noise and requires certain level of interpretation from the buyers. When the quality is public, the information is clear and no signal is necessary, which reduces the inefficiency derived from asymmetric information. As a result, average price and average quality are higher in both transactions making both buyers and suppliers better off.

These results explain the current trends relative to supplier awards observed in the industry. We observe that suppliers strive to loudly announce their awards, while buyers devise creative ways of generating informal awards which, have the ability of inducing incentive effects, but are also hard to publicize. Our results suggest that when implementing a "Supplier of the Year Award" initiative, firms should carefully contemplate the visibility of the award as it may have negative effects on the initiative's profitability. First, by increasing competition for good suppliers (and therefore the competitive nature of their particular industry is relevant) and second, by crowding out the incentive effects to provide higher quality to get the award.

CHAPTER III

The Impact of Decision Rights and Long Term Relationships on Innovation Sharing

3.1 Introduction

Manufacturers often benefit from innovations and process improvements discovered by their suppliers. This happens particularly in industries where suppliers are involved in important parts of the manufacturing process, such as new product development. Technologies developed by suppliers resulted in cost reductions for the manufacturer in the automotive industry for example, where suppliers take a significant role in R&D[1]. *Klier* (2006) describes the case of one of General Motors' door hinges supplier who allowed GM to get significant cost savings by simplifying the hinge design, making it stamped rather than welded. Another example is from one of Chrysler's suppliers, Becker Manufacturing Inc. Becker started molding pre-existing hooks onto the interior trim panels that are attached to the car's door body. This allowed the door panels to be fastened to the door frame by the pre-existing hooks, eliminating excessive fastener parts and screws and resulting in major savings for Chrysler [2]. Since these discoveries usually happen spontaneously, the supplier has no

[1]Neil De Koker, president of the Original Equipment Supplier Association reported in 2006 that in the automotive industry, suppliers are taking a bigger role in R&D, providing up to two thirds of the value added in the production of the car (*Klier* 2006).

[2]Source: http://www.allpar.com/corporate/score.html.

contractual obligation to share them with the manufacturer. Therefore, an important question for manufacturers is how to incentivize suppliers to share these innovations with them[3].

Suppliers can benefit greatly from sharing innovations with a manufacturer. For example, a recent trend in the automotive industry was to transition from the usage of solvent based paint to waterborne paint, which is less toxic and of easier disposal and cleanup. Both GM and its suppliers were transforming their paint booths as part of the transition. Implementing this technology involved making changes to processes and learning new procedures - for example, painters had to be retrained to paint more evenly, new techniques were needed for taping up to prevent bleeding, new equipment had to be installed to blow large volumes of clear air to enhance drying, etc. If the suppliers discovered a better way of implementing this new technology and shared it with GM, it could be mutually beneficial: for example, they could both reduce costs or they could make colors match better.

However, for small innovations or process improvements (which are usually not subject to patents) suppliers may be concerned that once they share the innovation with the buyer, the buyer will pass the technique along to other suppliers, looking for a lower price. Sharing the innovation would then make suppliers vulnerable to manufacturers taking away their competitive advantage. These acts of untrustworthiness are not uncommon among U.S. car manufacturers. *McMillan* (1990) already reported them several years ago. American automakers' procurement strategies have historically focused on achieving cost reductions rather than on building suppliers' trust (*Burt* 1989, *Liker and Choi* 2004)[4]. Industry reports suggest that this trend

[3]In some cases, car manufacturers explicitly express their expectations that suppliers will innovate on their behalf. *Burt* (1989) reports that Ford and other automakers explicitly sent their suppliers lists of technologies they would have liked to have developed. TRW's air bag and Bosch's antilock brakes (ABS) are examples of technologies that resulted from these sort of initiatives.

[4]We refer the reader to *Helper and Henderson* (2014) and *Liker and Choi* (2004) for a comprehensive review of the cultural differences between Japanese and American automakers between 1980 and 2009.

still continues. While in recent years American automakers have been improving in supplier working relations indexes, their ratings are still quite poor[5].

The prospect of a long term relationship can make both suppliers and manufacturers more likely to collaborate[6]. Relational contracts, defined by *Gibbons and Henderson* (2012) as "informal agreements enforced by the shadow of the future", can provide enough incentives for collaboration[7]. The manufacturer has incentives to keep the supplier's trust as long as his benefit from future innovation sharing exceeds the short-run gain from re-sharing the innovation. This, in turn, results in higher incentives for the supplier to share the innovations. However, even when the firms have the right incentives to collaborate, the individuals making decisions on behalf of the firms may not. Conversations conducted within GM and with GM's suppliers suggest that, which employees manage the relationship with the supplier impacts the nature of the relationship that develops. Some divisions within GM assign this function to procurement managers, while others leave supplier relations in the hands of the engineers[8]. Procurement managers are in essence short-run focused employees. Their incentives are driven by short-run cost reduction goals and they usually have higher rotation rates within the organization. On the other hand, engineers care

[5]Planning Perspectives, Inc. develops one of the most reputable indexes in the industry, the Working Relations Index (WRI). The WRI is based on interviews with American automotive suppliers, and measures aspects such as trust and overall working relationship, communication, supplier profit opportunities, help company gives to suppliers, etc. The 2015 report can be found online at http://www.ppi1.com/wp-content/uploads/2015/05/2015-WRI-Press-Release-May-19.pdf

[6]In the automotive industry, the length of supplier relations cover a large spectrum, from switching suppliers after each sale period, to strategic partnerships and preserving a supplier of a part for the entire length of a car model (typically, five to seven years), and even beyond the life-cycle of a model. *Dyer* (1996b) reports that Chrysler's efforts to increase their commitments to their suppliers increased the length of the contracts from an average of 2.1 years to 4.4 years. *McMillan* (1990) reports that contracts of three to five years are generally considered long-term in the automotive industry.

[7]For seminal work on the theory of relational contracts see *Gibbons* 1998, *Gibbons* 2001, *Gibbons* 2005, *Baker et al.* 2002. In our setting, *Helper and Henderson* (2014) describe the importance of relational contracts to understand the difference between Japanese and American automakers in terms of managing their supplier relations.

[8]Conversations within GM and with three of GM's top tier-one suppliers in the automotive industry were conducted by students of University of Michigan during the Spring-Summer semester of 2011. In other industries, such as electronics, there is also anecdotal evidence of engineers being involved in the development of a supplier base, particularly for new products (*Monczka* 2000).

about quality and design which are intrinsically long-run objectives. In addition, they have specific technical expertise on certain parts. This makes them more likely to be assigned to manage the relationship with the same supplier again the following year. Suppliers' trust on the buying firm is therefore influenced by who is the decision maker on the buyer side, beyond the length of the relationship between the firms. Consequently, buyers have to pay particular attention to the level of involvement they assign to the employees managing their supplier relations.

We conduct a laboratory experiment to study how firms' actions are affected by the buyers' allocation of decision rights to short-run and long-run focused employees. Specifically, our research questions are: (1) Does who controls the relationship affect a firm's strategy and the equilibrium outcomes? Does it matter if the buyer is a single or a dual decision maker? Does it matter if it is the procurement manager or the engineer that makes the decision for the buyer? (2) Does the duration of the contract matter? Our results show that the allocation of decision rights has a strong impact on the firms' actions and on supply chain efficiency. Having a procurement manager and an engineer making recommendations for what the buyer should do already leads to higher innovation sharing from the supplier, even when the procurement manager is the final decision maker. This results in higher supply chain efficiency relative to the case where the buyer is a single decision maker and has a short term supplier relation. On the hand, buyers only become significantly more collaborative when the engineer's recommendation is implemented. As a result, while buyers benefit from either allocation of decision rights among its employees (as long as both employees are involved in the decision making process), suppliers only get significantly higher profits when the engineer is the final decision maker. Finally, we observe that employees may be influenced by their peers' recommendations beyond their own monetary incentives.

3.2 Literature Review

There is a broad literature in operations management concerned with the issue of collaboration in buyer-supplier relations. Empirical papers show that efforts to cooperate with the other party can lead to improved performance and higher profits. For example, an empirical study of U.S. automotive suppliers by *Dyer and Hatch* (2006) found that greater knowledge sharing from the automakers' side resulted in a faster rate of learning within the suppliers and ultimately in a lower rate of defects. *Stallkamp* (2005) analyzes collaboration with business partners in terms of strategy, communication, information and responsibility sharing and reveals that remarkable cost and quality improvements can result from strategic collaboration. This collaboration between firms and their suppliers can be implemented in practice through different initiatives, which have been studied analytically for example in *Aviv* (2001), *Aviv* (2007), who focus on collaborative forecasting, and *Zhu et al.* (2007) who study how firms can coordinate quality-improvement efforts. Our paper focuses on collaboration in terms of innovation sharing with the final goal of reducing costs. Cost reduction is one of the main drivers of outsourcing decisions (*Gray et al.* 2009) and is an important part of supply chain relations (*Rudzki* 2004). There are several papers in the Operations Management literature that address the problem of incentives to invest in cost reduction initiatives in a supply chain. *Kim and Netessine* (2013) study the manufacturer's and supplier's interest in a collaborative effort to lower expected cost in the development phase of an innovative product. *Iyer et al.* (2005) focus on how buyers can allocate their resources to help suppliers transform specifications into finished components and reduce total costs. *Bernstein and Kök* (2009) study suppliers' incentives to invest in cost reduction solutions over the life cycle of the product under different procurement approaches, and consider gradual investment in process improvement (i.e. Lean Production, Six Sigma Programs). Our paper aims to address this topic from an experimental perspective, in order to understand how behavioral factors affect supply chain collaboration. *Brinkhoff et al.* (2015) provide empirical evidence that for supply chain projects to be successful trust is important, but it is mediated by project-level factors such as between-firm communication and within-firm commitment. This suggests that firms' organizational-level decisions may play a role in supply chain collaboration. Our experimental results show that in an innovation sharing setting, the allocation of decision rights to employees with different incentives is important in determining both firms' collaboration.

Arguments in favor of procuring from a reduced number of suppliers and preserving long-term supplier relations became popular during the late eighties and early nineties. Several studies reported a trend of shifting towards single sourcing (*Han et al.* 1993, *Newman* 1988), and assessed the benefits of this trend in terms of reducing costs and improving quality (*Kalwani and Narayandas* 1995, *Treleven* 1987). More recent papers in the OM literature have identified settings where longer relationships are beneficial for buyers. *Swinney and Netessine* (2009) model a non-cooperative supplier-buyer relationship in which the buyer is concerned with the failure of a supplier since switching suppliers in case of supplier default is costly. They find that, when they consider the possibility of default by the suppliers, buyers prefer long-term contracts and in particular, dynamic long-term contracts allow the buyer to coordinate the supply chain. *Taylor and Plambeck* (2007) analyze a setting where a firm is developing an innovative product and requires a supplier to invest in capacity for the product without being able to contract on it. They show that with long-term supplier relations, relational contracts provide enough incentive for the supplier to invest. Similarly, *Li and Debo* (2009) also find that committing to a longer relationship with a supplier can be more beneficial than running an auction in every period to select a supplier, since longer relationships incentivize suppliers to bid more aggressively. We provide further evidence in this direction: our experimental results show that longer relationships are also beneficial (for both, buyers and suppliers) in a setting with supplier innovation sharing.

There are a few recent experimental papers in operations management which study how buyer-supplier relationships are affected by the length of the relationship. *Özer et al.* (2011) find that trust and trustworthiness in forecast information sharing are enhanced with repeated interactions, resulting in lower forecast inflations, higher capacity and investment, and higher supply chain efficiency. *Beer et al.* (2014) show that when suppliers can signal trustworthiness by making an upfront buyer-specific investment, more collaborative relationships arise. With repeated interactions, the signaling effect of the investment is strengthened leading to higher profits and efficiency. To the extent of our knowledge the only experimental paper with a focus on the comparison between short run and long run incentives in a related setting, is *Hyndman et al.* (2014). They study a setup where two firms simultaneously invest in capacity to meet demand (think of a buyer and a supplier both making their respective investment decisions), and sales are the minimum of the two chosen capacities and realized demand. Firms have private information about demand, and need to coordinate on the optimal investment level. Their experiment compares behavior when subjects are in fixed pairs and when they are randomly re-matched after every round. They find that, while fixed pairs have higher alignment on average, they do not achieve higher efficiency. With fixed matching, the alignment reached in the initial rounds of play has a strong impact on the overall profits throughout the relationship. Therefore, pairs with higher alignment in the initial rounds ended up with higher profits than those who started misaligned. On the other hand, with random rematching, the initial rounds do not have much impact on overall average profits. *Hyndman and Honhon* (2014) find in a similar setting that when players are free to dissolve the relationship after every round, they earn higher average profits than when they are matched indefinitely. The object of study in our paper is quite different. First, because our stage game resembles more a trust game (innovation sharing) than a coordination game (capacity alignment). Second, because our focus is on the firms' allocation of decision rights, and therefore our setup is more of a hybrid of those in *Hyndman et al.* (2014): while procurement managers have random rematching after every round, engineers have fixed matching as long as the relationship between the firms lasts. With this setup, we capture the different incentives of the employees' working for the buyer, beyond the length of the relationship between then firms.

In order to study experimentally firms' actions in long-term supplier relations (which we model as infinitely repeated games), we implement an experimental design where subjects play an indefinitely repeated game. The literature in experimental economics has been using this methodology extensively. *Murnighan and Roth* (1983) and *Roth and Murnighan* (1978) were the firsts to induce infinitely repeated games using randomly terminated games, where the continuation probability is equated to the discount factor [9]. Since then, indefinitely repeated games have been used to understand the evolution of cooperation in a prisoners dilemma game (*Camera and Casari* 2009, *Aoyagi and Fréchette* 2009, *Bó* 2005, *Bó and Fréchette* 2011, *Fréchette and Yuksel* 2013, *Dal Bó and Fréchette* 2013, *Honhon and Hyndman* 2015), in a two-period Bertrand game (*Cooper and Kuhn* 2009), in a veto game (*Cabral et al.* 2014), and in a trust game (*Engle-Warnick and Slonim* 2004, *Engle-Warnick and Slonim* 2006a, *Engle-Warnick and Slonim* 2006b), among others (*Engle-Warnick* 2007). For the most part, the focus of these papers has been on inferring subjects' strategies from their actions in the game. This is not a trivial task since: (1) the set of possible strategies is infinite, and (2) strategy choices are not observable – the experimenter only observes the player's choice for the history that actually happened and not what the player would have done for any possible history (*Dal Bó and Fréchette* 2013).

[9]Recent experimental work by *Fréchette and Yuksel* (2013) finds evidence that games with random termination can be used to induce infinitely repeated games in the laboratory, as they generate behavior that is consistent with the theoretical predictions for these games.

Fortunately, there is evidence that relatively few basic strategies seem to explain players' actions quite well, and furthermore, these strategies are best responses to the opponent strategies. *Dal Bó and Fréchette* (2013) find that the most popular strategies in indefinitely repeated prisoners dilemma, are Always Defect, Tit-For-Tat, and Grim Trigger. Similarly, *Engle-Warnick and Slonim* (2006a) find that in the indefinitely repeated trust game, relatively few strategies explain vast majority of behavior. For the trustor both Grim Trigger and Tit-For-Tat are relevant strategies, while the trustee conditions behavior on round number rather than on the history of play with the opponent. Rather than directly recovering players' strategies, our focus is on the comparison of players' actions across treatments with different allocations of decision rights to the employees of the buying firm.

3.3 Model

We model a relationship between a manufacturer and his supplier. We first consider a benchmark case where the manufacturer and the supplier have a short-term relationship and model it as a single-period game. We then consider the case where the firms engage in a long-term relationship and model it as an infinitely repeated game with discounting, where the stage game is the single-period benchmark case. Finally, we analyze a case where the firms have a long-term relationship but the decision makers within the manufacturer are two employees, one short-run and one long-run focused.

3.3.1 Single period game

The single period game consists of a one-time transaction between a manufacturer and a supplier. The supplier produces a component that the manufacturer uses to produce a good. Let $C_i \geq 0$ be firm i's variable cost, $i \in \{m = \text{manufacturer}, s =$

supplier, a = alternative supplier} [10]. The supplier has a per unit production cost of C_{s1} and sells each unit of component to the manufacturer at a wholesale price w. The manufacturer has a per unit manufacturing cost of C_{m1} and a total per unit production cost of $C_{m1} + w$ and sells the product to the end customer at a retail price p. For simplicity, we model demand as a linear function of p, $Q(p) = a - bp$, where $a, b \geq 0$ and $a - bp > 0$ and assume the manufacturer can always meet demand. The manufacturer's profit from the transaction is $\Pi_m(p, w) = Q(p - w - C_m)$ and the supplier's profit is $\Pi_s(p, w) = Q(w - C_s)$.

At the beginning of the game, an innovation occurs with an exogenous probability π, which results in a reduction of the supplier's cost to C_{s2}, with $C_{s2} < C_{s1}$. Consider again the waterborne paint example. If the supplier discovered a way to set up the blowers so that the waterborne paint dries faster, this would allow the supplier to process more parts per shift. As a result, his cost would decrease, and surplus would increase. The supplier can also choose to share the innovation with the manufacturer. If he does, the manufacturer can also implement the new blowers in his own painting booths and thus, the manufacturer's production cost is reduced to C_{m2}, $C_{m2} \leq C_{m1}$.

After the supplier decides whether to share the innovation with the manufacturer, the manufacturer can choose to open up competition and bring in a new supplier (we call this decision "to compete") or to single source ("not to compete"). We assume that the alternative supplier has production cost C_{a1}, $C_{a1} > C_{s1}$. If the original supplier shared the cost reduction with the manufacturer and the manufacturer chooses "to compete", then the manufacturer shares the cost reduction with the new supplier, whose cost is reduced to C_{a2}, with $C_{a2} < C_{a1}$, and $C_{a2} = C_{s2}$. After the manufacturer chooses whether to compete or not, trade occurs. We assume that the supplier and the manufacturer simultaneously choose the wholesale price, w^*, and retail price, p^*

[10] As in *Bernstein and Kök* (2009), we assume complete information about cost structures: suppliers know the manufacturer's complementary assembly costs and the manufacturer knows the suppliers' production costs. This is a common assumption in the automotive industry, where suppliers share technical information with the manufacturer in the design phase.

that maximize surplus. We further assume the surplus is split between the supplier and the manufacturer according to Nash bargaining (*Nash Jr* 1950): the manufacturer earns a fraction α, $\alpha \in [0,1]$ of the surplus and the supplier earns $(1-\alpha)$ of the surplus[11]. The quantity sold Q^* is determined as a function of p^*, which determines the supplier's and manufacturer's profits. In the case where the manufacturer chooses not to compete (bilateral bargaining case), the Nash bargaining solution predicts equal splits of the surplus, that is $\alpha = \frac{1}{2}$, and the manufacturer's and supplier's profits are given by [12]:

$$\Pi_m = \Pi_s = \frac{(a - b(C_s + C_m))^2}{8b}. \tag{3.1}$$

In the case where the manufacturer chooses to compete, the Nash bargaining solution predicts $\alpha^* = \frac{1}{2} + \frac{(p - C_a - C_m)}{2(p - C_s - C_m)}$. Note first, that this requires $p > C_a + C_m$ and $p > C_s + C_m$. Second, that if the innovation does not happen, or if it happens and the supplier does not share, $C_a > C_s$ and the manufacturer's share of surplus is greater than $\frac{1}{2}$. In the particular case where the supplier shares and the manufacturer competes, we have $C_s = C_a$ and the manufacturer takes all the surplus. Thus, when the manufacturer chooses to compete, the original supplier is still the one who gets the deal, but now the manufacturer's profit is:

$$\Pi_m = \frac{[a - b(C_m + C_s)][a - b(C_m + C_a)]}{4b}, \tag{3.2}$$

and the supplier's profit is:

$$\Pi_s = \frac{[C_a - C_s][a - b(C_s + C_m)]}{4}. \tag{3.3}$$

[11]While there are several models of supply chain bargaining, we choose this approach for simplicity. For a more detailed study of bargaining in supply chains we refer the reader to *Lovejoy* (2010).

[12]We assume that the parameters are such that business makes sense, that is, $p - C_m \geq C_s$ and at $w^* = (1-\alpha)(p - C_m) + \alpha C_s$, both firms choose to transact.

The detailed calculations are presented in the Appendix.

3.3.2 Numerical Example

Figure 3.1 shows the game in extensive form. The game has six possible outcomes: If the innovations occurs, the possible outcomes are Share-Compete (ISC), Share-Do not Compete (ISN), Not Share-Compete (INC) and Not Share- Do not Compete (INN). If the innovation does not occur, the possible outcomes are Compete (NC) and Do not Compete (NN). In order to provide tractability for the experiment, we assign values to the parameters and generate payoffs for the supplier and the manufacturer for each possible outcome. The parameters used are presented in the Table C.1 in the Appendix and the firms' payoffs resulting from these parameters are presented in Figure 3.1.

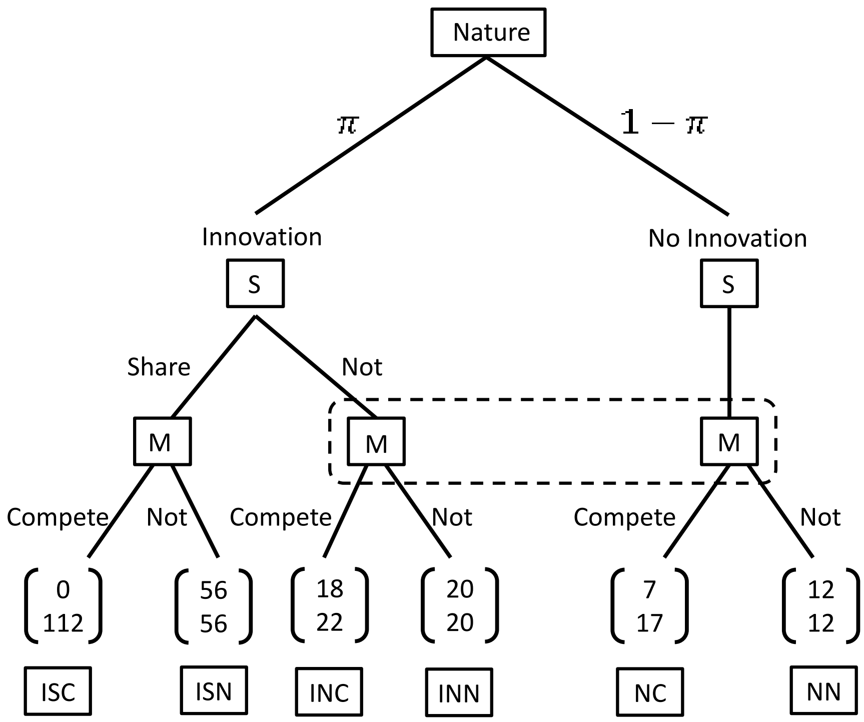

Figure 3.1: Stage Game

With these payoffs, if the innovation occurs, total surplus increases relative to the case where the innovation does not occur. In addition, if the supplier shares the innovation with the manufacturer, total surplus increases even further. However, sharing the innovation makes the supplier more vulnerable to competition: the minimum possible payoff from not sharing is 7 and from sharing it is zero. The manufacturer's decision does not affect the total surplus in size, but affects the allocation of this surplus between the two firms. Thus, we consider the supplier choosing to share and the manufacturer choosing not to compete as "collaborative" actions since both actions have a positive effect on their counterpart's payoffs. Note that at the moment of making a decision, the manufacturer cannot distinguish the case where the innovation occurred and the supplier chose not to share it, form the case where the innovation did not happen at all. This captures the information set of a manufacturer in the waterborne paint example. After making his decision, the manufacturer can infer from the payoffs which of the two scenarios happened. This exact same information structure is reproduced in the laboratory experiment.

The game on Figure 3.1 resembles the widely studied trust game (*Kreps* 1996) with two differences. First, in our setup the decision to trust is preceded by a random innovation. Second, the manufacturer makes his decision even if he was not trusted. In the original trust game, if the first decision was not to trust, the game ends and the second player is not called upon to play. These two differences are important to characterize our setting, and may affect our experimental results making them not directly comparable to those of the trust game. However, the main dynamics captured in our game are those of a trust game.

3.3.3 Firms' Decisions

We first analyze the most simple case where the supplier and the manufacturer have a short term relationship and model it as a single-shot game. Since firms interact only once, there are no incentives to play collaborative actions based on strategic concerns about future play. Thus, this case serves as a benchmark for the lowest theoretical level of collaboration. We then analyze the case where the firms have a long term relationship and model it as an indefinitely repeated game with discounting.

3.3.3.1 Single-Period Game

For the single-period game we solve by backward induction, starting with the manufacturer's strategy. The manufacturer's profit from choosing to compete is given by Equation 3.1 and from choosing not to compete is given by Equation 3.2. Since $p > C_a + C_m$, then $a - b(C_m + 2C_a - C_s) > 0$ and the manufacturer's profit is always higher if he chooses to compete than if he chooses not to compete. Given that the manufacturer chooses to compete, the supplier's profit is always given by Equation 3.3. Rolling back to the supplier's strategy, if he choose to share, then $C_{a2} = C_{s2}$ and the supplier earns zero profit, while if he does not share, $C_{a1} > C_{s2}$ and the supplier earns positive profits. As a result, the supplier does not share and the only Nash equilibrium in a one-period play of the game in Figure 3.1 are INC if the innovation occurs and NC if the innovation does not occur.

3.3.3.2 Repeated Interactions

We now consider the infinitely repeated play of the stage game depicted in Figure 3.1. We assume firms discount their payoffs across periods with a discount factor δ per period, $\delta \in [0, 1]$. That is, a dollar to be received next period is worth today δ and a dollar to be received n periods from today is worth today δ^n. This implies that the smaller δ is, the more impatient the player is. Another interpretation of the discount factor δ is the continuation probability of the indefinitely repeated game. We resort to this interpretation later on for experimental design purposes, as is common in the experimental economics literature.

Consider the six different possible outcomes of the stage game presented in Figure 3.1 and let $ISC_i, ISN_i, INC_i, INN_i, NC_i, NN_i$ be player i's payoffs, $i \in \{m =$ manufacturer, $s =$ supplier$\}$, from each possible outcome. The next proposition characterizes the conditions for the collaborative actions (the supplier shares and the manufacturer does not compete) to be part of a Nash Equilibrium of the infinitely repeated game. In preparation, we define a threshold $\hat{\delta}_1 = \frac{ISC_m - ISN_m}{ISC_m - (\pi INC_m + (1-\pi) ISN_m)}$.

Proposition III.1. *If $\delta \geq \hat{\delta}_1$, the following pair of grim trigger strategies is a subgame perfect Nash equilibrium of the infinitely repeated game. For the supplier, in the first period where there is an innovation play Share. Thereafter, if all moves in previous periods where there was an innovation have been Share and Do not Compete, play Share when there is an innovation. Otherwise, play Not Share when there is an innovation. For the manufacturer, play Do not Compete when there is an innovation if all moves in all previous periods where there was an innovation have been Share and Do not Compete. Otherwise, play Compete. In all periods where there is no innovation, play Compete.*

The proof follows the same logic as *Gibbons* (2001). Grim trigger strategies dictate that both players play the collaborative action and, in case of collaboration breakdown, they never collaborate again. Thus, both players' grim trigger strategies are a Nash Equilibrium of the game as long as the present value from collaboration outweighs the gains from a one-time deviation from collaboration followed by a perpetuity of defection by both players. For the supplier this happens if $ISN_s \geq INC_s$, that is if the profit in Equation 3.1 with C_{s2} and C_{m2} is greater than the profit in Equation 3.3 with C_{s2}, C_{m1}, and C_{a1}. This always holds since $p > C_{a1} + C_{m1}$ and $C_{m2} < C_{m1}$. For the manufacturer the condition requires $\delta \geq \hat{\delta}_1$, which in our numerical example means $\delta \geq \frac{112-56}{112-(0.75(22)+(1-0.75)56)} = 0.69$. Detailed calculations are presented in the Appendix. Based on the two interpretations of δ, this means that the manufacturer needs to care enough about his future payoff (be patient enough) or that the relationship needs to be likely enough to continue after each round of play.

The Folk Theorems for infinitely repeated games show that many strategies other than trigger strategies can support equilibria with collaborative outcomes [13]. We focus on trigger strategies since they provide the highest disincentive to deviate from collaboration. Thus, the conditions above provide the largest set of parameters under which collaboration can be sustained in equilibrium. In addition, trigger strategies are the less risky of the collaborative strategies for suppliers when matched with manufacturers playing always compete, which is a very common strategy based on previous experimental evidence.

3.3.4 Employees' Decisions

We focus now on the setting where the firms have a long term relationship. We assume that the manufacturer employs a short-run focused procurement manager and a long-run focused engineer. The procurement manager works for the firm for only one period, whereas the engineer works for the firm to infinity. We further assume that both employees make recommendations for what the manufacturer should do and that their compensation is the manufacturer's profit.

Consider first the procurement manager's recommendation. Since the procurement manager works for the buyer for only one period, the game between the supplier and the procurement manager resembles that of two firms playing a single period game. Thus, in a setting where the procurement manager's recommendation is always implemented, the procurement manager recommends to always compete and the supplier always chooses not to share. The only Nash equilibrium in this case are Not share-Compete (INC) when the innovation occurs, and Compete (NC) when the innovation does not occur.

Consider now the engineer's recommendation. Since the engineer works for the buyer to infinity and the firms have a long term relationship, the game between the

[13] *Fudenberg and Maskin* (1986), *Rubinstein* (1979). For an application of the Folk Theorem to problems similar to ours, refer to *Miller* (2001); *Miller and Smith* (1993).

supplier and the engineer resembles an infinitely repeated game. Thus, when the engineer's recommendation is always implemented, Proposition III.1 applies: trigger strategies can sustain a repetition of the collaborative outcome Share-Do not Compete (ISN) in every period where there is an innovation. When the innovation does not occur, the engineer will choose compete (NC)[14].

Finally, consider the case where if both employees' recommendations agree, their recommendation is implemented and if they disagree, one of the two recommendations is implemented at random, both with equal probability. We will assume that the supplier can perfectly observe both employees' recommendations [15]. In this case, trigger strategies analogous to those in Proposition 1 can sustain the collaborative outcome, Share - Not compete. The result is presented in the next proposition, for which we define $\hat{\delta}_2 = \frac{ISC_m - ISN_m}{(1+\frac{\pi}{2})ISC_m - \pi INC_m - (1-\frac{\pi}{2})ISN_m}$.

Proposition III.2. *If $\delta \geq \hat{\delta}_2$, the following set of trigger strategies is a subgame perfect Nash equilibrium of the infinitely repeated game. The supplier chooses to share in the first period where there is an innovation. Thereafter, if all moves in previous periods where there was an innovation have been Share and the engineer recommended Do not Compete, play Share when there is an innovation. Otherwise, play Not Share when there is an innovation. The engineer chooses Do not Compete when there is an innovation if all moves in all previous periods where there was an innovation have been Share and the engineer's recommendation was Do not Compete. Otherwise, he chooses Compete. In all periods where there is no innovation, he chooses Compete.*

[14]We focus only on pure strategies that lead to an equilibrium with high sharing rates. In mixed strategies, the engineer could induce the supplier to share by using, for example, a strategy where he does not compete only with some probability when the supplier shares. This would result in more sophisticated review strategies as the supplier needs to gather probabilistic evidence of the engineer's actions across several periods.

[15]We make the assumption that suppliers can observe both employees' recommendations for simplicity. It captures, for example, a setup where the supplier is present during the buyer's internal deliberation. If the supplier cannot observe both recommendations, a collaborative equilibrium can be reached if the supplier resorts to review strategies (*Radner* 1985) by which he can assess probabilistically if the engineer is recommending to not-compete, after observing several rounds of play.

The procurement manager chooses to Compete in every round.

The proof is analogous to that of Proposition III.1 and is relegated to the Appendix. The supplier's incentive compatibility requires $\frac{ISN_s+ISC_s}{2} \geq INC_s$, and the engineer's requires $\delta \geq \hat{\delta}_2$. Intuitively, the supplier's condition is tighter than in Proposition III.1 because, if the supplier shares in every period, half the times the procurement manager's recommendation will be chosen and the buyer will compete. Thus, the supplier's expected profit in equilibrium is lower than when the buyer chooses not to compete in every round. On the other hand, the engineer's expected payoff in equilibrium is now higher. This is because as long as the engineer always chooses not to compete when the innovation occurred, the supplier will always choose to share. Thus, the engineer enjoys the monetary benefits of the procurement manager's recommendation to compete without facing the supplier's punishment. In our numerical example, the supplier's incentive compatibility condition holds since $ISN_s = 56$, $ISC_s = 0$, $INC_s = 18$, and $\frac{56+0}{2} > 18$. The engineer's incentive compatibility requires $\delta \geq 0.55$.

3.4 Experimental Design

The experiment consists of two sets of treatments, the first one is Firms-as-a-Monolith (two treatments) and the second one is Firms-as-Employees (three treatments).

In the Firms-as-a-Monolith treatments, subjects are assigned a role representing a firm, which they keep throughout the experiment: supplier or buyer. To induce short term relationships between the firms (single-shot games), we use single-round games with random re-matching after every game. We call this treatment "Short Run" (SR). To induce long term relationships between the firms (indefinitely repeated games), we use a random continuation rule: after each round, the relationship continues with probability δ and ends with probability $(1-\delta)$. To implement this, the computer randomly generates a number between 0 and 100 (which is shown to the participants) and continues to a new round in the same relationship if the number is lower or equal than 100δ or ends the relationship otherwise. We define a *relationship* as a sequence of consecutive rounds where the buyer and the supplier remain matched until the random end occurs. A *period* counts each play of the stage game in a session, while *round* counts the periods within a relationship (it is re-set for each subject every time a new relationship begins). After a relationship ends, subjects are randomly re-matched as long as the time has not yet reached the session's pre-set time limit. Subjects know in advance the continuation and rematching rules. We denote this treatment "Long Run" (LR).

In the Firms-as-Employees treatments, there exist hypothetical firms (a buyer and a supplier) however now, subjects' roles do not represent the firms but particular employees working for these firms. Specifically, subjects are assigned one of the following three roles: supplier, procurement manager working for the buyer (from now on *procurement manager*), or engineer working for the buyer (from now on *engineer*). The "supplier" and the "buyer" have a long term relationship but the participants in the experiment play the roles of employees which have different life spans within the firms. While the subjects playing in the roles of supplier and engineer remain working for the same firm (supplier and buyer respectively) as long as the relationship between the two firms lasts, the procurement manager works for a different buyer in every round. We implement this as follows: When a new relationship begins, groups of one supplier, one procurement manager, and one engineer are formed. After each round, a random number is drawn to determine if the relationship between the firms continues. If the relationship continues, the supplier and engineer remain matched for the following round and the procurement manager is randomly and anonymously re-matched with a new supplier-engineer pair. If the relationship between the firms ends, all players are re-matched into new groups. Suppliers keep their role throughout the experiment, while procurement managers and engineers are randomly re-assigned a role at the beginning of each new relationship.

In the Firms-as-a-Monolith treatments, the sequence of events and payoffs in each round follow the stage game presented in Figure 3.1. In order to elicit complete strategies from the participants, we use the strategy method in which participants make conditional decisions for each possible scenario that may arise. As depicted on Figure 3.1, in the first stage the innovation occurs with probability π. The computer randomly determines whether the innovation occurs but does not inform the participants the outcome. Second, suppliers are asked whether, in case the innovation has occurred, they want to share it with the buyer. In the third stage, buyers are asked whether they would want to compete or not in case the supplier shared the innovation and in case he did not (at this point the buyer does not know in case of not sharing, whether the innovation occurred and the supplier did not share it, or if the innovation did not happen at all). After suppliers and buyers have made their decisions, all the subjects in the group are informed whether the innovation occurred and if so, the supplier's decision, and the buyer's decision for the realized case of innovation and supplier's decision. They are also presented with the resulting payoffs. In the SR treatment, subjects are re-matched after each round (and they know they are being re-matched). In the LR treatment, subjects are shown a randomly generated number and if it is lower than 100δ, the relationship continues and subjects keep their partners and if a number higher than 100δ shows up, they are randomly re-matched for a new relationship.

In the Firms-as-Employees treatments, the first and second stages are as in the Firms-as-a-Monolith treatments. In the third stage, both the engineer and the procurement manager make recommendations for what the buyer should do. Engineers and procurement managers answer whether the buyer should compete if the supplier shared and if the supplier did not share. Since the engineer has been matched with the same supplier starting from the first round of the relationship between the firms, he knows all the previous history of play within the relationship. The procurement manager on the other hand, joins a new relationship in every round and does not know the history of play in the relationship he is joining in. To allow for strategies that are contingent on previous play, before procurement managers make their recommendations, they are informed of the last round history in the relationship they have joined. All subjects know that this information is provided to procurement managers. Three treatments allow for different allocation of decision rights between the engineer and the procurement manager: the procurement manager's decision treatment (denoted PM treatment) where the procurement manager's recommendation is always implemented, the engineer's decision treatment (denoted Eng treatment) where the engineer's recommendation is always implemented, and the joint decision treatment (denoted $50-50$ treatment) where if both employees' recommendations coincide, the firm implements their recommendation and if they disagree, the computer randomly picks one recommendation (both with equal probability). These implementation policies are public information in all treatment. After all players made their choices, all subjects in the group learn whether the innovation happened and if so, the supplier's decision, the engineer's and procurement manager's recommendations for the scenario that happened, and which recommendation was implemented. In the three treatments, all group members know at the end of the round both employees' recommendations and which one was implemented. The payoffs for the round are presented to all players and a new number is drawn to determine if the relationship between the firms continues for another round. Subjects playing as suppliers get the payoff the supplier firm and subjects playing as procurement managers and engineers each get the payoff of the buying firm. Note that in the PM treatment, subjects playing as engineers spend a whole relationship making recommendations which are never implemented (and the same happens with procurement managers in the Eng treatment). However, since after each relationship engineers and procurement managers are randomly re-assigned a new role, most subjects get to play the role whose recommendation is implemented in the corresponding treatment at some point during the session.

In total we have five treatments, SR, LR, PM, Eng, and $50 - 50$ and follow a between-subjects design (each subject is exposed to one treatment). To ensure the subjects' understanding of the game, there are three examples included in the instructions, and the table with payoffs (Figure C.1 in the appendix) is shown to participants throughout the experiment. In particular, to avoid biases relative to the continuation probabilities, in the LR, PM, $50-50$, and Eng it was made explicit that after each round, the probability that the relationship would continue for at least another round was exactly the same. To avoid reputation effects, participants only learn the outcomes and payoffs of their own relationships. In addition, there is a minimum of four relationships playing simultaneously in any given session, so that it is unlikely that subjects can track their partners after random re-matching. The parameters used in the experiment match those in the numerical example (Section 3.2): the probability of innovation is $\pi = 0.75$[16] and the continuation probability is $\delta = 0.75$ [17].

3.4.1 Procedures

The experiments were conducted in z-Tree (*Fischbacher* 2007) between March and September of 2014 at the behavioral laboratory of the School of Information at

[16]Setting $\pi = 0.75$ allowed us to get a high frequency of the interesting outcome where the innovation happens. It captures the occurrence of small innovations or process improvements, rather than big events such as disruptive new technologies (which in reality happen less frequently).

[17]Propositions III.1 and III.2 show that with the payoffs in Figure 3.1, cooperation can be supported as part of a subgame perfect equilibrium for values of continuation probability greater than 0.69 and 0.55 respectively. If we consider the automotive industry, we can assume that firms make these decisions on an annual basis. A 0.75 continuation probability implies average relationship lengths of four years, which is consistent with the industry (*McMillan* 1990).

University of Michigan. A total of 372 undergraduates participated in four sessions of each of the Firms-as-a-Monolith treatments and six sessions of each of the Firms-as-Employees treatments. The maximum number of subjects per session was 18 and the minimum was 10 for the Firms-as-a-Monolith treatments and 12 for the Firms-as-Employees treatments. Each session lasted approximately one hour, the SR treatment ended after 40 rounds, all other treatments ended after 50 minutes (including the time for reading the instructions) to allow some time for payment [18]. The average number of rounds per relationship was 3.9, with a minimum of 1 and a maximum of 11. Average payoffs were $11, consisting of a $5 show up fee plus the payoffs of two randomly selected rounds at a conversion rate of $0.10 per point earned [19].

3.5 Hypotheses

We derive the following experimental hypotheses from the theoretical results. The first hypothesis is derived from the equilibrium outcomes of the games between firms. In the one-shot game the only equilibrium is non-collaborative while in the infinitely repeated game, collaboration can be supported in equilibrium. Thus, we expect collaboration to be lower when the firms have a short term relationship than when they have a long term relationship.

HYPOTHESIS 1. [Firms-as-a-Monolith Treatments] *In the SR treatment there is less collaboration between the firms than in the LR treatment. Specifically,*

1.a - in the SR treatment the supplier chooses to share less frequently than in the LR treatment,

1.b - in the SR treatment the buyer chooses to compete (if shared) more often than in

[18] We dropped all observations after period 30, which is the maximum number of periods reached in all sessions. The results do not change significantly if we use all observations.

[19] Some previous experimental papers chose to pay for performance on randomly chosen full relationships rather than rounds. Comparing both, *Sherstyuk et al.* (2013) find that per-round payment slightly biases subjects towards short-term focus (present-period bias). In our setup this effect would only bias against finding treatment differences. In addition, the effect seems to be more prominent in the first round of a relationship, while our results show bigger differences in later rounds.

the LR treatment, and

1.c - the frequency of collaborative outcomes (both firms collaborate simultaneously) is lower in the SR treatment than in the LR treatment.

The next hypothesis is for the Firms-as-Employees treatments. Since the procurement manager works for the buyer for only one period, his relationship with the supplier resembles a one-shot game. Thus the play in the PM treatment should map onto the SR treatment. On the other hand, the engineer remains working for the same buyer as long as the relationship with the supplier lasts. Thus, the Eng treatment should map onto the LR treatment. Finally, for the 50-50 treatment, the theory prescribes an equilibrium where the supplier always shares and the buyer implements half the times the engineer's recommendation and half the times the procurement manager's recommendation. The engineer recommends to not compete and the procurement manager recommends to compete. Thus, the frequency of rounds with collaborative outcomes in the 50-50 treatment should be higher than in the SR treatment but lower than in the LR treatment.

HYPOTHESIS 2. [Firms-as-Employees Treatments] *In the Firms as Employees treatments, collaboration is in between the SR and LR benchmarks:*

2.a - the PM treatment maps onto SR treatment,

2.b - the Eng treatment maps onto LR treatment, and

2.c - the 50-50 treatment is in between the SR and LR treatments: the supplier shares as in the LR and the buyer competes more than in the LR treatment and less than in the SR treatment.

The two Firms-as-a-Monolith treatments serve as benchmarks for the lowest and highest collaboration between the firms. In the Firms-as-Employees treatments, the buyer allocates decision rights to the procurement manager and the engineer and therefore, we expect collaboration to be in between the two benchmarks. Based on Hypotheses 1 and 2, if we order the treatments SR - PM - 50-50 - Eng - LR, we should see a gradient of increased collaboration from SR to LR.

HYPOTHESIS 3. [Trends across treatments] *There is a trend of increasing collaboration from SR to LR:*

3.a - the frequency of sharing increases,

3.b - the frequency of compete (if shared) decreases, and

3.c - the frequency of collaborative outcomes increases.

The procurement manager joins a relationship for only one round and then rotates to another firm. Thus, the procurement manager should always choose to compete, regardless of the engineer's previous recommendation. Similarly, the engineer should not condition his recommendation on the recommendation of the previous procurement manager. Trigger strategies prescribe that the engineer's strategy is only contingent on the supplier's and his own previous history of play.

HYPOTHESIS 4. [Interplay between employees] *The engineer's recommendation is independent of the procurement manager's recommendation in the previous round. The procurement manager's recommendation is independent of the engineer's recommendation in the previous round.*

3.6 Experimental Results

In the first two sections, we compare the supplier's and the buyer's actions across the five treatments and analyze the outcomes and resulting profits. In the third section, we analyze in depth each of the between-employees treatments and analyze the interplay between engineers and procurement managers.

3.6.1 Descriptive Results

We look first at the supplier's decision to share or not to share the innovation. Hypothesis 1 predicts that suppliers will choose to share more often in the LR treatment than in the SR treatment. Hypothesis 2 predicts that suppliers will share in the PM treatment with the same frequency as in the SR treatment and in the $50-50$ and Eng treatments, as in the LR treatment. Finally, Hypothesis 3 predicts an increasing trend of sharing from SR to LR when treatments are ordered SR, PM, $50-50$, Eng, and LR. The average results presented in Table 3.1 show that suppliers' decision to share becomes more frequent as we go from the SR treatment to the LR treatment. A non-parametric test for trends shows that sharing increases from SR to LR ($p < 0.001$) [20]. However, pair-wise comparisons across treatments show that the only significant difference is between all treatments and the SR treatment (Wilcoxon rank-sum test $p = 0.041$ for difference between SR and PM treatments). Average sharing is not significantly different across the PM, $50-50$, Eng, and LR treatments. We observe similar results in the regression presented on Table 3.2. The regression of the decision to share on treatment dummies, controlling for round within a relationship, period of play in the session, and subject fixed effects, shows that all treatments (including the PM treatment) have higher frequency of sharing than the SR treatment baseline. The coefficients for all treatment dummies are not significantly different.

The buyer's decision to compete or not to compete, also presents a significant trend of increased collaboration from SR to LR (test for trends: $p < 0.001$), which supports Hypothesis $3.b$. Table 3.1 shows that while in the SR treatment buyers compete 77.6% of the times, in the PM treatment they compete 71% of the times ($p = 0.009$). As in the case of the suppliers' sharing decisions, this suggests that the PM treatment does not fully map into the SR case, and instead it already shows

[20] The non-parametric test for trends across ordered groups is an extension of the Wilcoxon rank-sum test. We consider subject level data (each subject's average share decision across all rounds is considered as one observation for the test).

some increased collaboration relative to the SR benchmark. Nonetheless, the largest difference relative to the SR benchmark in the Eng and LR treatments (as predicted by Hypotheses 1.*b* and 2.*b*). Table 3.1 shows that the buyers' decisions to compete drop to 62.7% and 58.5% in the Eng and LR treatments respectively. Similarly, Table 3.2 shows that the buyers compete less frequently than the SR benchmark in the PM treatment (marginal effects: -0.087) and even less frequently in the Eng and LR treatments (marginal effects LR vs. SR: -0.212). So far, we considered the buyer's decision that was actually implemented. This includes cases where the innovation does not happen and where the supplier does not share. Since we used the strategy method, we collected data on whether the buyer chooses to compete *if the supplier shared* in every round. When looking at the buyer's compete-if-shared decisions, Column 2 on Table 3.2 shows that the only significant decrease relative to the SR benchmark occur in the Eng and LR cases (marginal effects LR vs. SR: -0.112), the PM treatment is not significantly different. While the trend of decreasing compete-if-shared from SR to LR is as predicted in Hypothesis 3.*b*, we observe that the $50-50$ treatment deviates significantly from the predictions and shows higher compete rates than expected. Further analysis of the $50-50$ treatment is presented in the following sections.

The results so far indicate that suppliers and buyers react differently to the joint decision making cases (Firms-as-Employees treatments). Hypothesis 2 predicts that in the PM treatment the suppliers' frequency of sharing is as low as in the SR benchmarck and that in the $50-50$ and Eng treatments it is as high as in the LR treatment. We observe that suppliers' sharing actually increases even *sooner* than predicted by hypothesis 2. Even in the PM treatment suppliers share significantly more frequently than in the SR benchmark. On the other hand, the frequency of buyers' compete decisions decreases *later* than predicted by the hypotheses. Hypothesis 2 predicts that in the $50-50$ treatment competition should already be lower than in the SR benchmark but still higher than in the LR treatment and that the Eng treatment should map onto the LR treatment. We find that in fact in the $50-50$ treatment the buyers choose to compete even more frequently than in the PM treatment ($p = 0.094$) and it is only in the Eng treatment that buyers star competing less frequently than in the SR benchmark.

Lastly, we analyze the frequency with which collaborative outcomes occur in each treatment. Recall that we defined a collaborative outcome as a play of the stage game where the supplier chooses to share and the buyer chooses not to compete. A test for trends shows that the frequency of collaborative outcomes increases from SR to LR ($p < 0.001$), as predicted by Hypothesis 3.c. Table 3.1 shows that collaborative outcomes occur 5.3% of the times in the SR treatment, 9.5% of the times in the PM treatment, and 10.8% of the times in the $50-50$ treatment (all these are not significantly different). In the Eng treatment, it is significantly higher than in the $50-50$ treatment ($24.6, p = 0.007$) and not different from the LR treatment (21.3). The regression presented on Table 3.2 shows that all treatments are associated with higher frequency of collaborative outcomes than the SR baseline (PM vs. SR marginal effects: 0.072). In particular, the Eng and LR treatments are correlated with an even higher frequency of collaborative outcomes (LR vs. SR marginal effects: 0.183). In addition, the coefficient for the Eng treatment is significantly higher than for the $50-50$ treatment ($p = 0.005$). These results seem, for the most part, to confirm the hypotheses. They depart from the hypotheses in two ways: collaboration in the PM treatment is (marginally) higher than in the SR treatment, and collaboration in the $50-50$ treatment is not higher than in the PM treatment. We explore these results in the following sections by analyzing each of the Firms-as-employees treatment in more detail.

Table 3.1: General Results - Frequency of Collaborative Outcomes

Treatment	Supplier's Decision (Share) (%)	Buyer's Decision (Compete) (%)	Collaborative Outcome (Share/Not Compete) (%)
SR	18.2	77.6	5.3
PM	29.2	71	9.5
$50-50$	33.5	74.6	10.8
Eng	40.2	62.7	24.6
LR	38.5	58.5	21.3

Table 3.2: General Results

Coefficients	Supplier's Decision (Share)	Buyer's Decision (Compete if Shared)	Buyer's Decision (Compete)	Collaborative Outcomes (Share/Not Compete)
PM	0.624***	-0.192	-0.247**	0.389*
	(0.206)	(0.150)	(0.126)	(0.233)
$(50-50)$	0.778***	0.979***	-0.102	0.462**
	(0.201)	(0.156)	(0.124)	(0.227)
Eng	0.964***	-0.571***	-0.465***	1.055***
	(0.207)	(0.152)	(0.127)	(0.230)
LR	0.897***	-0.315**	-0.574***	0.849***
	(0.201)	(0.147)	(0.123)	(0.225)
Period	-0.008***	-0.018***	-0.017***	0.013***
	(0.003)	(0.003)	(0.003)	(0.003)
Round	-0.064***	0.041***	0.013	-0.009
	(0.012)	(0.012)	(0.011)	(0.014)
Constant	-0.938***	0.734***	1.045***	-2.085***
	(0.154)	(0.114)	(0.098)	(0.181)
Observations	4286	4286	4286	4286
Nr. of Subjects	143	143	143	143

Probit regression with subject random effects. Standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Note:

3.6.2 Profits

Table 3.3 presents average profits for suppliers, buyers, and the total profits of both players combined. A test for trends shows that all three present an increasing trend from SR to LR ($p = 0.001, 0.009$, and < 0.001 for suppliers', buyers', and total profits respectively). The results also show that, while the PM treatment leads to profits that are slightly higher than the expected profits of no collaboration (16.14 vs. 15.25, one sided t-test p = 0.077), suppliers only earn significantly higher profits than in the SR benchmark in the Eng and LR treatments. Buyers' profits on the hand, are higher than in the SR benchmark in all the other treatments. This results are consistent with the previous findings; suppliers become more trusting in the Firms-as-Employees treatments, even if the procurement managers' recommendation is the one that is always implemented (PM treatment). On the other hand, the buyers' decision to compete is only significantly lower than in the SR benchmark when the engineer's recommendation is the one that is always implemented. These two effects combined explain why buyers benefit from all the treatments where both employees make recommendations.

Table 3.4 confirms the previous results with a regression of suppliers', buyers', and total profits on treatment dummies controlling for period, round and subject fixed effects. Suppliers' profits only increase relative to the SR benchmark in the Eng and LR treatments, while buyers' profits increase in all the treatments where there are engineers' and procurement managers' recommendations, as well as in the LR treatment. Total profits are higher in the PM and $50-50$ treatments than in the SR benchmark, and even higher in the Eng and LR treatments. Recall that total surplus increases if the innovation occurs and, even further, if the supplier shares the innovation. The buyer's decision affects only the allocation of total surplus between the supplier and the buyer. Since the innovation occurs with the same probability in all treatments, the results in total profits shows the pattern of increased sharing from the suppliers from SR to LR.

Table 3.3: General Results - Profits

Treatment	Supplier	Buyer	Total	Supplier's Fraction of Total Surplus (%)
SR	15.48	30.68	46.16	34
PM	16.14	35.23	51.37	31
$50-50$	15.63	36.93	52.55	30
Eng	21.83	37.60	59.43	38
LR	19.93	36.97	56.90	35
Non-Collaboration Expected Profit	15.25	20.75	36	42.4

Table 3.4: Profits

Coefficients	Supplier's Profits	Buyer's Profits	Total Profits
PM Treatment	0.137	7.455***	7.613***
	(0.773)	(2.675)	(2.794)
$50-50$ Treatment	-0.377	9.132***	8.775***
	(0.725)	(2.596)	(2.758)
Eng Treatment	5.843***	9.801***	15.661***
	(1.506)	(2.834)	(3.653)
LR Treatment	3.962***	8.993***	12.975***
	(1.332)	(2.674)	(3.291)
Period	0.164***	-0.281***	-0.117
	(0.035)	(0.064)	(0.076)
Round	0.224	-1.242***	-1.027***
	(0.142)	(0.265)	(0.316)
Constant	12.720***	36.278***	49***
	(0.596)	(2.138)	(2.250)
Observations	4286	4286	4286
Nr. of Subjects	143	143	143

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

3.6.3 Firms-as-Employees treatments

The previous results show that the trends of increased collaboration are present for the supplier's decision to share, the buyer's decision to compete, and the frequency of collaborative outcomes. We next describe in depth the experimental results for the Firms-as-Employees treatments. We find that: (1) the PM treatment does not exactly map onto the SR treatment, (2) the $50-50$ treatment is not exactly "in between" the PM and Eng treatments as predicted by Hypothesis 2, and (3) the Eng treatment presents no significant differences with the LR treatment. Finally, our results show that there exists interplay between the employees beyond what the theory predicts.

While the SR and PM treatments should be identical, the results on Table 3.5 shows important differences. The first and second columns in Table 3.5 show the fraction of times buyers chose to compete when the supplier shared and when the supplier did not share respectively. The next two columns show the supplier's expected profit from sharing and from not sharing given how the buyers responded to these two actions in the experiment. The difference between the two is negative in the SR treatment and positive in the PM treatment. This implies that, in expectation, sharing is profitable in the PM treatment and not in the SR treatment. Table C.2 in the Appendix confirms this result. A regression of the average profit per round within a relationship on the average sharing in that relationship in the PM treatment shows a positive correlation between the two ($\beta = 5.043$, $p = 0.01$). This means that, for example, for a supplier who shared 10% of the times, a increase in collaboration to sharing 60% of the times would be associated with an expected increase in profits per round of 2.52 points. Since the average profits per round for suppliers in the PM treatment is 16.59 points, this would imply a 15% increase in profits. In the experiment, suppliers seem to acknowledge this difference: they share 18.2% of the times in the SR treatment and 29.2% of the times in the PM treatment.

In the $50-50$ treatment, the hypotheses predicted that engineers would recommend not to compete and procurement managers would recommend to compete. With a continuation probability of 0.75, this allows for an equilibrium where the supplier always shares. Since one recommendation is chosen at random, collaboration should be higher than in the PM treatment and lower than in the Eng treatment. The results show that in the $50-50$ treatment, both the engineer and the procurement managers

Table 3.5: Supplier's Decision

Treatment	Compete if shared* (%)	Compete if not shared* (%)	E[Profit from sharing]	E[Profit from not sharing]	Diff.	Share* (%)
SR	67.6	78.5	15.63	15.84	−0.22	18.2
PM	64.3	73.2	17.08	15.99	1.09	29.2
50 − 50	71.2	76.4	14.14	15.90	−1.76	33.4
Eng	50.7	68.6	22.85	16.11	6.74	40.3
LR	59.0	56.6	19.51	16.44	3.07	38.5

Note: the columns marked with (*) present data from the experiment. The others present the suppliers' expected profits given the buyer's choices in the experiment.

compete more than in the PM and Eng treatments. Figure C.2 in the Appendix shows that the engineers' decision to compete when the supplier shared is higher in the 50 − 50 treatment than in the PM treatment ($p = 0.013$) and the Eng treatment ($p = 0.033$). Similarly, the procurement managers' decision is higher in the 50 − 50 treatment than in the PM treatment ($p = 0.055$) and the Eng treatment ($p = 0.071$). In addition, a high proportion of the collaborative outcomes in the 50 − 50 treatment is generated by procurement managers (engineers' decisions account for 63% of the collaborative outcomes and procurement managers' for 37%).

We also study how the dynamics of play in the 50 − 50 treatment compare to those in the PM and Eng treatments. Table 3.6 shows the probability that any round will result in a collaborative outcome for each treatment, partitioned into the following cases: collaboration that happens in the first round of a relationship, collaboration that happens in any round after the first one of a relationship when there was a collaboration in the immediate previous round, and collaboration that happens in any round after the first one of a relationship when there was not a collaborative outcome in the round immediate before. First, we note that in all treatments the probability of having a collaborative outcome when there was no collaboration in the period immediate before is very low (approximately, 0.06) and does not vary significantly by treatment. Second, the probability of having a collaborative in the first round of a relationship is higher in the Eng and LR treatments (0.23 and 0.21 respectively) relative to the $50 - 50$ and PM treatments (0.12 and 0.14 respectively). However, the largest difference across treatments resides in the probability of a collaborative outcome in rounds 2 onwards when there was a collaborative outcome in the round immediate before (0.31 and 0.43 in the PM and $50 - 50$ treatments vs. 0.77 and 0.79 in the Eng and LR treatments). This suggests that the $50 - 50$ treatment is more similar to the PM treatment than to the Eng treatment in terms of sustaining collaboration, as can be seen graphically in Figure 3.2. While in the Eng and LR treatments, a relationship that starts with a collaborative outcome is likely to result in collaborative outcomes in the following periods, in the PM and $50 - 50$ treatments this is less likely to occur. Furthermore, in all treatments, once the collaborative outcome is not reached, it is very unlikely that it will be reached again in a subsequent round.

Table 3.6: Collaboration throughout Relationships

		SR	PM	50 − 50	Eng	LR
Pr (CO)*		0.053	0.095	0.108	0.246	0.213
(A)	Pr (CO $\mid R = 1$)	0.053	0.14	0.12	0.23	0.21
			(37%)	(30%)	(24%)	(25%)
(B)	Pr (CO | CO prev round, $R > 1$)		0.31	0.43	0.77	0.79
			(24%)	(33%)	(59%)	(58%)
(C)	Pr (CO | no CO prev round, $R > 1$)		0.06	0.06	0.08	0.06
			(39%)	(37%)	(17%)	(17%)

Note: (*) CO refers to collaborative outcome. (%) represents the percentage of all collaborative outcomes that occur in a particular treatment corresponding to cases (A), (B), and(C). $R = 1$ represents the first round of a relationship, $R > 1$ represents all rounds other than the first one in a relationship.

In the Eng treatment the engineers recommendation is implemented in every period. Since the engineer is matched with the supplier as long as the relationship between the firms lasts, the Eng treatment should resemble the play in the LR treatment (hypothesis 2). The results show that there are no significant differences between the two treatments in terms of the suppliers sharing, and the buyers compete decisions

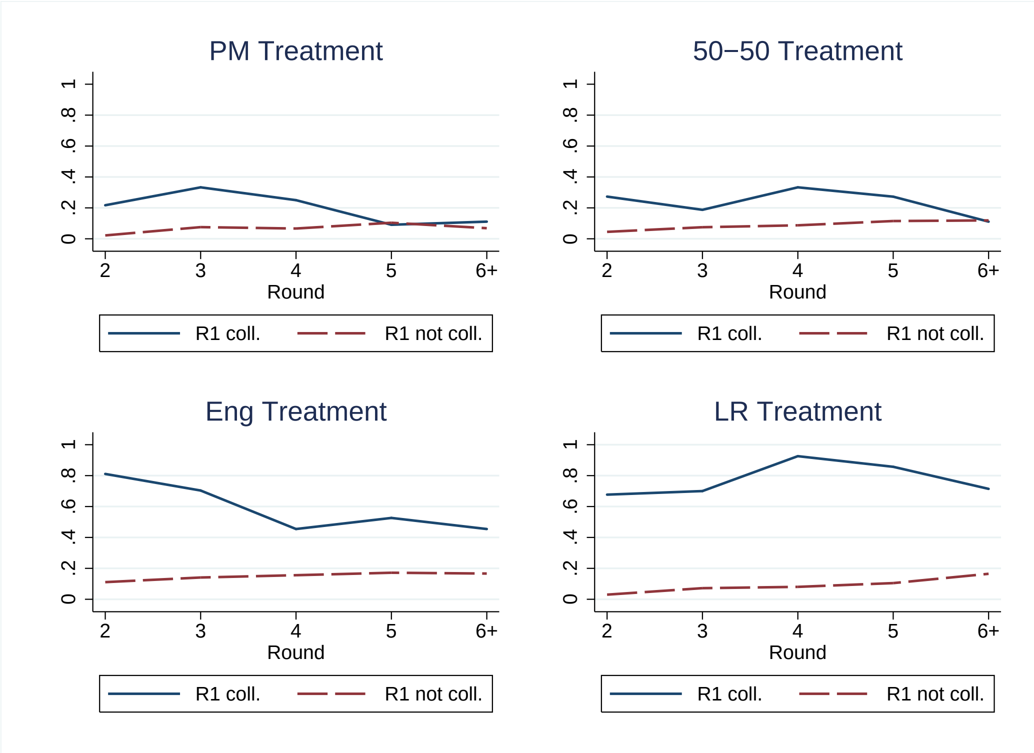

Figure 3.2: Effect of first round collaboration of subsequent rounds

in aggregate.

3.6.3.1 Interplay between employees

Figure 3.3 shows how the engineers play is affected by his interaction with the procurement manager in the Eng treatment. The bar chart on the left shows how the engineers choice to "compete if the supplier shared" in every round of the relationship (except the first one) differs depending on his own action in the round immediate before. We observe that if an engineer competed in one round, he is more likely to compete again in the following round within the same relationship than if he did not compete in the previous round. This difference in behavior is present even in later rounds within a relationship [21]. This results is not surprising: it is consistent with

[21] We considered the first six rounds of a relationship since it is the longest relationship that every supplier got to play. Thus, for this particular plot, we eliminated the observations from rounds 7 onwards.

hypothesis 2.b, which was derived assuming trigger strategies, but can also be the result of other common strategies such as "tit-for-tat" and "always compete". The bar chart on the right of Figure 3.3 shows a more surprising result: engineers seem to be more likely to compete if the previous procurement manager in the relationship chose to compete than if the previous procurement manager chose not to compete. Since this result can be intertwined with the engineer's own choice in the previous round, we conduct the regression presented in Table 3.7. Column 6 shows that, even after controlling for the engineer's own decision in the previous round, the engineer's decision is correlated with the previous procurement manager's decision in the round immediate before. This suggests that in the Eng treatment, where the procurement manager has no say in the final decision, the engineer takes into account the procurement manager and incorporates it into his own decision making. On the contrary, procurement managers do not consider the engineer's previous recommendation in the PM treatment (column 4). Procurement managers ignore the previous round play in the relationship and are only consistent with their own previous actions (note that the regression only considers the cases where in the previous round the innovation did occur, so that the procurement manager is informed of all the players' actions in the previous round). Columns 5 and 7 show that in the $50-50$ treatment, where both the engineer and the procurement manager have input on the final decision, both players ignore the previous recommendation of the player in the other role. Finally, columns 1, 2, and 3 show the supplier's actions in each of the Firms-as-Employees treatments. We find that the suppliers, as the engineers, care about the actions of the previous procurement manager even after he has left the relationship – and they correctly only do so when the procurement manager has a say in the final decision.

Table 3.7: Strategy Analysis

	Supplier's Decision			PM's Decision		Eng's Decision	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment	PM	$50-50$	Eng	PM	$50-50$	Eng	$50-50$
Coefficients	Share	Share	Share	Compete if Shared	Compete if Shared	Compete if Shared	Compete if Shared
Prev Shared Grp	0.586***	0.420***	0.757***	-0.064	-0.266	-0.005	-0.512***
	(0.152)	(0.136)	(0.165)	(0.172)	(0.169)	(0.174)	(0.174)
Prev Comp Eng Grp	-0.135	-1.545***	-2.099***	0.103	0.185	0.629***	0.697***
	(0.137)	(0.140)	(0.171)	(0.145)	(0.157)	(0.174)	(0.171)
Prev Comp PM Grp	-1.125***	-0.488***	-0.180	0.224	-0.003	0.355**	-0.017
	(0.136)	(0.143)	(0.169)	(0.157)	(0.170)	(0.181)	(0.179)
PM's own prev Comp				1.294***	0.786***		
				(0.175)	(0.176)		
Period	-0.019**	-0.014*	-0.003	-0.006	-0.020**	-0.014	-0.041***
	(0.008)	(0.008)	(0.010)	(0.010)	(0.010)	(0.010)	(0.011)
Round	0.032	0.021	-0.046	-0.015	-0.007	0.090**	-0.013
	(0.035)	(0.033)	(0.040)	(0.038)	(0.040)	(0.039)	(0.040)
Constant	0.030	0.786***	0.918***	-0.260	0.828**	-0.756**	1.351***
	(0.251)	(0.242)	(0.304)	(0.300)	(0.376)	(0.362)	(0.384)
Observations	620	687	596	620	687	596	687
Nr. of Subjects	28	31	27	55	62	52	61

Probit regression with subject random effects. Standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Note: the variable "Prev Shared Grp" takes value one if in the previous period the innovation happened and the supplier shared.

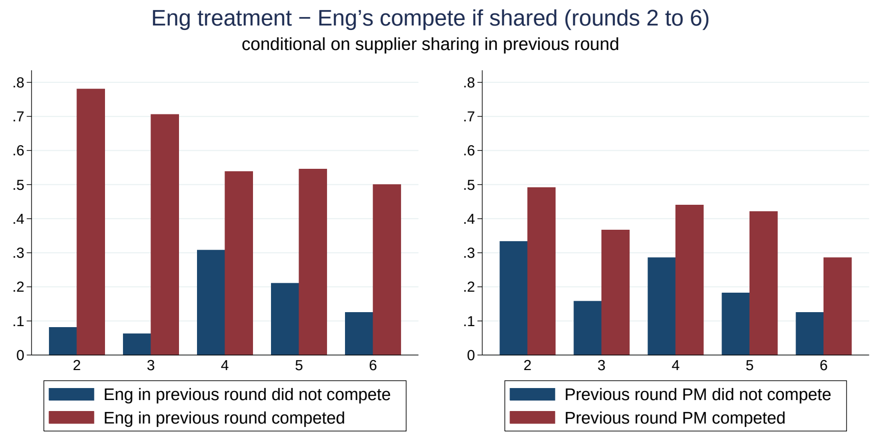

Figure 3.3: Engineer's Decision I

3.7 Discussion

The results of the $50-50$ treatment show the highest deviation from our theoretical predictions. While the theory predicts that procurement managers will always compete (as in the PM treatment) and engineers will never compete (as in the Eng treatment), we find that both procurement managers and engineers compete more than in the PM and Eng treatments. This result suggests that when there is uncertainty about whether their recommendation will be implemented, both types of players become less trustworthy. In a somewhat related setup, *Fudenberg et al.* (2012) study experimentally the play in a repeated prisoner's dilemma when the intended actions are implemented with noise. Their results show that introducing noise makes subjects slower to resort to punishment and more prone to forgive. With the uncertainty introduced by the $50-50$ treatment in our experimental design, suppliers do not seem to become more forgiving. The frequency of sharing after a round where the innovation happened, the supplier shared, and the buyer chose to compete was 0.31 in the $50-50$ treatment versus 0.30 in the PM treatment and 0.27 in the Eng treatment. Similarly, the overall level of sharing is not significantly lower in the $50-50$ treatment than in the PM and Eng treatments.

We also observe that, even when collaborative outcomes are infrequent in the $50-50$ treatment, a relatively high proportion of those collaborative outcomes (37%) occur when the procurement manager's recommendation is implemented. This result seems to be driven by a number of procurement managers who choose with a certain frequency not to compete when the supplier shared. While 50% of the procurement managers in the $50-50$ treatment choose to compete 90% of the times or more, another 23% choose to compete 50% of the times or less. Previous literature has shown that collaboration can be supported in a sub-game perfect equilibrium in a setting where subjects change partners over time. *Kandori* (1992b) and *Kandori* (1992a) extend the Falk Theorem result, which holds for infinitely repeated games between the same subjects, to a setting where a community can sustain collaboration if defection against one subject triggers punishment by other subjects, or if the subject who leaves overlaps with his successor for a long enough period of time. This would explain why some procurement manager's "compete if shared" choices are not 100% in the $50-50$ treatment. It could also explain why in the PM treatment procurement managers' "compete-if-shared" is lower than 100% and why in the SR treatment buyers' "compete-if-shared" is lower than 100% (they are 64% and 68% respectively).

Another surprising result is that procurement managers influence their group members' actions *after* they have left a relationship. Previous experimental literature on group decision making has found somehow similar results. *Ambrus et al.* (2013) study how individual preferences get aggregated in groups, where subjects choose how much to reciprocate as a second mover in a sequential gift exchange game (*Fehr et al.* 1993, *Brandts and Charness* 2004)[22]. In their setup, subjects freely discuss in groups of five individuals before making a group decision. Their results show

[22]The gift exchange game is similar in structure and incentives to the trust game. It captures the dynamics of an incomplete labor contract where the employee's effort is non-contractible or verifiable. Both players start with an initial endowment. The first mover sends a gift to the second mover where the gift is deducted from the first mover's endowment and is tripled by the experimenter. The second mover then decides whether to send a gift to the first mover under the same conditions.

that the relative position of an individual in a group is correlated with his influence on the other members of the group: median members and non-median members who are closer to the mean have significant influence on others. Interestingly, they find that deliberation causes that once subjects move on to play in another group, individual opinions move towards the previous decision of the group individuals participated in [23]. In our setting procurement managers influence the actions of the supplier and the engineer after they have left a relationship. This occurs in a setting which is less favorable for social influence: first, because our setting does not allow for discussion and deliberation [24], and second because in our setting subjects playing as procurement mangers and engineers differ in their monetary incentives in addition to their potential differences in preferences about the distribution of payoffs [25].

3.8 Conclusions

We analyze a case where a supplier has to decide whether to share an innovation with a buyer when sharing the innovation increases supply chain efficiency but makes the supplier vulnerable if the buyer re-shares the innovation with the supplier's competitors. The buyer decides what type of procurement policy he will follow: single source, which protects the suppliers' intellectual property rights for the innovation and distributes total profits more evenly between the firms, or to open up competition among suppliers, which takes advantage of the supplier's innovation sharing and gives

[23] *Ambrus et al.* (2013) reference two social psychology mechanisms which explain why subjects may behave different in group contexts. Social comparison theory proposes that individuals want to perceive and present themselves in a socially desirable way, and therefore they react in a way that is closer to a social norm. The identifiability explanation proposes that in a group setting others' ability to assign responsibility is more limited, allowing them to behave more selfishly.

[24] Related to this point, *Kocher and Sutter* (2007) find that results of a one-shot gift exchange game are closer to the standard game-theoretic prediction when the experiment is computerized and group members anonymously reach consensus by voting on proposals, than when group members can discuss face to face.

[25] Further literature on group decision making in trust games has focused for the most part in comparing how individuals and groups make decisions as senders and as receivers. *Cox* (2002) finds that groups in the role of responders send back smaller amounts than individuals, while *Kugler et al.* (2007) find that groups are just as trustworthy as individuals.

the buyer a larger share of total profits. As it is common in the automotive industry, the buyer may allocate decision rights to short-run and long-run focused employees. Anecdotal evidence from automotive suppliers tells that in different occasions it is either the short-run or the long-run focused employees that has more power in the decision making process. To study how this impacts firms' decisions, we conduct a laboratory experiment where both an engineer and a procurement manager make recommendations for what the buyer should do. As predicted by the theory, both the level of collaboration between the firms and supply chain efficiency are higher when the firms have a long term relationship. We also observe that, in addition to the length of the relationship between the firms, the allocation of decision rights to employees also matters. Having both short- and long-run focused employees involved in the decision (by making a recommendation), increases collaboration and efficiency, even if it is the short-run focused employee who has the final decision rights or if there is uncertainty about which recommendation will be chosen. However, the highest increase in collaboration and efficiency is reached when the decision rights are allocated to the long-run focused employee. When we analyze separately suppliers' and buyers' profits, we find that suppliers benefit only from long-run focused employees, while buyers benefit from any of the joint decision cases.

Most importantly, our results show that subjects' may be influenced by their peers' recommendations. In particular, it is the short run focused employee who has the strongest impact on the future play within the relationship: his actions are correlated with those of both the supplier and the long-run focused employee, but not those of his short-run focused successor. Understanding this interplay between employees is important for a buyer deciding whether (and how) to build teams to manage his supplier relations. Our experimental results suggest that: First, if the relationship is being managed by a short-run focused procurement manager, the buyer can benefit from introducing a long-run focused employee to the team. This can lead to increased efficiency without hurting the supplier. Second, if the long-run focused employee is in charge of making the decision, introducing a short-run focused employee may influence the decision maker's actions but does not lead to significantly worse outcomes in terms of efficiency or buyer's profits. Lastly, our results show that introducing uncertainty about which employee will be the final decision maker, leads to significantly lower collaboration by both types of employees. This is particularly detrimental for suppliers' profits.

APPENDICES

APPENDIX A

Proofs and Additional Figures and Tables for Chapter I

A.1 Proofs for the Theoretical Results

Proof of Lemma I.1: The proof is algebraic, therefore omitted.

Proof of Lemma I.2: Notice from Lemma 1 that offering a null contract to the selfish supplier is optimal. This is because, regardless of the contract, the selfish supplier will either reject the offer or exert zero effort. For the trustworthy supplier, notice also from Lemma 1 that, if $p \geq \gamma + \bar{u}_i$, the trustworthy supplier accepts the offer and exerts $e^{t*}(p,i) = c'^{-1}(\alpha_i \phi)$. If $\bar{u}_i < p \leq \gamma + \bar{u}_i$, he accepts the offer and exerts zero effort. If $p \leq \bar{u}_i$, he rejects the buyer's offer and earns $\overline{u}_i$. Thus, if the buyer makes an offer p, such that $p \geq \gamma + \bar{u}_i$, he will earn $\alpha_i e_i^{t*} - p$ and if his offer is such that $p < \gamma + \bar{u}_i$, he earns zero profit. From equation 3, the minimum price that yields an effort level of $e^{t*}(p,i) = c'^{-1}(\alpha_i \phi)$ is $T_i = \bar{u}_i + \gamma$. If $\alpha_i e_i^{t*} - \bar{u}_i - \gamma \geq 0$, doing so results in a strictly positive payoff for the buyer, therefore, it is optimal. Otherwise, offering the null contract is optimal.

Proof of Theorem I.3: The proof utilizes Lemma 1 and 2 and derives the optimal contract for both selfish and trustworthy suppliers.

Proof of Theorem I.4: To show sufficiency, first notice that condition *i*), $\bar{u}_g \geq \bar{u}_b + \gamma$, guarantees that the selfish supplier chooses the general investment as his profit with the general investment is greater than the profit he earns when he mimics the trustworthy supplier and chooses the specific investment. Likewise, condition *ii*), $(1-\phi)(\bar{u}_b+\gamma)-c(e_b^{t*})+\phi\alpha_b e_b^{t*} \geq \bar{u}_g$, guarantees that the trustworthy supplier chooses the specific investment. Conditions *i*) and *ii*) together imply condition C_b, thus it is optimal for the buyer to offer a contract T_b to the trustworthy supplier when he chooses the buyer-specific investment.

To show necessity, we show that the separating equilibrium breaks apart when any of conditions (i)-(iii) does not hold. If condition *i*) is not met, both types of suppliers choose the specific investment. If condition *ii*) does not hold, both types of suppliers choose the general investment. If condition *iii*) does not hold, the set of parameters under which the separating equilibrium exists becomes empty.

Proof of Theorem I.5:

We show that the three pooling equilibria described in Theorem I.5 exist and survive the intuitive criterion. We focus on the beliefs that give greater disincentive for deviation, except in the case where it violates the intuitive criterion. In this case, we consider alternative beliefs. Together, they characterize the full set of parameters under which the equilibrium can be supported with some beliefs.

We first prove the equilibrium described in part a) – both suppliers choose the buyer-specific investment and both suppliers accept the buyer's trusting contract T_b– exists and survives the intuitive criterion. Under this equilibrium, the selfish supplier exerts no effort, and the trustworthy supplier exerts $e_b^{t^*} = c'^{-1}(\alpha_b\phi) > 0$. We will show that this equilibrium arises if (i) $\theta \geq \tilde{\theta}_b$, with $\tilde{\theta}_b = \frac{\bar{u}_b+\gamma}{\alpha_b(c'^{-1}(\alpha_b\phi))}$; (ii) $\bar{u}_g < \bar{u}_b + \gamma$; and (iii) $\bar{u}_g < (1-\phi)(\bar{u}_b+\gamma) - c(e_b^{t^*}) + \phi\alpha_b e_b^{t^*}$.

Note that this equilibrium cannot exist when $\theta < \tilde{\theta}_b$. This is because the selfish supplier will exert no effort regardless of the contract. Hence, at low θ, the buyer finds it better off to deviate and offer a null contract. We now show that this equilibrium exists and survives the intuitive criterion with two sets of beliefs: (1) $\theta(t|b) = \theta$ and $\theta(t|g) = 0$ and (2) $\theta(t|b) = \theta$ and $\theta(t|g) = 1$.

Suppose first that $\theta(t|b) = \theta$ and $\theta(t|g) = 0$. Condition (i) implies condition C_b (in equation (5)). Hence, from Lemma 2, it is optimal for the buyer to offer a reciprocal contract, T_b under the buyer-specific investment. Note that under these beliefs, it is optimal for the buyer to offer the null contract under the general investment. Both suppliers choose the buyer-specific investment, and both suppliers accept the trusting contract T_b. The selfish supplier exerts no effort and the trustworthy supplier exerts effort $e_b^{t^*} = c'^{-1}(\alpha_b \phi) > 0$. The selfish supplier gets price $\bar{u}_b + \gamma$ and exerts zero effort under the specific investment and gets $\bar{u}_g$ under the general investment. Condition (ii), $\bar{u}_g < \bar{u}_b + \gamma$, makes choosing the specific investment incentive compatible and rational for the selfish supplier. The trustworthy supplier gets price $\bar{u}_b + \gamma$ and chooses effort $e_b^{t^*}$ under the specific investment. Thus, he derives utility $(1-\phi)(\bar{u}_b+\gamma) - c(e_b^{t^*}) + \phi\alpha_b e_b^{t^*}$ under the specific investment and $\bar{u}_g$ under the general investment. Condition (iii), $\bar{u}_g < (1-\phi)(\bar{u}_b+\gamma) - c(e_b^{t^*}) + \phi\alpha_b e_b^{t^*}$, guarantees that choosing the specific investment is incentive compatible and rational for the trustworthy supplier.

To show that this equilibrium survives the intuitive criterion, we need to consider two cases separately depending on whether condition C_g is met or not (Note that C_b already holds if $\theta \geq \tilde{\theta}_b$). Suppose C_g holds. Note that the intuitive criterion is violated if a deviating strategy is equilibrium-dominated [1] for the selfish supplier (i.e., $\bar{u}_b + \gamma > \bar{u}_g + \gamma$) but not for the trustworthy supplier (i.e., $(1-\phi)(\bar{u}_b + \gamma) - c(e_b^{t^*}) + \phi\alpha_b e_b^{t^*} \leq (1-\phi)(\bar{u}_g + \gamma) - c(e_g^{t^*}) + \phi\alpha_g e_g^{t^*})$[2]. However the first inequality cannot hold

[1] For the intuitive criterion a deviating strategy is defined as *equilibrium-dominated* if it gives the player a lower payoff than his equilibrium payoff for any belief the uninformed party may have following deviation.

[2] We set these conditions using the highest off-equilibrium payoffs a supplier can earn under any

for any $\gamma > 0$ since $\bar{u}_g \geq \bar{u}_b$. Thus, the equilibrium in part a) survives the intuitive criterion.

Now suppose that condition C_g does not hold. Thus, it is optimal for the buyer to offer a null contract to a trustworthy supplier when he makes the general investment. Again, the intuitive criterion is violated when the deviating strategy is equilibrium-dominated for the selfish supplier (i.e., $\bar{u}_b + \gamma > \bar{u}_g$) but not for the trustworthy supplier (i.e., $(1-\phi)(\bar{u}_b+\gamma) - c(e_b^{t^*}) + \phi\alpha_b e_b^{t^*} \leq \bar{u}_g$). But, this violates condition (iii). Thus, the equilibrium survives the intuitive criterion.

A similar reasoning can be applied to show that the same equilibrium with belief $\theta(t|b) = \theta$ and $\theta(t|g) = 1$ also exists if, in addition to conditions (i)-(iii), condition C_g is not met, and it survives the intuitive criterion if $(1-\phi)(\bar{u}_b+\gamma) - c(e_b^{t^*}) + \phi\alpha_b e_b^{t^*} \leq (1-\phi)(\bar{u}_g+\gamma) - c(e_g^{t^*}) + \phi\alpha_g e_g^{t^*}$. Likewise, we can show the pooling equilibrium described in part b) arises if (i) $\theta \geq \tilde{\theta}_g$; (ii) $\bar{u}_b \leq (1-\phi)(\bar{u}_g+\gamma) - c(e_g^{t^*}) + \phi\alpha_g e_g^{t^*}$; and (iii) $(1-\phi)(\bar{u}_g+\gamma) - c(e_g^{t^*}) + \phi\alpha_g e_g^{t^*} > (1-\phi)(\bar{u}_b+\gamma) - c(e_b^{t^*}) + \phi\alpha_b e_b^{t^*}$, and the equilibrium described in part c) arises if (i) $\theta < \tilde{\theta}_g$; (ii)$\bar{u}_g > \bar{u}_b + \gamma$; and (iii) $\bar{u}_g > (1-\phi)(\bar{u}_b+\gamma) - c(e_b^{t^*}) + \phi\alpha_b e_b^{t^*}$. The proof follows a similar logic, thus it is omitted.

<u>Theoretical Model for Random Investment Case:</u>

Consider a game where nature moves first to assign investment randomly, the buyer acts second to offer a take-it-or-leave-it price offer, and the supplier acts last and decides if he should accept the offer, and, if so, selects an effort level. Because the investment is randomly chosen, the buyer does not update his prior belief about the probability of the supplier being trustworthy, θ, after observing the investment. Analogous to the result found in the pooling equilibrium case, there exists a threshold $\tilde{\theta}_i = \frac{\bar{u}_i+\gamma}{\alpha_i c'^{-1}(\alpha_i\phi)}$ for $i = b, g$ (with $\tilde{\theta}_g \geq \tilde{\theta}_b$) above which the buyer offers a trusting contract or offers a null contract instead. When he is offered a trusting contract, the

belief. By doing this, we ensure that the equilibrium refinement by intuitive criterion is robust even under the highest possible incentives to deviate from equilibrium.

selfish supplier accepts and chooses $e = 0$ and when he is offered a null contract he rejects the offer. The trustwirthy supplier accepts a trusting contract and chooses $e_i^{r*} = c'^{-1}(\alpha_i \phi)$ and rejects a null contract.

One of the following three outcomes arises:

1) If $\theta \geq \tilde{\theta}_b$ and $\theta \geq \tilde{\theta}_g$, the buyer offers a trusting contract T_b under the specific investment and a trusting contract T_g under the general investment. With probability θ, the supplier will be trustworthy and will choose effort e_i^{t*} and with probability $(1-\theta)$ the supplier will be selfish and choose $e = 0$. Under the specific investment, the trustworthy supplier earns $\bar{u}_b + \gamma - c(c'^{-1}(\alpha_b \phi))$ and the selfish supplier earns $\bar{u}_b + \gamma$. The buyer's expected profit is $\theta[\alpha_b c'^{-1}(\alpha_b \phi) - (\bar{u}_b + \gamma)] + (1-\theta)[-(\bar{u}_b + \gamma)]$. Under the general investment, the trustworthy supplier earns $\bar{u}_g + \gamma - c(c'^{-1}(\alpha_g \phi))$ and the selfish supplier earns $\bar{u}_g + \gamma$. The buyer's expected profit is $\theta[\alpha_g c'^{-1}(\alpha_g \phi) - (\bar{u}_g + \gamma)] + (1-\theta)[-(\bar{u}_g + \gamma)]$.

2) If $\theta \geq \tilde{\theta}_b$ and $\theta < \tilde{\theta}_g$, the buyer offers a trusting contract under the specific investment $p = \bar{u}_b + \gamma$ and a null contract under the general investment. Under the general investment both types of supplier reject the offer. Under the specific investment, with probability θ the supplier is trustworthy and chooses effort e_b^{t*} and with probability $(1-\theta)$ the supplier is selfish and chooses $e = 0$. Under the general investment both types of supplier earn $\bar{u}_g$ and the buyer earns zero profits. Under the specific investment, the trustworthy supplier earns $\bar{u}_b + \gamma - c(c'^{-1}(\alpha_b \phi))$ and the selfish supplier earns $\bar{u}_b + \gamma$. The buyer earns $\theta[\alpha_b c'^{-1}(\alpha_b \phi) - (\bar{u}_b + \gamma)] + (1-\theta)[-(\bar{u}_b + \gamma)]$.

3) If $\theta < \tilde{\theta}_b$ and $\theta < \tilde{\theta}_g$, the buyer offers a null contract under both investments. Thus, both types of supplier reject the offer. The supplier earns $\bar{u}_b$ under the specific investment and $\bar{u}_g$ under the general investment. The buyer earns zero profits under both investments.

Repeated Interactions Model:

We analyze a repeated interactions model in which the supplier first makes the investment decision and then the supplier and the retailer engage in N transactions. In each round of transaction, the buyer makes an offer and the supplier decides whether he accepts the offer and, if so, how much effort to make.

The buyer's utility is the sum of profits over N periods:

$$U^B(p|e,i) = \sum_{t=1}^{N} \begin{cases} \alpha_i e_t - p_t & \text{if supplier accepts in period } t \text{ and exerts } e_t \\ 0 & \text{if supplier rejects in period } t. \end{cases} \tag{A.1}$$

The selfish supplier's utility is the total profit accrued:

$$U^s(e|i,p) = \sum_{t=1}^{N} [p_t - c(e_t)]. \tag{A.2}$$

The trustworthy supplier's utility is the sum of her utilities for the N periods [3]:

$$\begin{aligned} U^r(e|i,p) &= \sum_{t=1}^{N} \begin{cases} [p_t - c(e_t)] & \text{if } p_t < \gamma + \bar{u}_i \\ [(1-\phi)(p_t - c(e_t)) + \phi(\alpha_i e_t - c(e_t))] & \text{if } p_t \geq \gamma + \bar{u}_i \end{cases} \\ &= \sum_{t=1}^{N} \mathbf{1}_{\{p_t < \gamma + \bar{u}_i\}}[p_t - c(e_t)] + (1 - \mathbf{1}_{\{p_t < \gamma + \bar{u}_i\}})[(1-\phi)(p_t - c(e_t)) + \phi(\alpha_i e_t - c(e_t))] \end{aligned} \tag{A.3}$$

We first show that the same separating equilibrium where the trustworthy supplier chooses the specific investment and the selfish supplier chooses the general investment, exists in a finitely repeated game as well. Furthermore, a sufficient condition is exactly the same as that in the single period game (Theorem 2). In this equilibrium, the buyer always offers a trusting contract to a supplier who chose the specific investment, the trustworthy supplier accepts and exerts effort $e_b^{t*} = c'^{-1}(\alpha_b \phi)$. On the other hand, the

[3] We assume that the supplier's reciprocity concerns depend only on the latest offer received. It is not affected by offers received in previous transaction periods.

buyer offers a null contract to a supplier who chose a general investment, which the selfish supplier rejects. This result is formalized in Theorem 4, which stipulates that in a finitely repeated interaction game; there exists a separating equilibrium that is the same as in Theorem I.4.

Proof of Theorem I.6: We will show that the sufficient conditions for a separating equilibrium to arise in the single period case, $i)\bar{u}_g \geq \bar{u}_b + \gamma$ and $ii)(1-\phi)(\bar{u}_b + \gamma) - c(e_b^{t*}) + \alpha_b \phi e_b^{t*} \geq \bar{u}_g$, are also sufficient for the separating equilibrium to arise in N-period games.

Condition $i)$ guarantees that the selfish supplier does not want to deviate from the equilibrium. To see why, consider first the case where the selfish supplier chooses the general investment. In this case, the buyer believes that the supplier is selfish and offers a null contract in every period. Thus, the selfish supplier rejects the contract in every period and earns a total profit of $N\bar{u}_g$.

Consider now the case where the selfish supplier deviates and chooses the specific investment. We show that one of the following two strategies will dominate any other strategy. The first strategy is to exert effort e_b^{t*} for the first $N-1$ periods and exert effort $e = 0$ in period N. This strategy will result in a profit of $(N-1)[\bar{u}_b + \gamma - c(e_b^{t*})] + \bar{u}_b + \gamma$. The other strategy is to exert effort $e = 0$ in the first period and then reject the contract throughout, which results in a profit of $\bar{u}_b + \gamma + (N-1)\bar{u}_b$. To show this, note first that in any period the supplier's optimal effort is always one of the following two: $e = 0$ or $e_b^{t*} = c'^{-1}(\alpha_b \phi)$. This is because any other effort leads the buyer to believe that the supplier is selfish and offer him a null contract, so the supplier prefers $e = 0$. If $c(e_b^{t*}) > \gamma$, then $e = 0$ is optimal in every period. We show this by backward induction. In the last period the selfish supplier always chooses $e = 0$. In period $N-1$, choosing $e = 0$ has no cost and results in a profit of $\bar{u}_b$ in the following period. On the other hand, choosing e_b^{t*} has a cost of $c(e_b^{t*})$ but earns him $\bar{u}_b + \gamma$ in the following period. Thus, if $c(e_b^{t*}) > \gamma$, the supplier chooses $e = 0$. The same logic applies to every period prior to $N-1$ (all the way up to period 1), so the supplier chooses $e=0$ in every period. This strategy earns the supplier a total profit of $\bar{u}_b+\gamma+(N-1)\bar{u}_b$ which, under condition i), is dominated by his payoff under the general investment. If $c(e_b^{t*})\leq\gamma$, exerting effort e_b^{t*} for the first $N-1$ periods and $e=0$ in period N, dominates any other strategy. To see why, note first that in period N the selfish supplier always chooses $e=0$. In periods 1 through $N-1$ exerting effort e_b^{t*} has cost $c(e_b^{t*})$ earns the supplier $\bar{u}_b+\gamma$ in the following period, and choosing $e=0$ has no cost but earns the supplier $\bar{u}_b$ in the following period. Thus, if $c(e_b^{t*})\leq\gamma$, the supplier chooses e_b^{t*} in periods 1 to $N-1$. In this case, the supplier earns $(N-1)[\bar{u}_b+\gamma-c(e_b^{t*})]+\bar{u}_b+\gamma$ which, under condition i), is dominated by his payoff under the general investment.

Now we show that the trustworthy supplier does not want to deviate from the equilibrium. Consider first the case where the trustworthy supplier chooses the specific investment. In every period, the optimal effort is either e_b^{t*} or $e=0$. If in any period, the supplier chooses anything other than e_b^{t*}, then in every subsequent period she will be offered a null contract, which she rejects and gets utility $\bar{u}_b$. If the supplier is offered a trusting contract and exerts effort e_b^{t*}, then she gets utility $(1-\phi)(\bar{u}_b+\gamma)-c(e_b^{t*})+\alpha_b\phi e_b^{t*}$ in that period and is offered a trusting contract again in the following period. Thus, if conditions i) and ii) hold, it is optimal for the supplier to exert effort e_b^{t*} in every period, which gets her a total utility of $N[(1-\phi)(\bar{u}_b+\gamma)-c(e_b^{t*})+\alpha_b\phi e_b^{t*}]$. To see why, we solve by backward induction. In the last period, the utility from exerting effort e_b^{t*} is $(1-\phi)(\bar{u}_b+\gamma)-c(e_b^{t*})+\alpha_b\phi e_b^{t*}$ and the utility from $e=0$ is $\bar{u}_b+\gamma$. Thus, if conditions i) and ii) are met, the optimal effort is e_b^{t*}. In period $N-1$, if the supplier exerts effort e_b^{t*}, then she gets utility $(1-\phi)(\bar{u}_b+\gamma)-c(e_b^{t*})+\alpha_b\phi e_b^{t*}$ in periods $N-1$ and N. If she chooses $e=0$, she gets $\bar{u}_b+\gamma$ in period $N-1$ and $\bar{u}_b$ in period N. Thus, if conditions i) and ii) are met, the optimal effort is e_b^{t*}. If we continue solving backwards until period 1, the same logic shows that, if conditions i) and ii) are met, the optimal effort is e_b^{t*} in every period.

Consider now the case where the trustworthy supplier chooses the general investment. The supplier is offered a null contract in every period and she rejects the offer in every period. Thus, she derives utility $N\bar{u}_g$. Condition *ii*) guarantees that the utility she derives from the equilibrium payoff, $N[(1-\phi)(\bar{u}_b+\gamma)-c(e_b^{t*})+\alpha_b\phi e_b^{t*}]$, is greater or equal than $N\bar{u}_g$.

Lastly, conditions *i*) and *ii*) guarantee that the buyer does not want to deviate from equilibrium under the equilibrium beliefs. First, note that if a supplier chooses the general investment, the buyer believes the supplier is selfish and it is optimal to offer a null contract in every period. If the supplier chooses the specific investment, the buyer believes the supplier is trustworthy. Conditions *i*) and *ii*) guarantee that it is optimal for the buyer to offer a trusting contract to a trustworthy supplier in every period. To see why, note that conditions *i*) and *ii*) combined imply condition C_b, that is $\alpha_b(c'^{-1}(\alpha_b\phi))-\bar{u}_{-}\gamma \geq 0$. If the buyer offers a trusting contract to a trustworthy supplier in any given period, the trustworthy supplier will exert effort e_b^{t*} and the buyer gets profit $\alpha_b(c'^{-1}(\alpha_b\phi))-\bar{u}_b-\gamma$. If he offers a null contract, the trustworthy supplier will reject it and the buyer earns zero profit in every period thereafter. Thus, if conditions i) and ii) hold, it is optimal for the buyer to offer a trusting contract in every period. ■

In a finitely repeated game we also find that an interesting equilibrium exists (we call it semi-separating [4]) which leads to collaborative outcomes in every period except for the last. Under this equilibrium, both types of supplier choose the specific investment. The buyer offers a trusting contract under the specific investment and a null contract under the general investment. The trustworthy supplier chooses effort

[4]We call this equilibrium a "semi-separating equilibrium" since the selfish and trustworthy suppliers make the same investment choice and their actions coincide in all trading periods except for the last one. As a result, their actions coincide at every point where their actions affecting the buyer's beliefs has impact on the buyer's future actions. Both types of supplier's actions only differ in the last trading period, where they no longer affect the buyer's actions.

$e_b^{t*} = c'^{-1}(\alpha_b \phi)$ in every period, and the selfish supplier also chooses effort $e_b^{t*} = c'^{-1}(\alpha_b \phi)$ in every period, except the last period, where he accepts the offer and chooses zero effort. This result is formalized in Theorem I.7, which stipulates that in a finitely repeated game, there exists a semi-separating equilibrium under which both suppliers choose the specific investment, and the buyer offers a trusting contract T_b in every period. Upon receiving the contract, the trustworthy supplier exerts effort $e_b^{t*} = c'^{-1}(\alpha_b \phi)$ for all periods, and the selfish supplier exerts the same effort except in period N in which he exerts zero effort.

The semi-separating equilibrium exists if the buyer's belief that the supplier is trustworthy, θ, is high enough. In particular, it requires $\theta \geq \tilde{\theta}_b^N$. The threshold $\tilde{\theta}_b^N = \frac{\bar{u}_b + \gamma}{\alpha_b e_b^{t*}}$ is also the threshold above which a buyer offers a trusting contract in a pooling equilibrium of the single interaction model. Additional conditions on the parameters are necessary to guarantee that the equilibrium is incentive compatible and rational for all players. The proof is similar as in Theorem I.6 and thus, omitted.

A.2 Robustness Checks - High and Low α Treatments

In this section we analyze how the results of our main treatment change when we vary the values of the parameters. In particular, we are interested in changes to the value of α_b (the quality coefficient under the specific investment). Lower levels of the quality coefficient result in greater incentives for the trustworthy suppliers to choose the general investment. Thus, by reducing the value of α_b from 12 in the main treatment to 6 [5] in the additional *low benefit treatment*, we expect to find results that are more consistent with a pooling equilibrium on the general investment than with a separating equilibrium. In particular, we predict that the result observed in the main

[5]The new value of α_b for the low benefit treatment is picked so that the set of values of ϕ and γ that allow for the separating equilibrium to occur is considerably smaller that in the main treatment. For example, in the low benefit treatment, a separating equilibrium arises only when ϕ is between 0.84 and 0.91 and γ is between 0 and 2. See Figure A.1 in the Supplementary Tables and Graphs section for the set of parameters for which the separating equilibrium holds.

treatment, where there was a positive correlation between the specific investment and suppliers' reciprocal behavior, to fall apart in the low benefit treatment.

In addition, we conduct a new additional treatment where we raise the value of α_b to 18. In this case, we expect to find results that are consistent with a separating equilibrium, with the benefits of the specific investment being even more prominent relative to the general investment in terms of efficiency and profitability.

We conducted five sessions of the low benefit treatment and four sessions of the high benefit treatment with a total number of 56 and 46 participants respectively.

Our experimental results support the predictions. As opposed to the main the treatment, in the low benefit treatment we find that the general investment is chosen more often than the specific investment (60% vs. 40%), consistently with a higher probability of finding a pooling equilibrium on the general investment. We find that the low benefit treatment is less efficient than the main treatment in terms of overall effort and profits. Overall exerted effort is lower in the low benefit treatment than in the main treatment (1.22 in the low benefit treatment and 1.47 in the main treatment, $p = 0.069$). Similarly, overall buyer's profit is lower in the low benefit treatment relative to the main treatment (92.24 versus 95.6, $p = 0.03$). In addition, when we look at trends across the treatments with $\alpha_b = 6$, 12 and 18, a non parametric test for trends shows that the expected effort presents a positive trend as α_b increases ($p = 0.001$). A similar results holds for total profits. A non parametric test for trends shows that total profit presents a positive trend as α_b increases ($p < 0.001$) for the treatments with $\alpha_b = 6$, 12 and 18.

We also find that the profit premiums of the specific investment relative to the general investment for both buyers and suppliers are no longer present in the low benefit treatment. Average buyer's profit is 96.88 under the specific investment and 92.94 under the general investment in the main treatment ($p = 0.02$). In the low benefit treatment, it is 91.03 and 93.04 respectively ($p = 0.038$). In addition, in the low benefit treatment a buyer does not make a significantly higher profit at an individual level under the specific investment than under the general investment (Wilcoxon signed-rank test: $p > 0.20$), as was the case in the main treatment. Table A.3 presents the interaction effects of the specific investment and treatment on profits. While under the main treatment the specific investment leads to increased profits (both jointly and individually), in the low benefit treatment it consistently leads to lower profits. On the other hand, in the high benefit treatment ($\alpha_b = 18$) the difference between investments increases. Buyers' average profit was 101.968 under the specific investment and 67.333 under the general investment, a significant difference ($p < 0.01$).

When we examine the subject-level behavior in the low benefit treatment, we find that the mechanism driving our results in the main treatment is no longer present: there is no longer a positive correlation between the supplier being more reciprocal and choosing the specific investment more often. In the low benefit treatment, we find that there is still sorting - with some subjects choosing the specific investment more often than others. Subjects who choose the specific investment more often in the first five periods also choose it more often in the last five periods, both as measured by correlation ($\rho = 0.526$, $p < 0.05$), and a non-parametric trend test($p = 0.007$). Similarly, the permutation test also indicates a significantly larger number of subjects choosing the specific investment frequently ($p < 0.05$ for frequencies greater than seven). However, we do not find a difference in supply chain game play between subjects who choose the specific investment five or more times and subjects who choose it four of fewer times.[6] Unlike the main treatment, when we regress effort on price (with the specific investment) distinguishing these two groups, we find very similar price coefficients for high- and low-frequency subjects ($\beta = 0.097$ and $\beta = 0.093$ respectively) and the difference in coefficients is not significant ($p = 0.904$). In

[6]The cutoff point of five was chosen so that the fraction of suppliers above and below the cutoff point is the closest to that in the main treatment, where the cutoff point was eight. If we use the cutoff points derived from the permutation test, 6 (marginally significant) or 7 (significant), the results do not change.

the main treatment we also considered as a measure of a subject's reciprocity level, his behavior in the additional investment game. While in the main treatment we find that higher levels of reciprocity in the additional investment game are correlated with choosing the specific investment more often, neither trust nor reciprocity in the trust game are correlated with choosing the specific investment at least five times in the low benefit treatment (Table A.7 in the supplementary tables section).

Table A.1: Interaction Effects of Specific Investment and Treatment on Profit

Coefficients	Suppliers' Profit	Buyers' Profit	Total Profit
Specific x Main Treatment	3.167**	4.863***	5.357***
	(1.346)	(1.589)	(1.680)
Specific x Low Benefit Treatment	-2.806*	-1.170	-3.682**
	(1.691)	(1.440)	(1.481)
Specific x Random Treatment	-0.202	0.263	1.131
	(1.960)	(2.048)	(1.937)
Low Benefit Treatment	1.145	0.388	-0.387
	(1.401)	(2.116)	(1.028)
Random Treatment	2.147	-1.296	-1.550
	(1.443)	(1.858)	(1.059)
Constant	76.892***	92.319***	171.014***
	(0.865)	(1.072)	(0.766)
Observations	1220	1220	1220
Nr. of Subjects	122	122	122

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

A.3 Supplementary Tables and Graphs

Table A.2: Profit Comparison

Treatment	Investment	Average Supplier's Profit	Average Buyer's Profit	Average Total Profit
Main	Specific	80.23	96.88	177.11
	General	76.54	92.94	169.48
	Rank Sum test (p-value)*	0.97	0.02	0.99
Random	Specific	78.94	91.46	170.40
	General	78.91	90.80	169.71
	Rank Sum test (p-value)*	0.067	0.943	< 0.001

* Non-parametric test of difference in average profits between general and specific investments.

Table A.3: Interaction Effects of Specific Investment and Treatment on Profit

Coefficients	Suppliers' Profit	Buyers' Profit	Total Profit
Specific x Main Treatment	3.216**	4.725***	5.352***
	(1.341)	(1.565)	(1.681)
Specific x Random Treatment	-0.183	0.320	1.132
	(1.955)	(2.040)	(1.938)
Random Treatment	2.169	-1.421	-1.554
	(1.442)	(1.856)	(1.060)
Constant	76.859***	92.412***	171.018***
	(0.862)	(1.058)	(0.767)
Observations	940	940	940
Nr. of Subjects	94	94	94

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table A.4: Effect of Treatment on Overall Effort and Total Profit

Coefficients	Overall Effort	Total Profit
Low Benefit Treatment	-0.244	-5.461***
	(0.274)	(1.243)
Random Treatment	-0.673***	-4.532***
	(0.191)	(1.537)
Constant	1.466***	174.628***
	(0.144)	(1.093)
Observations	1220	1220
Number of Subjects	122	122

OLS with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table A.5: Risk and Trust Summary Statistics

Variable	Mean	Median	Std. Dev.	Min	Max
Trust	8.104	4	7.399	0	20
Trustworthiness	16.515	18	13.111	0	40
Risk aversion	0.518	0.533	0.145	0	1

Table A.6: Effect of Risk and Trust Measures

Coefficients	Chose Specific 8+			Price			Effort	
HT	1.298***			9.920***			-1.049	
	(0.357)			(3.773)			(0.646)	
HTS		0.630**			4.909			-0.501
		(0.314)			(3.603)			(0.641)
HRA			-0.047			-1.729		
			(0.313)			(4.587)		
P*HT							0.092***	
							(0.007)	
P*LT							0.062***	
							(0.010)	
P*HTS								0.094***
								(0.009)
P*LTS								0.072***
								(0.008)
P*HRA								
P*LRA								
Con	-1.044***	-0.508**	-0.180	16.880***	18.458***	21.968***	-0.619	-1.058**
	(0.297)	(0.219)	(0.238)	(1.830)	(2.252)	(4.179)	(0.502)	(0.442)
Obs	67	67	67	670	670	670	483	483
# of S.	67	67	67	67	67	67	66	66

Columns 1 to 3: Probit regression. Standard errors reported in parentheses. Columns 4 to 6: OLS with subject random effects. Robust standard errors reported in parentheses. Columns 7 and 8: Tobit with subject random effects, accepted offers only. Standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Variables denoted: HT = High Trust, HTS = High Trustworthiness, HRA = High Risk Aversion, LT = Low Trust, LTS = Low Trustworthiness, LRA = Low Risk Aversion.

Table A.7: Effect of Trust and Reciprocity on Investment Choice - Additional Treatments

Coefficients	Main Treatment Chose Specific 8+		Low Benefit Treatment Chose Specific 5+		Random Treatment Chose Specific 6+	
High Trust	1.298***		-0.053		-0.271	
	(0.357)		(0.482)		(0.492)	
High Reciprocity		0.630**		-0.275		-0.097
		(0.314)		(0.492)		(0.483)
Constant	-1.044***	-0.508**	-0.157	-0.074	0.157	0.097
	(0.297)	(0.219)	(0.315)	(0.304)	(0.315)	(0.348)
Observations	67	67	28	28	27	27
Nr. of Subjects	67	67	28	28	27	27

Probit regression. Standard errors reported in parentheses. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table A.8: Profit Comparison - Main vs. Repeated Interactions

Treatment	Main			Repeated Interactions		
Investment	Average Supplier's Profit	Average Buyer's Profit	Average Total Profit	Average Supplier's Profit	Average Buyer's Profit	Average Total Profit
Specific	80.23	96.88	177.11	136.549	110.961	247.51
General	76.54	92.94	169.48	125.510	63.75	189.26
p-value*	0.97	0.02	0.996	0.024	< 0.0001	< 0.0001

* Wilcoxon Rank-sum test of difference in average profits between general and specific investments.
Note: Profits include initial endowment (60 for suppliers and 100 for buyers).

Table A.9: Total Profit on Specific Investment Choice - Repeated Interactions

Coefficients	Supplier's Profit	Buyer's Profit	Total Profit
Specific	9.746	45.640***	44.931***
	(7.429)	(12.067)	(8.645)
Constant	126.389***	64.818***	198.317***
	(5.267)	(10.009)	(3.883)
Observations	150	150	150
Number of Subjects	25	25	25

OLS with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

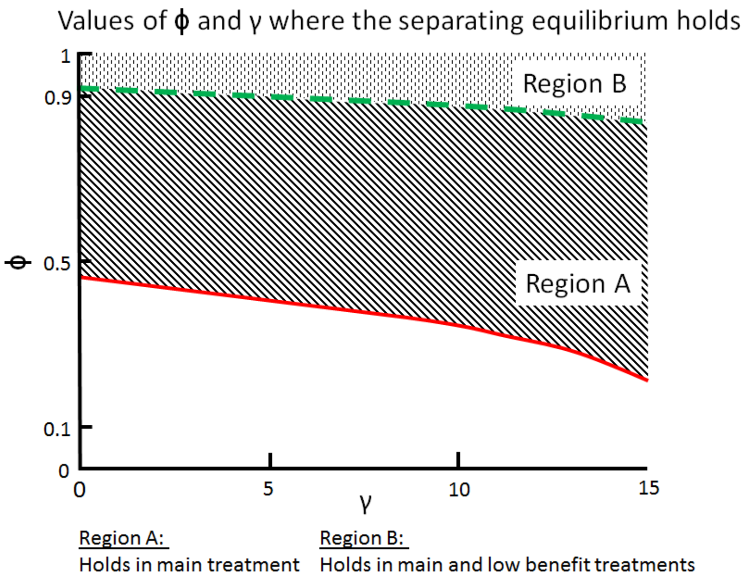

Figure A.1: Area where Separating Equilibrium Holds

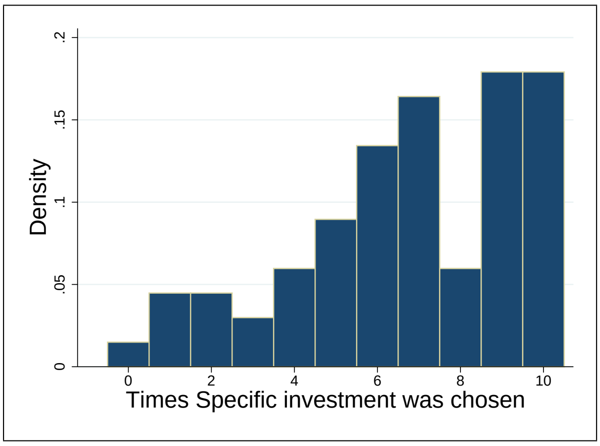

Figure A.2: Frequency of Choosing Specific Investment

APPENDIX B

Proofs and Additional Figures and Tables for Chapter II

Proof of Proposition II.1: We show that, when awards are not available ($A_1 = A_2 = 0$), all possible equilibria have $p_t = 0$ and $q_t = 0$ and any requested quality can be supported in equilibrium. We solve by backward induction starting with the last period. In the last stage of period 2, for any subgame starting at period 2, it is easy to see from equation (2.1) that $q_2 = 0$ is a dominant strategy. Knowing this, a buyer maximizes his utility (equation (2.2)) by offering price zero. Considering this as a subgame outcome, the first period transaction is identical as the second transaction. Each buyer makes an offer of $p_1 = 0$ and any requested quality, and suppliers choose $q_1 = 0$.

Proof of Proposition II.2: While in the no-award case all possible equilibria have $q_1 = q_2 = 0$, in the private award scenario there exists a separating equilibrium in which a low-type supplier exerts no effort and does not receive an award and a high-type supplier exerts strictly positive effort and receives an award in both periods.

First, let $\pi^i_k \in (0, 1)$ be a player's prior belief that supplier i's type is k, $k \in$ {h = high, l = low}, and $\pi(i, k|A_1)$ be the updated belief about supplier i's type at the beginning of period 2 when the supplier receives an award (or not) in period 1.

We show that the following is a truth-telling equilibrium of the game, where the buyer gives the supplier an award if and only if the supplier's quality is greater or equal to the buyer's requested quality. The buyers' initial beliefs are that both suppliers are of high type with probability π and of low type with probability $(1-\pi)$. In the first transaction, buyers offer $p_1 = 0$ and $\hat{q}_1 = C^{-1}(\phi)$, low type suppliers choose $q_1 = 0$ and do not get the award, and high type suppliers choose $q_1 = \hat{q}_1$ and get the award. After the first transaction period, buyer j's updated beliefs about his own supplier's type are $\pi(i = j, h|1) = 1$ and $\pi(i = j, h|0) = 0$, and about the other supplier are $\pi(i \neq j, h|A_1) = \pi_h^i = \pi$. That is, if his own supplier received and award, the buyer believes the supplier is of high type and otherwise, he believes the supplier is of low type. Since the buyer does not observe whether the other supplier received an award or not, his updated belief remains equal to the initial belief; the supplier is of hight type with probability π and of low type with probability $(1-\pi)$. In the second transaction, the buyer makes a continuing offer to his own supplier, a low type supplier chooses $q_2 = 0$ and does not get the award, and a high type supplier chooses $q_2 = \hat{q}_2$ and receives the award.

We show that none of the players has incentives to deviate from equilibrium. Consider first the low type supplier. In the second transaction he chooses $q_2 = 0$ since he does not derive any utility from the award. In the first transaction, if he chooses $q_1 = \hat{q}_1$ he induces the buyer to believe that he is high type, and if he chooses $q_1 = 0$ he induces the buyer to believe that he is low type. However, in either case, the buyer makes a continuing offer (which has price equal to zero). Thus, the low type supplier chooses $q_1 = 0$ in transaction 1. Consider now the high type supplier. In the second transaction, the high type supplier solves:

$$\max_{q_2 \geq 0} \begin{cases} \phi + p_2 - C(q_2) & \text{if } q_2 \geq \hat{q}_2 \\ p_2 - C(q_2) & \text{otherwise.} \end{cases}$$

If the supplier chooses q_2 such that $q_2 \geq \hat{q}_2$, then the optimal choice is $q_2 = \hat{q}_2$. If he chooses q_2 such that $q_2 < \hat{q}_2$ then the optimal choice is $q_2 = 0$. If the cost of getting the award, $C(\hat{q}_2)$ exceeds the utility the supplier gets from it, ϕ, the supplier chooses $q_2 = 0$ and does not get the award. Otherwise, he chooses $q_2 = \hat{q}_2$ and gets the award. Thus, the high type supplier will choose $q_2 = \hat{q}_2$ if and only if $\phi \geq C(\hat{q}_2)$ or equivalently, $\hat{q}_2 \leq C^{-1}(\phi)$. In equilibrium, the requested quality is $\hat{q}_2 = C^{-1}(\phi)$, therefore the high type supplier chooses $q_2 = \hat{q}_2$ and gets the award. In the first transaction, the high type supplier chooses either $q_1 = 0$, which induces the buyer to believe that he is of low type, or $q_1 = \hat{q}_1$, which induces the buyer to believe that he is of high type. Since, the buyers' actions in transaction 2 are not affected by their beliefs about the suppliers' type (the buyers make a continuing offer in period 2 with $p_2 = 0$ regardless of their updated beliefs), the suppliers solve the same problem as in transaction 2. Thus, in transaction 1 the high type supplier chooses $q_1 = \hat{q}_1$ and gets the award.

Finally, consider the buyer's incentives to deviate from equilibrium. In equilibrium, both buyers offer $p_1 = 0$ and $\hat{q}_1 = C^{-1}(\phi)$ in transaction 1, and make a continuing offer to their own supplier in transaction 2. We solve backwards starting with period 2. Consider first a buyer whose belief is that his supplier is of low type ($A_1 = 0$). If the offer in period 1 had $p_1 = 0$, making a continuing offer is a weakly dominating strategy. If p_1 was not zero, the optimal strategy is to make a new offer with $p_2 = 0$. To see why, note that if he makes a new offer to his own supplier, he can earn at most his equilibrium payoff, which is zero. If he makes an offer to the other supplier – whom he believes to be of high type with probability π and of low type with probability $(1 - \pi)$ – he can make at most zero profits. This is because any offer to the other supplier which could earn him positive profits (if the other supplier is of high type), would be matched by the other buyer. Consider now a buyer who believes his supplier is of high type ($A_1 = 1$). If in period 1 the offer had $p_1 = 0$, it is (weakly) optimal to make a continuing offer. If $p_1 > 0$, he makes a new offer to his own supplier with $p_2 = 0$, since in equilibrium the other buyer makes a continuing offer to his own supplier. Rolling back to the first period, making an offer with $p_1 = 0$ and $\hat{q}_1 = C^{-1}(\phi)$ grants the buyer the maximum profit he can earn considering it is never profitable to switch suppliers in period 2. If the supplier is of high type, the buyer earns $2\alpha C^{-1}(\phi)$ and if the supplier is low type, he earns zero profits.

We additionally note that when the awards are private, all equilibria have price equal to zero in both transactions. To see why, consider first a buyer who is matched with a low type supplier. In the second transaction, any offer to his own supplier will earn him zero quality and should therefore have $p_2 = 0$. If he makes an offer with price strictly greater than zero to the other supplier, with probability π the other supplier is of high type and any offer than earns him a positive profit (i.e. p_2 and $\hat{q}_2$ such that $\alpha\hat{q}_2 - p_2 > 0$) will be matched by the other buyer, and with probability $(1-\pi)$ the supplier is low type so he wins the deal but earns negative profit. Thus, the buyer who is matched with a low type supplier does not make an offer with strictly positive price to the other supplier. Consider now a buyer who is paired with a high type supplier. If $0 < p_1 < \alpha C^{-1}(\phi)$, making a continuing offer to his own supplier is a (weakly) dominating strategy since it allows him to match any profitable offer that the other buyer makes to his supplier. If $p_1 > \alpha C^{-1}(\phi)$, it is optimal to make a new offer with $p_2 = 0$. Rolling backwards, given that his supplier will not receive an offer with strictly positive price from the other buyer in transaction 2, the buyer has no incentive to offer a price higher than zero in transaction 1.

Proof of Proposition II.3: First, let us define $\bar{p}$ as the maximum price that allows the buyer to earn positive profits in transaction 2 when the requested quality is $\hat{q}_2 = C^{-1}(\phi)$, that is $\bar{p} = \alpha C^{-1}(\phi)$. We show that the following is an equilibrium of the game with public awards. In the first transaction, the buyer offers $p_1 = 0$ and $\hat{q}_1 = C^{-1}(\phi + (1-\pi)\bar{p})$. The high type supplier chooses $q_1 = \hat{q}_1$ and gets the award, and the low type supplier chooses $q_1 = 0$ and does not get the award. After the first transaction, buyer j updates his beliefs about the suppliers' types as follows: $\pi(i, h|A_1 = 1) = 1$ and $\pi(i, h|A_1 = 0) = 0$. That is, each buyer observes whether each supplier received an award or not in period 1. If the supplier received an award in period 1, they believe the supplier is of high type, and otherwise they believe the supplier is of low type. In the second transaction, if the buyer believes that both suppliers are of high type, the buyer makes a new offer to his own supplier with $p_2 = 0$ and $\hat{q}_2 = C^{-1}(\phi)$. If he believes that both suppliers are of low type, he makes a continuing offer. If he believes that one of the suppliers is of high type and the other supplier is of low type, he makes a new offer to the high type supplier with price $p_2 = \bar{p}$ and requested quality $\hat{q}_2 = C^{-1}(\phi)$. The high type supplier accepts the offer of the buyer he was matched with in period 1 and the other buyer makes a zero price offer to the low type supplier. The low type supplier chooses $q_2 = 0$ and does not get the award, and the high type supplier chooses $q_2 = \hat{q}_2$ and gets the award. We can see that in this equilibrium, if one supplier is of high type and the other supplier is of low type (occurs with probability $2\pi(1 - \pi)$), then in period 2 the high type supplier is offered $p_2 = \bar{p} > 0$.

We show that none of the players wants to deviate from equilibrium. Consider first a high type supplier's incentive to deviate. In transaction 2, a high type supplier chooses $q_2 = \hat{q}_2$ if and only if $\hat{q}_2 \leq C^{-1}(\phi)$. Since this holds in equilibrium, he chooses $q_2 = \hat{q}_2$. In transaction 1, requested quality is $\hat{q}_1 = C^{-1}(\phi + (1-\pi)\bar{p})$. The supplier in this case maximizes expected utility. If he chooses $q_1 = \hat{q}_1$, both buyers will update their beliefs to high type. Thus, with probability $(1 - \pi)$ the other supplier is of low type and he earns $p_2 = \bar{p}$ and with probability π, the other supplier is also of high type and he will be offered $p_2 = 0$. This results in an expected utility of $\phi - C(C^{-1}(\phi + (1 - \pi)\bar{p})) + (1 - \pi)\bar{p} + \phi - C(C^{-1}(\phi)) = 0$. If he chooses $q_1 = 0$, he induces the buyer to believe that he is of low type and this earns him zero utility.

If he chooses any $q_1 > 0$ and different from $\hat{q}_1$ this earns him negative utility. Thus, choosing $q_1 = \hat{q}_1$ is a (weakly) dominating strategy.

Consider now the low type supplier. In transaction 2, it is always optimal for the low type supplier to choose $q_2 = 0$. In the first transaction, the low type supplier could potentially have incentive to mimic the high type supplier and choose $q_1 = \hat{q}_1$ so that the buyer believes he is of high type. However, this is never optimal since it would earn him an expected utility of $[(1-\pi)\bar{p}] - C(C^{-1}(\phi + (1-\pi)\bar{p}))$. Choosing $q_1 = 0$ earns him zero profits and choosing any $q_1 > 0$ and different from $\hat{q}_1$ earns him strictly negative profits. Thus, choosing $q_1 = 0$ is a weakly dominating strategy.

Finally, consider the buyer's incentives to deviate. In the second transaction, if he believes that both suppliers are of high type, making a new offer to his own supplier with $p_2 = 0$ and $\hat{q}_2 = C^{-1}(\phi)$ gets the buyer the highest profit he can get. If he believes that both suppliers are of low type, he makes a continuing offer and earns zero profits. No other offer can earn him a positive profit since both suppliers will always choose $q_2 = 0$. If the buyer believes that his own supplier is of low type and the other supplier is of high type, the buyer is willing to offer up to $\bar{p}$ to win the high type supplier (since this supplier offers up to $q_2 = C^{-1}(\phi)$, he would set $\hat{q}_2 = C^{-1}(\phi)$). Thus, if the buyer's own supplier is of high type and the other supplier is of low type, the buyer also offers $\bar{p}$ and requests $\hat{q}_2 = C^{-1}(\phi)$ and keeps the high type supplier. In equilibrium, the buyer whose supplier is of low type ends up making a new zero price offer to his own supplier. Both buyers earn zero profits from transaction 2 in this case. Rolling back to the first transaction, any price greater than zero is dominated by price equal to zero, as p_1 does not affect the suppliers' actions. Requested quality $\hat{q}_1 = C^{-1}(\phi + (1-\pi)\bar{p})$ makes the high type supplier indifferent from choosing $q_1 = 0$ and is therefore, the maximum quality the buyer can get in transaction 1.

Proof of Proposition II.4: First, note that the buyers' expected profit in the no-awards case is always zero, since prices and quality are always equal to zero.

Additionally, we showed that with private awards, price is always equal to zero and there exists an equilibrium where quality is strictly positive. Thus, buyers' expected profit is higher in the private awards case than in the no awards case. In the public awards case, buyers' expected profit is $\pi[\alpha C^{-1}(\phi + (1-\pi)\alpha C^{-1}(\phi)) + \pi\alpha C^{-1}(\phi)]$, which is strictly positive. Finally, we note that buyers' expected profits are higher when awards are private than when awards are public if $2\alpha C^{-1}(\phi) > \alpha C^{-1}(\phi + (1-\pi)\alpha C^{-1}(\phi)) + \pi\alpha C^{-1}(\phi)$. Whether this condition holds depends on the values of the problem's parameters. For the canonical quadratic cost function used in the experiment, α and π need to be sufficiently low, and ϕ needs to be sufficiently large.

When we focus on truth-telling equilibria, where a buyer gives the supplier an award if and only if the supplier chooses $q_t \geq \hat{q}_t$, the equilibrium outcomes of the scenario with public quality and private award are equivalent to those in the scenario with public award and private quality. The only difference between the two settings is that when quality is public (and the award is not), buyers update their beliefs based on the quality the suppliers chose, $\pi(i,k|q_1)$, instead of updating their beliefs based on whether the supplier received an award or not, $\pi(i,k|A_1)$. If the supplier chose a quality greater or equal to the requested quality, they believe the supplier is of high type and otherwise they believe the supplier is of low type (i.e. $\pi(i,h|q_1 \geq \hat{q}_1) = 1$, and $\pi(i,h|q_1 < \hat{q}_1) = 0$). Note that in equilibrium, if suppliers choose $q_i \geq \hat{q}_i$, they receive an award and otherwise they do not receive an award. Thus, observing whether the supplier receives an award or not is equivalent to observing whether $q_i \geq \hat{q}_i$, and the analysis of the two cases is analogous.

Table B.1: Differences between Private and Public Award

Coefficients	Suppliers' accepted price (Transaction 2)	Suppliers' accepted price if received award in transaction 1 (Transaction 2)	Buyers' Profit (Total)
Public Award	7.017**	8.908**	-8.293*
	(3.176)	(4.343)	(4.719)
Constant	28.495***	31.417***	196.080***
	(2.149)	(2.853)	(2.900)
Observations	181	76	228
Nr. of Subjects	73	46	75

OLS regression with subject random effects. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

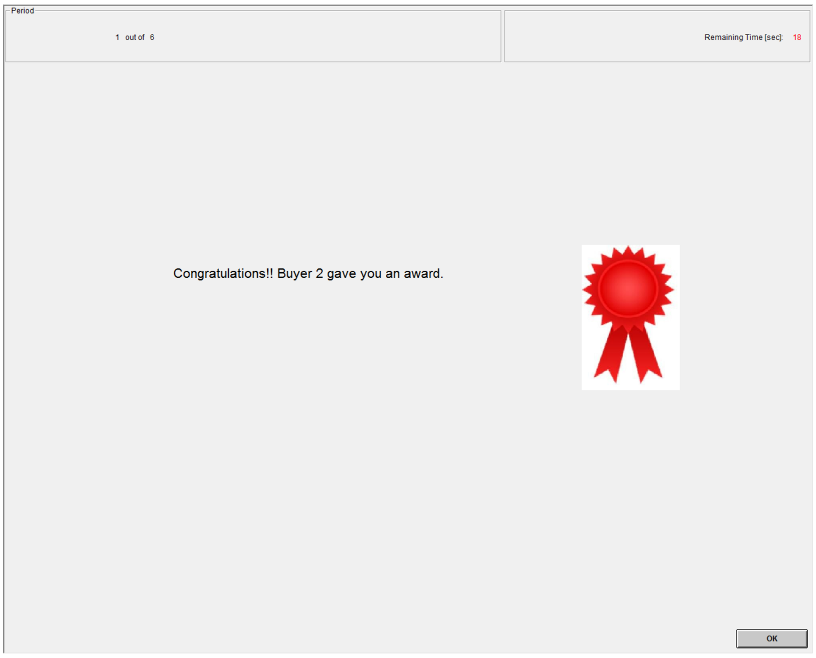

Figure B.1: Award Screenshot

APPENDIX C

Proofs and Additional Figures and Tables for Chapter III

Table C.1: Values of parameters for the numerical example

Parameter	Description	Value
a	Demand parameter	50
b	Demand parameter	2
C_{s1}	Supplier 1's cost before innovation	7
C_{s2}	Supplier 1's cost if innovation occurs	5
C_{m1}	Manufacturer's cost before supplier shares	11
C_{m2}	Manufacturer's cost if innovation occurs and supplier shares	5
C_{a1}	Supplier 2's cost before manufacturer shares	9
C_{a2}	Supplier 2's cost if manufacturer chooses to compete	5

	Innovation does not Occur	Innovation Occurs	
		Supplier: Share	**Supplier**: do not Share
Buyer: Compete	Supplier = **7** Buyer = **17**	Supplier = **0** Buyer = **112**	Supplier = **18** Buyer = **22**
Buyer: do not Compete	Supplier = **12** Buyer = **12**	Supplier = **56** Buyer = **56**	Supplier = **20** Buyer = **20**

Figure C.1: Payoffs table shown in the experiment

Table C.2: Supplier's sharing

Coefficients	Average profit per round (Supplier)
Average sharing in a relationship	5.043***
	(1.950)
Match	0.358
	(0.291)
Constant	13.43***
	(1.769)
Observations	220
Nr. of Subjects	28

Tobit regression with subject random effects. Standard errors reported in parentheses. Each relationship represents one observation. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Note:

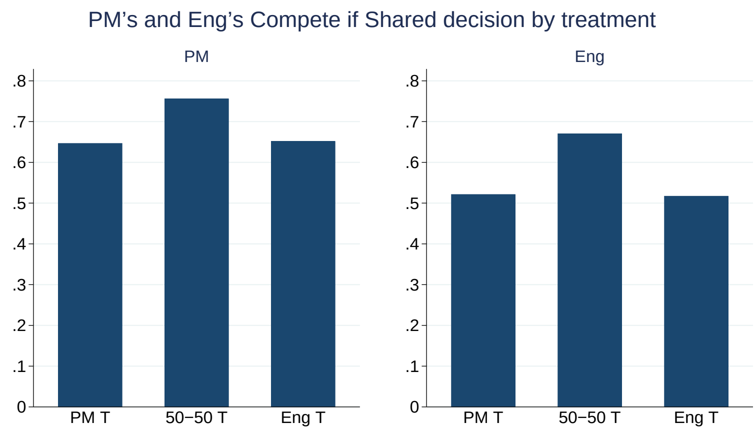

Figure C.2: Procurement Manager's and Engineer's Decision

C.1 Proofs and Numerical Example

Manufacturer chooses not to compete: Bilateral bargaining We consider the case where both firms agree on a contract that splits profits according to some parameter α. In the bilateral case, we assume that the total surplus in the supply chain will be split in such a way that the manufacturer earns a fraction α of the total surplus, $\Pi_m = Q\alpha(p - C_s - C_m)$, and the supplier earns a fraction $(1 - \alpha)$ of the total surplus $\Pi_s = Q(1 - \alpha)(p - C_s - C_m)$.

The manufacturer and the supplier simultaneously choose p^* and w^* that maximize

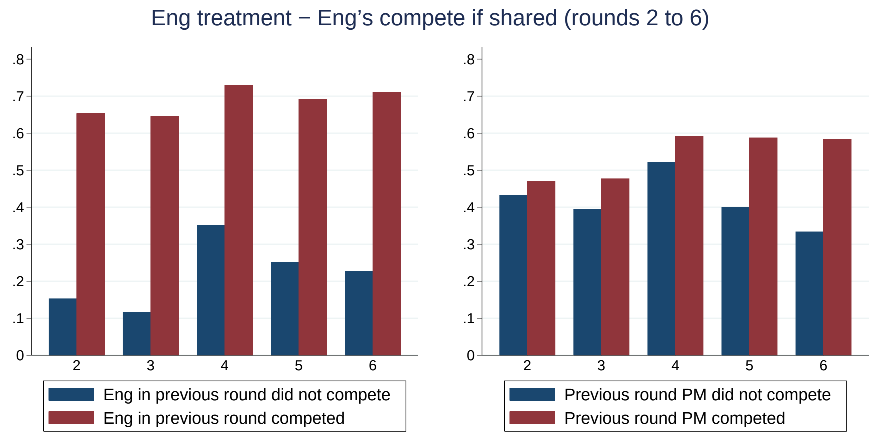

Figure C.3: Engineer's Decision II

total surplus while keeping the Nash Bargaining allocation of surplus between them. That is, they solve:

$$\max_{p,w} \quad Q(p - C_m - C_s)$$

$$\text{s.t. } \Pi_m(p,w) = \alpha Q(p - C_m - C_s)$$

$$\text{and } \Pi_s(p,w) = (1-\alpha) Q(p - C_m - C_s)$$

Taking FOC, the optimal retail price is $p^* = \frac{a+b(C_m+C_s)}{2b}$. At this retail price, the quantity sold is $Q^* = \frac{a-b(C_m+C_s)}{2}$. The supplier's wholesale price w^*, is such that earns the supplier $(1-\alpha)$ times total surplus. That is, w^* such that $Q(w - C_s) = (1-\alpha)Q(p - C_m - C_s)$. Then $w^* = (1-\alpha)(p - C_m) + \alpha C_s$ [1] .

It is a commonly known result that in the case where the manufacturer and the supplier have the same bargaining power and they both get zero profits in case of disagreement, the Nash Bargaining Solution predicts equal splits of the surplus, that

[1]For more on the surplus split in case of monopolies with exogenous bargaining power, see *Lovejoy* (2010).

is $\alpha = \frac{1}{2}$. Thus, replacing for the manufacturer's and supplier's profits with $\alpha = \frac{1}{2}$, we get:

$$\Pi_m = \Pi_s = \frac{(a - b(C_s + C_m))^2}{8b} \tag{C.1}$$

Manufacturer chooses to compete: Bargaining with supplier competition

Consider now the case where the manufacturer chooses to compete. If the supplier shared, then the firms have costs $C_{s2} = C_{a2}$ and if the innovation did not occur or if occurs and the supplier chose not to share, then the firms have costs $C_{a1} > C_{s1}$ and $C_{a1} > C_{s2}$ respectively. We assume that, in either case, the original supplier wins the deal. The Nash Bargaining solution dictates that the manufacturer and the supplier find the split α^* that solves

$$\underset{\alpha}{\text{argmax}}[(u_s - t_s)(u_m - t_m)]$$

where u_s is the supplier's agreement payoff, $(1 - \alpha)Q(p - C_s - C_m)$; t_s is the supplier's disagreement payoff, 0; u_m is the manufacturer's agreement payoff $\alpha Q(p - C_s - C_m)$; and t_m is the manufacturer's disagreement payoff, $\beta Q(p - C_a - C_m)$. We assume $\beta = 1$ since the manufacturer can extract the whole surplus from the high-cost supplier.

Thus, Nash Bargaining dictates that the total surplus will be allocated according the α that solves:

$$\underset{\alpha}{\text{argmax}}[(1 - \alpha)Q(p - C_s - C_m)][\alpha Q(p - C_s - C_m) - Q(p - C_a - C_m)].$$

The solution to this problem is $\alpha^* = \frac{1}{2} + \frac{(p - C_a - C_m)}{2(p - C_s - C_m)}$. Given this split of surplus, the manufacturer and supplier simultaneously find the optimal p^* and w^* that result in maximum total surplus while splitting it according to α. They solve:

$$\max_{p,w} \quad Q(p - C_m - C_s)$$

$$\text{s.t. } \Pi_m(p,w) = \alpha Q(p - C_m - C_s)$$

$$\text{and } \Pi_s(p,w) = (1-\alpha)Q(p - C_m - C_s)$$

Taking FOC, the optimal retail price is $p^* = \frac{a+b(C_m+C_s)}{2b}$. At this retail price, the quantity sold is $Q^* = \frac{a-b(C_m+C_s)}{2}$. The wholesale price w^8, is such that earns the supplier $(1-\alpha)$ times total surplus. That is, such that $Q(w - C_s) = [\frac{1}{2} - \frac{(p-C_a-C_m)}{2(p-C_s-C_m)}][Q(p - C_s - C_m)]$, which yields $w^* = \frac{C_a+C_s}{2}$. At this retail price, the quantity sold is $Q^* = \frac{a-b(C_m+C_s)}{2}$.

Replacing α^* in the manufacturer's and supplier's profits, we get that the manufacturer's profit is:

$$\Pi_m = \frac{[a - b(C_m + C_s)][a - b(C_m + C_a)]}{4b}, \tag{C.2}$$

and the supplier's profit is:

$$\Pi_s = \frac{[C_a - C_s][a - b(C_s + C_m)]}{4}. \tag{C.3}$$

Proof of Proposition II.1:

For the supplier, the strategy described in Proposition II.1 is an equilibrium strategy if the present value from a perpetuity of collaborative actions when the innovation happens is greater than the payoff of a one time "Not Share" decision plus the present value of a perpetuity of non-collaborative payoffs. That is, if $\frac{1}{(1-\delta)}(\pi ISN_s + (1-\pi)NC_s) \geq \frac{1}{(1-\delta)}(\pi INC_s + (1-\pi)NC_s)$, or equivalently $ISN_s \geq INC_s$. For the manufacturer the strategy described in Proposition II.1 is an equilibrium strategy if his present value from a perpetuity of collaborative outcomes is greater than his profit from a one-time Share-Compete outcome plus the present value of a perpetuity of non-collaborative actions by both firms. That is, if $ISN_m + \frac{1}{(1-\delta)}[\pi ISN_m + (1-\pi)NC_m] - [\pi ISN_m + (1-\pi)NC_m] \geq ISC_m + \frac{1}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m] - [\pi INC_m + (1-\pi)NC_m]$, or equivalently $\delta \geq \frac{ISN_m - ISC_m}{\pi INC_m + (1-\pi)ISN_m - ISC_m}$.

Proof of Proposition II.2:

Analogous to the proof of Proposition II.1, the supplier does not want to deviate from collaboration if $\frac{1}{(1-\delta)}[\pi(\frac{ISN_s + ISC_s}{2}) + (1-\pi)NC_s] \geq \frac{1}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$, or equivalently $\frac{ISN_s + ISC_s}{2} \geq INC_s$. In the numerical example, $ISN_s = 56$ is given by equation 3.1 with C_{s2} and C_{m2}, $ISC_s = 0$ is given by equation 3.3 with C_{s2}, C_{m2}, and C_{a2}, and $INC_s = 18$ is given by equation 3.3 with C_{s2}, C_{m1}, and C_{a1}.

For the engineer does not want to deviate from collaboration if $ISN_m + \frac{1}{(1-\delta)}[\pi(\frac{ISN_m + ISC_m}{2}) + (1-\pi)NC_m] - [\pi(\frac{ISN_m + ISC_m}{2}) + (1-\pi)NC_m] \geq ISC_m + \frac{1}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m] - [\pi INC_m + (1-\pi)NC_m]$, or equivalently $\delta \geq \frac{ISN_m - ISC_m}{\pi INC_m + (1-\frac{\pi}{2})ISN_m - (1+\frac{\pi}{2})ISC_m}$. In the numerical example, $ISN_m = 56$ is given by equation 3.1 with C_{s2} and C_{m2}, $ISC_m = 112$ is given by equation 3.2 with C_{s2}, C_{m2}, C_{a2}, and $INC_m = 22$ is given by equation 3.2 with C_{s2}, C_{m1}, and C_{a1}.

Lastly, the procurement manager will always choose compete. This is because he maximizes his profit from a one-shot game, which dictates to compete regardless of the other players' actions in the current or previous periods.

Numerical Example

We fix the values of the parameters ($a, b, C_{s1}, C_{s2}, C_{m1}, C_{m2}, C_{a1}, C_{a2}$, and π) and calculate the payoffs under these conditions. Figure 3.1 shows the extensive form of the game with these payoffs [2].

Suppose that supplier 1 started off with production cost $C_{s1} = 7$ and the manu-

[2]Note that this game resembles the structure of the basic Trust Game (Kreps 1990).

facturer with a cost $C_{m1} = 11$. If the innovation occurs, the supplier's cost is reduced to $C_{s2} = 5$. If the supplier chooses to share the technology with the manufacturer, the manufacturer's cost is reduced to $C_{m2} = 5$. If the supplier does not share the technology, the manufacturer's cost remains $C_{m1} = 11$.

In the second period, the manufacturer can choose to bring in another supplier, who initially has cost $C_{a1} = 9$. In the case where the first supplier (supplier 1) shared the technology with the manufacturer, if the manufacturer chooses to bring in a new supplier (alternative supplier, "a"), the manufacturer then shares the technology with the alternative supplier whose cost is reduced to $C_{a2} = 5$. Otherwise, the alternative supplier's cost remains $C_{a1} = 9$.

We assume that the innovation occurs with probability $\pi = 0.75$. In the case where the innovation does happen, we observe the following:

Consider first the case where the supplier shared the technology. If the manufacturer chooses to compete, then the optimal retailer price is $p = 17.5$, the wholesale price is $w = 5$ and the quantity sold is $Q = 15$. This results in a total surplus of 112.5. In this case, we get $\alpha = 1$, that is, the manufacturer keeps all the surplus and the supplier gets nothing [3]. If the manufacturer chooses not to compete, then the optimal retailer price is $p = 17.5$, the wholesale price is $w = 8.75$ and the quantity sold is $Q = 15$. This results in a total surplus of 112.5. In this case, we get $\alpha = \frac{1}{2}$ (the manufacturer and the supplier split profits equally and earn 56 each).

Consider now the case where the supplier did not share the technology with the manufacturer. In this case, if the manufacturer chooses to compete, then the optimal retailer price is $p = 20.5$, the wholesale price is $w = 7$ and the quantity sold is $Q = 9$. This results in a total surplus of 40.5. In this case, we get $\alpha = 0.556$, that is, the manufacturer keeps 55.6% of the surplus and earns 22.5, and the supplier gets 44.4% and earns 18. If the manufacturer chooses not to compete, then the optimal retailer

[3] We assume that if supplier 1 and the alternative supplier have the same costs, supplier 1 wins the deal.

price is $p = 20.5$, the wholesale price is $w = 7.25$ and the quantity sold is $Q = 9$. This results in a total surplus of 40.5. In this case, we get $\alpha = \frac{1}{2}$ (the manufacturer and the supplier split profits equally and earn 20.25 each).

In the case where the innovation does not happen, we observe the following:

If the manufacturer chooses to compete, then the optimal retailer price is $p = 21.5$, the wholesale price is $w = 8$ and the quantity sold is $Q = 7$. This results in a total surplus of 24.5. In this case, we get $\alpha = 0.714$, that is, the manufacturer keeps 71.4% of the surplus and earns 17.5, and the supplier gets 28.6% and earns 7. If the manufacturer chooses not to compete, then the optimal retailer price is $p = 21.5$, the wholesale price is $w = 8.75$ and the quantity sold is $Q = 7$. This results in a total surplus of 24.5. In this case, we get $\alpha = \frac{1}{2}$ (the manufacturer and the supplier split profits equally and earn 12.25 each).

BIBLIOGRAPHY

BIBLIOGRAPHY

Ambrus, A., B. Greiner, and P. Pathak (2013), How individual preferences get aggregated in groups-an experimental study, *Economic Research Initiatives at Duke (ERID) Working Paper*, (158).

Andreoni, J., and J. Miller (2002), Giving according to garp: An experimental test of the consistency of preferences for altruism, *Econometrica*, pp. 737–753.

Andreoni, J., and L. Vesterlund (2001), Which is the fair sex? gender differences in altruism, *Quarterly Journal of Economics*, pp. 293–312.

Aoyagi, M., and G. Fréchette (2009), Collusion as public monitoring becomes noisy: Experimental evidence, *Journal of Economic theory*, *144*(3), 1135–1165.

Ariely, D., A. Bracha, and S. Meier (2009), Doing good or doing well? image motivation and monetary incentives in behaving prosocially, *The American Economic Review*, pp. 544–555.

Asanuma, B. (1989), Manufacturer-supplier relationships in japan and the concept of relation-specific skill, *Journal of the Japanese and international economies*, *3*(1), 1–30.

Ashraf, N., O. Bandiera, and K. Jack (2012), No margin, no mission, *A Field Experiment on Incentives for Pro-Social Tasks, CEPR Discussion Papers*, *8834*.

Aviv, Y. (2001), The effect of collaborative forecasting on supply chain performance, *Management science*, *47*(10), 1326–1343.

Aviv, Y. (2007), On the benefits of collaborative forecasting partnerships between retailers and manufacturers, *Management Science*, *53*(5), 777–794.

Baker, G., R. Gibbons, and K. Murphy (2002), Relational contracts and the theory of the firm, *The Quarterly Journal of Economics*, *117*(1), 39–84.

Baron, J. N., and D. M. Kreps (1999), *Strategic human resources: Frameworks for general managers*, Wiley New York.

Becker-Peth, M., E. Katok, and U. Thonemann (2011), Designing contracts for irrational but predictable newsvendors, *Available at SSRN 1567614*.

Beer, R., H.-S. Ahn, and S. Leider (2014), Can trustworthiness in a supply chain be signaled?, *Ross School of Business Paper*, (1251).

Berg, J., J. Dickhaut, and K. McCabe (1995), Trust, reciprocity, and social history, *Games and economic behavior*, *10*(1), 122–142.

Bernstein, F., and A. G. Kök (2009), Dynamic cost reduction through process improvement in assembly networks, *Management Science*, *55*(4), 552–567.

Besley, T., and M. Ghatak (2008), Status incentives, *The American Economic Review*, pp. 206–211.

Bó, P. D. (2005), Cooperation under the shadow of the future: experimental evidence from infinitely repeated games, *American Economic Review*, pp. 1591–1604.

Bó, P. D., and G. R. Fréchette (2011), The evolution of cooperation in infinitely repeated games: Experimental evidence, *The American Economic Review*, *101*(1), 411–429.

Bohnet, I., and R. Zeckhauser (2004), Trust, risk and betrayal, *Journal of Economic Behavior & Organization*, *55*(4), 467–484.

Bohnet, I., B. Frey, and S. Huck (2001), More order with less law: On contract enforcement, trust, and crowding, *American Political Science Review*, *95*(1), 131–144.

Bohnet, I., F. Greig, B. Herrmann, and R. Zeckhauser (2008), Betrayal aversion: Evidence from brazil, china, oman, switzerland, turkey, and the united states, *The American Economic Review*, pp. 294–310.

Bolton, G., and E. Katok (2008), Learning by doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback, *Manufacturing & Service Operations Management*, *10*(3), 519–538.

Bolton, G., and A. Ockenfels (2000), Erc: A theory of equity, reciprocity, and competition, *American economic review*, pp. 166–193.

Bolton, G. E., E. Katok, and A. Ockenfels (2004), How effective are electronic reputation mechanisms? an experimental investigation, *Management science*, *50*(11), 1587–1602.

Bradler, C., R. Dur, S. Neckermann, and A. Non (2013), Employee recognition and performance: A field experiment, *ZEW-Centre for European Economic Research Discussion Paper*, (13-017).

Brandts, J., and G. Charness (2004), Do labour market conditions affect gift exchange? some experimental evidence*, *The Economic Journal*, *114*(497), 684–708.

Bresnahan, T., and P. Reiss (1985), Dealer and manufacturer margins, *The RAND Journal of Economics*, pp. 253–268.

Brinkhoff, A., Ö. Özer, and G. Sargut (2014), All you need is trust? an examination of inter-organizational supply chain projects, *Production and Operations Management.*

Brinkhoff, A., Ö. Özer, and G. Sargut (2015), All you need is trust? an examination of inter-organizational supply chain projects, *Production and Operations Management*, *24*(2), 181–200.

Burt, D. N. (1989), Managing suppliers up to speed, *Harvard Business Review*, *67*(4), 127–135.

Cabral, L., E. Y. Ozbay, and A. Schotter (2014), Intrinsic and instrumental reciprocity: An experimental study, *Games and Economic Behavior*, *87*, 100–121.

Cachon, G. (2003), Supply chain coordination with contracts, *Handbooks in operations research and management science*, *11*, 229–340.

Cachon, G., and M. Lariviere (2005), Supply chain coordination with revenue-sharing contracts: strengths and limitations, *Management science*, *51*(1), 30–44.

Camera, G., and M. Casari (2009), Cooperation among strangers under the shadow of the future, *The American Economic Review*, *99*(3), 979–1005.

Camerer, C. (2003), *Behavioral game theory: Experiments in strategic interaction*, Princeton University Press.

Charness, G., D. Masclet, and M. C. Villeval (2013), The dark side of competition for status, *Management Science*, *60*(1), 38–55.

Cho, I., and D. Kreps (1987), Signaling games and stable equilibria, *The Quarterly Journal of Economics*, *102*(2), 179–221.

Cooper, D. J., and K.-U. Kuhn (2009), Communication, renegotiation, and the scope for collusion.

Cox, J. C. (2002), Trust, reciprocity, and other-regarding preferences: Groups vs. individuals and males vs. females, in *Experimental Business Research*, pp. 331–350, Springer.

Cui, T., and P. Mallucci (2010), Quantal response equilibrium in a fair channelan experimental investigation, *Tech. rep.*, Working paper.

Cui, T., and P. Mallucci (2012), Fairness ideals in distribution channels, *Available at SSRN 2021248.*

Cui, T., J. Raju, and Z. Zhang (2007), Fairness and channel coordination, *Management Science*, *53*(8), 1303–1314.

Cuzick, J. (1985), A wilcoxon-type test for trend, *Statistics in medicine*, *4*(4), 543–547.

Dal Bó, P., and G. R. Fréchette (2013), Strategy choice in the infinitely repeated prisoners dilemma, *Available at SSRN 2292390.*

Dohmen, T., and A. Falk (2011), Performance pay and multidimensional sorting: Productivity, preferences, and gender, *The American Economic Review*, *101*(2), 556–590.

Doney, P., and J. Cannon (1997), An examination of the nature of trust in buyer-seller relationships, *The Journal of Marketing*, pp. 35–51.

Dufwenberg, M., and G. Kirchsteiger (2004), A theory of sequential reciprocity, *Games and Economic Behavior*, *47*(2), 268–298.

Dyer, J. (1996a), Specialized supplier networks as a source of competitive advantage: Evidence from the auto industry, *Strategic management journal*, *17*(4), 271–291.

Dyer, J., and H. Singh (1998), The relational view: Cooperative strategy and sources of interorganizational competitive advantage., *Academy of management review*, *23*(4), 660–679.

Dyer, J. H. (1996b), How chrysler created an american keiretsu, *Harvard Business Review.*

Dyer, J. H., and N. W. Hatch (2006), Relation-specific capabilities and barriers to knowledge transfers: creating advantage through network relationships, *Strategic management journal*, *27*(8), 701–719.

Ellingsen, T., and M. Johannesson (2008), Anticipated verbal feedback induces altruistic behavior, *Evolution and Human Behavior*, *29*(2), 100–105.

Engle-Warnick, J. (2007), *Five indefinitely repeated games in the laboratory*, Citeseer.

Engle-Warnick, J., and R. L. Slonim (2004), The evolution of strategies in a repeated trust game, *Journal of Economic Behavior & Organization*, *55*(4), 553–573.

Engle-Warnick, J., and R. L. Slonim (2006a), Inferring repeated-game strategies from actions: evidence from trust game experiments, *Economic theory*, *28*(3), 603–632.

Engle-Warnick, J., and R. L. Slonim (2006b), Learning to trust in indefinitely repeated games, *Games and Economic Behavior*, *54*(1), 95–114.

Englmaier, F., and S. Leider (2012), Contractual and organizational structure with reciprocal agents, *American Economic Journal: Microeconomics*, *4*(2), 146–83, doi: 10.1257/mic.4.2.146.

Eriksson, T., and M. C. Villeval (2012), Respect and relational contracts, *Journal of Economic Behavior & Organization*, *81*(1), 286–298.

Falk, A., and U. Fischbacher (2006), A theory of reciprocity, *Games and Economic Behavior*, *54*(2), 293–315.

Fehr, E., and A. Falk (1999), Wage rigidity in a competitive incomplete contract market, *Journal of political Economy*, *107*(1), 106–134.

Fehr, E., and S. Gächter (2000), Cooperation and punishment in public goods experiments, *American Economic Review*, pp. 980–994.

Fehr, E., and K. Schmidt (1999), A theory of fairness, competition, and cooperation, *The quarterly journal of economics*, *114*(3), 817–868.

Fehr, E., G. Kirchsteiger, and A. Riedl (1993), Does fairness prevent market clearing? an experimental investigation, *The Quarterly Journal of Economics*, *108*(2), 437–459.

Fehr, E., S. Gächter, and G. Kirchsteiger (1997), Reciprocity as a contract enforcement device: Experimental evidence, *Econometrica: journal of the Econometric Society*, pp. 833–860.

Fehrler, S. (2010), *Charity as a Signal of Trustworthiness*, IZA.

Fischbacher, U. (2007), z-tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics*, *10*(2), 171–178.

Fréchette, G. R., and S. Yuksel (2013), Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination, *Available at SSRN 2225331*.

Frey, B. S., and S. Neckermann (2008), Awards, *Zeitschrift für Psychologie/Journal of Psychology*, *216*(4), 198–208.

Fudenberg, D., and E. Maskin (1986), The folk theorem in repeated games with discounting or with incomplete information, *Econometrica*, *54*(3), 533–554.

Fudenberg, D., A. Dreber, and D. G. Rand (2012), Slow to anger and fast to forgive: Cooperation in an uncertain world, *American Economic Review*, *102*, 720–749.

Gächter, S., and E. Fehr (1999), Collective action as a social exchange, *Journal of Economic Behavior & Organization*, *39*(4), 341–369.

Ganesan, S. (1994), Determinants of long-term orientation in buyer-seller relationships, *the Journal of Marketing*, pp. 1–19.

Gibbons, R. (1998), Incentives in organizations, *Tech. rep.*, National Bureau of Economic Research.

Gibbons, R. (2001), Trust in social structures: Hobbes and coase meet repeated games, *Trust in society*, pp. 332–353.

Gibbons, R. (2005), Incentives between firms (and within), *Management Science*, *51*(1), 2–17.

Gibbons, R., and R. Henderson (2012), Relational contracts and organizational capabilities, *Organization Science*, *23*(5), 1350–1364.

Glaeser, E., D. Laibson, J. Scheinkman, and C. Soutter (2000), Measuring trust, *The Quarterly Journal of Economics*, *115*(3), 811–846.

Gneezy, U., and A. Rustichini (2000), Pay enough or don't pay at all, *Quarterly journal of economics*, pp. 791–810.

Gray, J. V., A. V. Roth, and B. Tomlin (2009), The influence of cost and quality priorities on the propensity to outsource production*, *Decision Sciences*, *40*(4), 697–726.

Green, J. R., and N. Stokey (1982), A comparison of tournaments and contracts.

Grossman, S., and O. Hart (1986), The costs and benefits of ownership: A theory of vertical and lateral integration, *The Journal of Political Economy*, pp. 691–719.

Haitao Cui, T., J. S. Raju, and Z. J. Zhang (2007), Fairness and channel coordination, *Management Science*, *53*(8), 1303–1314.

Han, S.-L., D. T. Wilson, and S. P. Dant (1993), Buyer-supplier relationships today, *Industrial Marketing Management*, *22*(4), 331–338.

Handfield, R., and C. Bechtel (2002), The role of trust and relationship structure in improving supply chain responsiveness, *Industrial Marketing Management*, *31*(4), 367–382.

Harackiewicz, J. M. (1979), The effects of reward contingency and performance feedback on intrinsic motivation., *Journal of Personality and Social Psychology*, *37*(8), 1352.

Heinrich, T., and J. Brosig-Koch (2011), Reputation and mechanism choice in procurement auctions-an experiment, *Ruhr Economic Paper*, (254).

Helper, S., and R. Henderson (2014), Management practices, relational contracts, and the decline of general motors, *Tech. rep.*, National Bureau of Economic Research.

Ho, T., and J. Zhang (2008), Designing pricing contracts for boundedly rational customers: Does the framing of the fixed fee matter?, *Management Science*, *54*(4), 686–700.

Honhon, D., and K. B. Hyndman (2015), Flexibility and reputation in repeated prisoners' dilemma games, *Available at SSRN 2562727*.

Hyndman, K., and D. Honhon (2014), Flexibility in long-term relationships: An experimental study, *preparation for submission to Management Science*.

Hyndman, K., S. Kraiselburd, and N. Watson (2014), Coordination in games with strategic complementarities: An experiment on fixed vs. random matching, *Production and Operations Management*, *23*(2), 221–238.

Iyer, A. V., L. B. Schwarz, and S. A. Zenios (2005), A principal-agent model for product specification and production, *Management Science*, *51*(1), 106–119.

Johnson, S., J. McMillan, and C. Woodruff (2002), Courts and relational contracts, *Journal of Law, Economics, and Organization*, *18*(1), 221–277.

Joskow, P. (1985), Vertical integration and long-term contracts: The case of coal-burning electric generating plants, *JL Econ. & Org.*, *1*, 33.

Kalwani, M. U., and N. Narayandas (1995), Long-term manufacturer-supplier relationships: do they pay off for supplier firms?, *The Journal of Marketing*, pp. 1–16.

Kandori, M. (1992a), Repeated games played by overlapping generations of players, *The Review of Economic Studies*, *59*(1), 81–92.

Kandori, M. (1992b), Social norms and community enforcement, *The Review of Economic Studies*, *59*(1), 63–80.

Katok, E., and V. Pavlov (2013), Fairness in supply chain contracts: A laboratory study, *Journal of Operations Management*, *31*(3), 129–137.

Katok, E., and D. Wu (2009), Contracting in supply chains: A laboratory investigation, *Management Science*, *55*(12), 1953–1968.

Kaya, M., and Ö. Özer (2009), Quality risk in outsourcing: Noncontractible product quality and private quality cost information, *Naval Research Logistics (NRL)*, *56*(7), 669–685.

Kim, S.-H., and S. Netessine (2013), Collaborative cost reduction and component procurement under information asymmetry, *Management Science*, *59*(1), 189–206.

King-Casas, B., D. Tomlin, C. Anen, C. Camerer, S. Quartz, and P. Montague (2005), Getting to know you: reputation and trust in a two-person economic exchange, *Science*, *308*(5718), 78–83.

Klein, B., R. Crawford, and A. Alchian (1978), Vertical integration, appropriable rents, and the competitive contracting process, *Journal of law and economics*, *21*(2), 297–326.

Klier, T. H. (2006), The supplier industry in transitionthe new geography of auto production, *Chicago Fed Letter*, *229*.

Kocher, M. G., and M. Sutter (2007), Individual versus group behavior and the role of the decision making procedure in gift-exchange experiments, *Empirica*, *34*(1), 63–88.

Kosfeld, M., and S. Neckermann (2011), Getting more work for nothing? symbolic awards and worker performance, *American Economic Journal: Microeconomics*, pp. 86–99.

Kosfeld, M., M. Heinrichs, P. Zak, U. Fischbacher, and E. Fehr (2005), Oxytocin increases trust in humans, *Nature*, *435*(7042), 673–676.

Kreps, D. M. (1996), Corporate culture and economic theory, *Firms, Organizations and Contracts, Oxford University Press, Oxford*, pp. 221–275.

Kreps, D. M., and R. Wilson (1982), Reputation and imperfect information, *Journal of economic theory*, *27*(2), 253–279.

Kreps, D. M., P. Milgrom, J. Roberts, and R. Wilson (1982), Rational cooperation in the finitely repeated prisoners' dilemma, *Journal of Economic theory*, *27*(2), 245–252.

Kube, S., M. A. Maréchal, and C. Puppea (2012), The currency of reciprocity: Gift exchange in the workplace, *The American Economic Review*, *102*(4), 1644–1662.

Kugler, T., G. Bornstein, M. G. Kocher, and M. Sutter (2007), Trust between individuals and groups: Groups are less trusting than individuals but just as trustworthy, *Journal of Economic psychology*, *28*(6), 646–657.

Lacetera, N., and M. Macis (2010a), Do all material incentives for pro-social activities backfire? the response to cash and non-cash incentives for blood donations, *Journal of Economic Psychology*, *31*(4), 738–748.

Lacetera, N., and M. Macis (2010b), Social image concerns and prosocial behavior: Field evidence from a nonlinear incentive scheme, *Journal of Economic Behavior & Organization*, *76*(2), 225–237.

Lariviere, M., and E. Porteus (2001), Selling to the newsvendor: An analysis of price-only contracts, *Manufacturing & service operations management*, *3*(4), 293–305.

Larson, A. (1992), Network dyads in entrepreneurial settings: A study of the governance of exchange relationships, *Administrative science quarterly*, pp. 76–104.

Lazear, E. P., and S. Rosen (1979), Rank-order tournaments as optimum labor contracts, *Working Paper 401*, National Bureau of Economic Research, doi: 10.3386/w0401.

Li, C., and L. G. Debo (2009), Strategic dynamic sourcing from competing suppliers with transferable capacity investment, *Naval Research Logistics (NRL)*, *56*(6), 540–562.

Liker, J. K., and T. Y. Choi (2004), Building deep supplier relationships, *Harvard business review*, *82*(12).

Loch, C., and Y. Wu (2008), Social preferences and supply chain performance: An experimental study, *Management Science*, *54*(11), 1835–1849.

Lovejoy, W. S. (2010), Bargaining chains (long version), *Ross School of Business Paper*, (1146).

Masclet, D., C. Noussair, S. Tucker, and M.-C. Villeval (2003), Monetary and non-monetary punishment in the voluntary contributions mechanism, *American Economic Review*, pp. 366–380.

Masten, S. (1984), The organization of production: Evidence from the aerospace industry, *Journal of Law and Economics*, *27*(2), 403–417.

McMillan, J. (1990), Managing suppliers: Incentive systems in japanese and us industry, *California Management Review*, *32*(4), 38.

Miller, G. (2001), Why is trust necessary in organizations? the moral, *Trust in society*, p. 307.

Miller, G. J., and R. W. Smith (1993), *Managerial dilemmas: The political economy of hierarchy*, Cambridge University Press.

Monczka, R. M. (2000), *New product development: strategies for supplier integration*, ASQ Quality Press.

Monteverde, K., and D. Teece (1982), Supplier switching costs and vertical integration in the automobile industry, *The Bell Journal of Economics*, pp. 206–213.

Morgan, R., and S. Hunt (1994), The commitment-trust theory of relationship marketing, *the journal of marketing*, pp. 20–38.

Murnighan, J. K., and A. E. Roth (1983), Expecting continued play in prisoner's dilemma games a test of several models, *Journal of Conflict Resolution*, *27*(2), 279–300.

Nalebuff, B. J., and J. E. Stiglitz (1983), Prizes and incentives: towards a general theory of compensation and competition, *The Bell Journal of Economics*, pp. 21–43.

Nash Jr, J. F. (1950), The bargaining problem, *Econometrica: Journal of the Econometric Society*, pp. 155–162.

Newman, R. G. (1988), Single source qualification, *Journal of Supply Chain Management*, *24*(2), 10.

Noussair, C., and S. Tucker (2005), Combining monetary and social sanctions to promote cooperation, *Economic Inquiry*, *43*(3), 649–660.

Özer, Ö., Y. Zheng, and K. Chen (2011), Trust in forecast information sharing, *Management Science*, *57*(6), 1111–1137.

Özer, Ö., Y. Zheng, and Y. Ren (2014), Trust, trustworthiness, and information sharing in supply chains bridging china and the united states, *Management Science*, *60*(10), 2435–2460.

Pasternack, B. (1985), Optimal pricing and return policies for perishable commodities, *Marketing science*, *4*(2), 166–176.

Pasternack, B. A. (2008), Optimal pricing and return policies for perishable commodities, *Marketing science*, *27*(1), 133–140.

Pavlov, V., and E. Katok (2009), Fairness and coordination failures in supply chain contracts, *working paper*.

Peeters, R., and M. Vorsatz (2013), Immaterial rewards and sanctions in a voluntary contribution experiment, *Economic Inquiry*, *51*(2), 1442–1456.

Piboonrungorj, P., and S. Disney (2012), Supply chain collaboration, inter-firm trust and logistics performance: Evidence from the tourism sector, *Available at SSRN 2050703*.

Rabin, M. (1993), Incorporating fairness into game theory and economics, *The American Economic Review*, pp. 1281–1302.

Radner, R. (1985), Repeated principal-agent games with discounting, *Econometrica: Journal of the Econometric Society*, pp. 1173–1198.

Roth, A. E., and J. K. Murnighan (1978), Equilibrium behavior and repeated play of the prisoner's dilemma, *Journal of Mathematical psychology*, *17*(2), 189–198.

Rousseau, D. M., S. B. Sitkin, R. S. Burt, and C. Camerer (1998), Not so different after all: A cross-discipline view of trust, *Academy of management review*, *23*(3), 393–404.

Rubinstein, A. (1979), Equilibrium in supergames with the overtaking criterion, *Journal of Economic Theory*, *21*(1), 1–9.

Rudzki, R. A. (2004), The advantages of partnering well, *SUPPLY CHAIN MANAGEMENT REVIEW, V. 8, NO. 2 (MAR. 2004), P. 44-51: ILL*, *8*(2).

Scharlemann, J., C. Eckel, A. Kacelnik, and R. Wilson (2001), The value of a smile: Game theory with a human face, *Journal of Economic Psychology*, *22*(5), 617–640.

Schweitzer, M., and G. Cachon (2000), Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence, *Management Science*, *46*(3), 404–420.

Sherstyuk, K., N. Tarui, and T. Saijo (2013), Payment schemes in infinite-horizon experimental games, *Experimental Economics*, *16*(1), 125–153.

Silver, E., D. Pyke, R. Peterson, et al. (1998), *Inventory management and production planning and scheduling*, vol. 3, Wiley New York.

Spence, M. (1973), Job market signaling, *The quarterly journal of Economics*, *87*(3), 355–374.

Spengler, J. (1950), Vertical integration and antitrust policy, *The Journal of Political Economy*, *58*(4), 347–352.

Stallkamp, T. T. (2005), *SCORE!: A Better Way to Do Busine : Moving from Conflict to Collaboration*, Pearson Education.

Swinney, R., and S. Netessine (2009), Long-term contracts under the threat of supplier default, *Manufacturing & Service Operations Management*, *11*(1), 109–127.

Taylor, T. A., and E. L. Plambeck (2007), Supply chain relationships and contracts: The impact of repeated interaction on capacity investment and procurement, *Management science*, *53*(10), 1577–1593.

Tomlin, B. (2003), Capacity investments in supply chains: Sharing the gain rather than sharing the pain, *Manufacturing & Service Operations Management*, *5*(4), 317–333.

Treleven, M. (1987), Single sourcing: a management tool for the quality supplier, *Journal of Purchasing and Materials Management*, *23*(1), 19–24.

Williamson, O. (1971), The vertical integration of production: market failure considerations, *The American Economic Review*, pp. 112–123.

Williamson, O. (1975), Markets and hierarchies: analysis and antitrust implications: a study in the economics of internal organization, *University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship.*

Williamson, O. (1979), Transaction-cost economics: the governance of contractual relations, *Journal of law and economics*, *22*(2), 233–261.

Zak, P., and S. Knack (2001), Trust and growth, *The economic journal*, *111*(470), 295–321.

Zhu, K., R. Q. Zhang, and F. Tsung (2007), Pushing quality improvement along supply chains, *Management science*, *53*(3), 421–436.