

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

Ruth Beer

Kelley School of Business, Indiana University, Bloomington, IN 47405, ruthbeer@indiana.edu

Hyun-Soo Ahn

Ross School of Business, University of Michigan, Ann Arbor, MI 48104, hsahn@umich.edu

Stephen Leider

Ross School of Business, University of Michigan, Ann Arbor, MI 48104, leider@umich.edu

While innovation sharing between a buyer and a supplier can increase the efficiency and total profit of a supply chain, many suppliers are reluctant to do so. Sharing innovations leaves the supplier in a vulnerable position if the buyer exploits the information (e.g. by re-sharing the supplier's innovation with competing suppliers). In this paper, we examine conditions under which a collaborative relationship can arise in this situation, with a supplier voluntarily sharing an innovation and a buyer repaying that trust by sharing the surplus increase rather than seeking competing bids from other suppliers. First we show, both theoretically and experimentally, that decisions to collaborate are affected by the length of the relationship between the firms - longer relationships lead to higher collaboration and higher total profits. We additionally show that collaboration depends not just on the firm-level relationship length, but also on the long- or short-run focus of the employees within the firms that make decisions. We model the buyer as a dual decision maker, with long-run and/or short-run focused employees ("engineers" and "procurement managers") determining the buyer's actions. We characterize the equilibrium of this model and show that collaborative outcomes depend on the level of control the long-run employee has within the buyer. Our experimental results verify this intuition. Collaborative relationships occur more often when the engineer has more control. However, the supplier's decision to share an innovation depends primarily on the firm-level relationship length, while the buyer's decision to seek competition depends more on the relationship focus of the controlling employees. Consequently, buyers' profits increase with long-run firm relationships (for any decision maker), while suppliers' profits only significantly increase with a long-run decision maker. Finally, while suppliers and engineers should theoretically ignore the actions of the previous procurement managers, we find that both suppliers' and engineers' actions are correlated with the previous procurement managers' decision.

1. Introduction

Manufacturers often benefit from innovations and process improvements discovered by their suppliers. This happens particularly in industries where suppliers are involved in research and development (R&D) and product design. In the automotive industry, a substantial share of cost reductions come

from part suppliers¹. For example, General Motors' suppliers developed hinges that did not need welding, which resulted in significant cost savings for GM (Klier 2006). Similarly, Chrysler had major cost savings when a supplier, Becker Manufacturing, eliminated excess fasteners by developing molded hooks in their interior trim panels so that door panels could be directly fastened to the frame².

Valuable innovations often involve process improvements and benefits for both the supplier and the manufacturer. For example, the automotive industry has been transitioning from solvent based paint to waterborne paint, which is less toxic and easier to dispose and clean up. Implementing this technology requires significant changes to the manufacturing process - for example, painters have to be retrained to paint more evenly, new taping techniques are needed to prevent bleeding, new equipment needs to be installed to blow large volumes of clear air for drying, etc. Both the manufacturer and its suppliers transformed their paint booths during the transition, and sharing painting process improvements could benefit both parties – resulting in cost reductions or allowing better paint matching between parts.

For some innovations or process improvements (which are usually not subject to patents) suppliers are often concerned that the buyer may pass the innovation on to other suppliers to increase competition and lower future prices. Sharing the innovation will then make the supplier vulnerable and, ultimately, take away the supplier's competitive advantage. Historically, these acts of untrustworthiness have been quite frequent among U.S. car manufacturers (McMillan 1990). U.S. automakers have commonly used procurement strategies primarily focused on cost reduction even at the expense of destroying supplier's trust (Burt 1989, Liker and Choi 2004)³. This focus on pushing for cost reductions is often associated with short-term supplier relationships and seeking competitive bids frequently (e.g. switching suppliers after each sale period)⁴. Recent initiatives, such as Ford's "Aligned Business Framework" and GM's "Strategic Supplier Engagement" program, have focused on building longer-run strategic partnerships with suppliers and encouraging innovation sharing.

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>.

³ In this regard, Helper and Henderson (2014) and Liker and Choi (2004) provide a comprehensive review of the cultural differences between Japanese and American automakers between 1980 and 2009.

⁴ McMillan (1990) reports that contracts of three to five years are generally considered long-term in the automotive industry. At times, U.S. manufacturers have tried to forge longer supplier relationships, 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 (1996) reports on Chrysler's efforts in the 1990s to increase their commitments to their suppliers, which increased the average contract length from 2.1 years to 4.4 years.

While this has led to recent improvements in measures of the quality of the working relationship between American automakers and their suppliers, their ratings are still quite poor⁵.

The prospect of a long term relationship can make both suppliers and manufacturers more likely to collaborate. Relational contracts - defined by Gibbons and Henderson (2012) as “informal agreements enforced by the shadow of the future” - can provide enough incentives for collaboration, as poor behavior sacrifices future gains⁶. In our setting, this means that the manufacturer has an incentive to keep the supplier’s trust as long as his benefit from future innovation sharing exceeds the short-run gain from betraying the supplier’s trust by bringing in a competing supplier. In turn, the supplier is incentivized to maintain the collaborative relationships by sharing innovations. However, even when the firms have sufficient incentives to collaborate, the individuals making decisions for the firms may not. Conversations conducted with GM’s suppliers suggest that the nature of a buyer-supplier relationship heavily depends on which employees within the firm manages the relationship. For some divisions within GM this responsibility is primarily with the procurement managers, while in others engineers have extensive control over supplier relations⁷. Procurement managers are often evaluated by performance metrics that focus on short-run (immediate) cost savings. Thus, their incentives are driven by these KPIs. This is further exacerbated by the fact that, in many organizations, procurement managers are rotated through an organization to source different parts or negotiate with different suppliers. On the other hand, engineers’ performance depends on quality and design, both of which are intrinsically long-term oriented objectives. In addition, since they have specific technical expertise, engineers are less likely to be rotated and commonly specialize in a certain auto part. A supplier’s trust in the buying firm depends, then, on which employee manages the relationship. This indicates that, if a buyer wants to build a long-term and collaborative relationship with a supplier, the buyer needs to be careful in assigning the roles and responsibility for managing the relationship.

The importance of trust and trustworthiness in buyer-supplier relations is well established in the operations management literature (see Özer and Zheng 2016 for an extensive review). For example,

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 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.

⁷ 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).

suppliers benefit when they can rely on a trustworthy buyer's demand forecast report to build capacity (Özer et al. 2011, Özer et al. 2014, Spiliotopoulou et al. 2015). Similarly, buyers benefit when a trustworthy supplier provides assistance about a product (Özer et al. 2016) or exerts effort towards a product's non-contractible quality (Beer et al. 2017). In most of these cases, when trust and trustworthiness arise, more collaborative relationships develop which are mutually beneficial. In our setting, a trusting supplier shares the innovation with a buyer, and a trustworthy buyer splits the surplus with the supplier rather than re-share the innovation to increase supplier competition. In order to study this situation, we analyze a theoretical model and identify conditions supporting a collaborative equilibrium. We then develop hypotheses and conduct a laboratory experiment to examine how the allocation of decision rights within the buying firm affects firms' actions. We answer the following research questions: (1) Does the length of a relationship between firms affect the likelihood of collaborative outcomes? (2) Within a long-run firm-level relationship, does who has the decision-making right (e.g., a single individual vs. a team, a short-sighted employee or a long-run focused employee) affect the buyer's and supplier's strategies and the likelihood of a collaborative relationship?

We find that a long-run relationship matters, both at the firm and individual level. When buyers are represented by a single decision maker, suppliers are more likely to share innovations and buyers are more likely to be trustworthy when they have a long-run relationship than a one-shot interaction. When buyers are represented by a team (an engineer and a procurement manager), who controls the relationship further affects the firms' strategies and relationship outcomes. The supplier's decision is primarily driven by the firm-level relationship characteristics - suppliers share innovations more often than in the pure one-shot interaction even if a procurement manager is in charge. Shifting control from the procurement manager to the engineer does not significantly change the supplier's decision. Buying firms, however, are significantly more likely to be trustworthy and not seek a competing supplier when the engineer is in charge (compared to procurement manager control and to the one-shot interaction). As a result, while buyers earn higher profits in a long-run firm-level relationship than in the one-shot interaction (with any form of decision control), suppliers only earn significantly higher profits when a long-run decision maker is in control. Our experiment provides some additional interesting insights. First, we find that a buyer's action when no innovation is shared also matters. In our setup, this may happen either when the supplier did not have an innovation or when the innovation happened but the supplier did not share it. We observe that some buyers are skeptical in this case and choose to open up competition right away while others wait until the end of the round to have confirmation that the innovation happened before punishing the supplier in a following round. Relationships where the buyer is skeptical lead to significantly lower total profits, as it strongly hurts the supplier's profit without making the buyer better off.

Unsurprisingly, suppliers are for the most part very unforgiving when the buyer takes the skeptical action. Second, we find that when there is uncertainty about which employee’s recommendation will be implemented –decision rights are allocated randomly– this randomness results in both players being significantly less trustworthy. Finally, we observe that employees may be influenced by their peers’ recommendations beyond their own monetary incentives.

2. Literature Review

There is a broad literature in operations management studying collaboration in buyer-supplier relations. Empirical papers show that cooperation between firms in a supply chain 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 by automakers resulted in faster learning and fewer defects by suppliers. Stallkamp (2005) analyzes several forms of collaboration: strategy, communication, information, and responsibility sharing. They find that strategic collaboration yields substantial cost and quality improvements. Firms’ organizational-level decisions may play a role in supply chain collaboration, as shown in Brinkhoff et al. (2015). They provide empirical evidence that trust is important for supply chain projects to be successful. Özer and Zheng (2016) provide an extensive study of when, how, and why trust and trustworthiness can arise to support collaboration between supply chain partners. They emphasize the importance of the market environment: for example, in forecast information sharing, the level of trust and trustworthiness that develops is affected by investment costs and demand volatility (Özer et al. 2011), the managers’ country of origin (Özer et al. 2014), and inventory competition among the managers (Spiliotopoulou et al. 2015). The settings in our paper vary along two important market characteristics: length of the relationship between the firms and allocation of decision rights within the buying firm.

Our paper focuses on collaboration via 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). A number of papers in the operations management literature analyze the problem of providing incentives to invest in cost reduction in a supply chain. Kim and Netessine (2013) study collaborative effort by the manufacturer and supplier 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 over the life cycle of the product under different procurement approaches, and consider gradual investment in process improvement (e.g. 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. Our experimental results show that in an

innovation-sharing setting, the allocation of decision rights to employees with different incentives is important in determining the level of trust between the firms, and both firms' willingness to collaborate.

During the late eighties and early nineties arguments in favor of procuring from a reduced number of suppliers and preserving long-term supplier relations became popular. 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.

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. This method of representing an infinitely repeated game has been used extensively in experimental economics. Roth and Murnighan (1978) and Murnighan and Roth (1983) were the firsts to induce infinitely repeated games using randomly terminated games, where the continuation probability is equated to the discount factor. Since then, indefinitely repeated games have been used to understand the evolution of cooperation in a prisoner's 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 directly 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). 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 prisoner's 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.

On a broader level, we contribute to the literature on behavior in supply chain management. Several papers have studied the effects of social preferences and decision biases on supply chain contracting (Bolton and Katok 2008, Katok and Wu 2009, Ho and Zhang 2008, Becker-Peth et al. 2013). Social preferences, such as fairness concerns, trust, and trustworthiness, play an important role in supply chain performance (Cui et al. 2007, Loch and Wu 2008, Özer et al. 2011, Katok and Pavlov 2013, Özer et al. 2014, Spiliotopoulou et al. 2015). A few experimental papers in operations management study how buyer-supplier relationships are affected by relationship length. Loch and Wu (2008) find that inducing a positive relationship before the game leads to more collaborative actions by suppliers and buyers, which persist over many rounds of the game. Özer et al. (2011) find that repeated interactions enhance trust and trustworthiness in forecast information sharing, resulting in lower forecast inflations, higher capacity investment, and higher supply chain efficiency. Beer et al. (2017) show that when suppliers can signal trustworthiness by making an upfront buyer-specific investment, more collaborative relationships arise. In that setting, repeated interactions strengthen the impact of signaling investments, leading to higher profits and efficiency. Davis and Hyndman (2016) show that relational incentives (i.e. a long term relationship where there is a threat of punishment) lead to increased quality and supply chain efficiency. Hyndman et al. (2014) conduct an experimental study that directly compares short run and long run incentives in a related supply chain setting. They study a setup where two firms simultaneously invest in capacity, and sales are the minimum of the two chosen capacities and realized demand. In their setting firms have private information about demand, and need to coordinate on the optimal investment level. They find that while fixed pairs have higher alignment on average than pairs that are randomly re-matched after every round, 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. 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. Our paper is different in both setting and research focus. First, our stage game more closely resembles a trust game (innovation sharing) than a coordination game (capacity alignment). Second, we introduce joint decision making within the buyer. Thirdly, our focus is on the firms' allocation of decision rights to procurement managers and engineers. Procurement managers have random rematching after every round, while engineers have fixed matching as long as the relationship between the firms lasts. With this setup, we capture the different incentives the employees in the buyer firm face, beyond the firm-level relationship length.

3. Model

We examine a supply chain consisting of a buyer (he) and a supplier (she) that are engaged in a multi-period relationship. In each period, the supplier may have discovered a new innovative idea and, if so, she needs to decide whether to share this with a buyer or not. The buyer needs to decide to collaborate with the supplier (offer a generous price) or to make the supplier compete against another supplier (to lower the price). We consider several scenarios varying in (i) duration of a relationship, and (ii) who makes a decision for the buyer and how the decision is made. 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. Building on these two benchmark cases, 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. We consider several different settings by varying the decision making process (in terms of who has the decision right).

3.1. Setup

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, $i \in \{m = \text{manufacturer}, s = \text{supplier}, a = \text{alternative supplier}\}$ ⁸. 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

⁸ 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.

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 a 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 a single-stage game, the supplier may have a new innovation which can lower the supplier's cost from C_{s1} to C_{s2} . We assume that this innovation occurs with probability π , which is exogenously determined. For instance, in the waterborne paint example discussed in Introduction, the supplier's innovation was a change to the blower set up that led to faster paint drying and lower unit costs. The supplier can share the innovation with the manufacturer (and voluntarily reduce the unit cost). If the supplier shares this innovation, the manufacturer can also implement the same technology in his own painting booths which reduces the manufacturer's production cost to C_{m2} , $C_{m2} \leq C_{m1}$. Alternatively, the supplier can decide to not to share the information and just increase the unit-product margin.

After the supplier decides whether to share the innovation with the manufacturer, the manufacturer can choose to solicit bids from a new supplier (we call this decision "to compete") or to single source ("not to compete"). We assume that the alternative supplier has the initial 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 and bring in the alternative supplier, then the production cost is reduced from C_{a1} to C_{a2} , with $C_{a2} < C_{a1}$ and $C_{a2} = C_{s2}$, essentially taking away the competitive advantage of the supplier who had the innovation.

After the manufacturer chooses whether to compete or not, the supplier and the manufacturer negotiate the terms of trade. As a result of this negotiation, the wholesale price, w^* , and retail price, p^* , are set to maximize the surplus of the supply chain. We assume the surplus is split between the supplier and the manufacturer according to Nash bargaining (Nash 1950): the manufacturer earns a fraction α , $\alpha \in [0, 1]$ of the surplus and the supplier earns a fraction $(1 - \alpha)$ of the surplus⁹.

In the case where the manufacturer chooses not to compete (case of *bilateral bargaining*), the Nash bargaining solution predicts equal splits of the surplus ($\alpha^* = \frac{1}{2}$) and the manufacturer's and supplier's profits are given by ¹⁰:

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

⁹ 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).

¹⁰ We assume that the parameters are such that transacting is efficient, that is, $p - C_m \geq C_s$ and at $w^* = (1 - \alpha)(p - C_m) + \alpha C_s$, both firms choose to transact.

In the case where the manufacturer chooses to compete, the Nash bargaining solution predicts $\alpha^* = \frac{1}{2} + \frac{(p-C_a-C_m)}{2(p-C_s-C_m)}$. This instance is referred to in the literature as *bargaining with supplier competition* (Lovejoy 2010). In this case, the original supplier is still selected but now the manufacturer's profit is:

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

and the supplier's profit is:

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

Note first that this requires $p > C_a + C_m$ and $p > C_s + C_m$. Second, if the innovation does not happen, or if it happens and the supplier does not share, $C_a > C_s$ and therefore the manufacturer's share of surplus, α^* , is greater than $\frac{1}{2}$. In the case where the supplier shares and the manufacturer competes, we have $C_s = C_a$ and the manufacturer takes all the surplus ($\alpha^* = 1$).

The detailed calculations are presented in sections 9.1 and 9.2 in the Appendix.

3.2. Numerical Example

As Figure 1 shows, the game has six possible outcomes: If the innovation 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). We analyze the equilibrium and draw hypotheses from a canonical example whose parameter values and payoffs are carefully chosen to facilitate our resulting lab experiments. The parameters used are presented in the Table 8 in the Appendix and the firms' payoffs resulting from these parameters are presented in Figure 1. Section 9.3 in the Appendix explains in detail how the payoffs in Figure 1 are derived given the chosen values for parameters C_{s1} , C_{s2} , C_{m1} , C_{m2} , C_{a1} , C_{a2} , and demand parameters, a and b .

With these payoffs, if the innovation occurs, the total surplus increases relative to the case where the innovation does not occur. In addition, if the supplier shares the innovation with the manufacturer, the total surplus (the size of a pie that can be shared between the two parties) is the largest. However, sharing the innovation makes the supplier more vulnerable to competition. Specifically, the minimum possible payoff from not sharing is 7 and from sharing is zero. The manufacturer's decision does not affect the total surplus 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 firms benefit from their counter part's action. Since the innovation occurs in each period with probability π , when choosing his action the manufacturer cannot distinguish between the case where the innovation has occurred but the supplier decided to




Figure 1 Stage Game

not share and the case where the innovation has not occurred in the first place. This is captured in the information set of a manufacturer (from now onwards, we will denote this decision node as the manufacturer’s “ambiguous node”). When the outcomes are realized, the manufacturer can learn whether an innovation had occurred, and therefore whether the supplier had shared. We make this simplification so that a manufacturer can condition his action directly on the supplier’s action in the previous round. Without this simplification, the manufacturer may implement a more elaborate review strategy (Radner 1985) by which he conditions his action on his probabilistic assessment of the supplier’s strategy after observing several rounds of play. This exact same information structure is reproduced in the laboratory experiment¹¹.

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

¹¹ Note that the game in Figure 1 resembles the widely studied trust game (Kreps 1990) with two differences. First, in our setup, the supplier’s decision to trust the manufacturer 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.

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.

3.2.1. Single-Period Game For the single-period game we solve by backward induction, starting with the manufacturer’s strategy. The manufacturer’s profit from the action “compete” is given by Equation 1 and from “not compete” is given by Equation 2. Since $p > C_a + C_m$, then $a - b(C_m + 2C_a - C_s) > 0$ which implies that the manufacturer’s profit is always higher when he chooses to compete (than not to compete). Given that the manufacturer chooses “compete”, the supplier’s profit is given by Equation 3. Rolling back to the supplier’s strategy, if she chooses to share the innovation, then $C_{a2} = C_{s2}$ and the supplier earns zero profit, if she does not share, then $C_{a1} > C_{s2}$ 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 1 are INC if the innovation occurs and NC if the innovation does not occur.

3.2.2. Repeated Interactions We now consider the infinitely repeated play of the stage game depicted in Figure 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 (game with random end). This interpretation is commonly used in the experimental economics literature, where it was first introduced by Murnighan and Roth (1983) and Roth and Murnighan (1978)¹². We resort to this interpretation later on in the Experimental Design section.

Consider the six different possible outcomes of the stage game presented in Figure 1 and let $ISC_i, ISN_i, INC_i, INN_i, NC_i, NN_i$ be player i ’s payoffs, $i \in \{m = \text{manufacturer}, s = \text{supplier}\}$, from each possible outcome. Additionally, recall that at the end of the stage game, the manufacturer learns whether the innovation occurred and whether the supplier shared. Therefore, the manufacturer can condition his strategies on the past sharing decisions of the supplier. Similarly, the supplier can condition her strategies on the manufacturer’s past decisions to compete.

We want to characterize collaborative equilibria of the infinitely repeated game – that is equilibria where the supplier shares the innovation, and the manufacturer does not compete when the innovation is shared. We will focus on trigger strategies, where both players will choose the collaborative

¹² Recent experimental work by Fréchet and Yuksel (2013) verify 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.

action until either of them “defects” (fails to share an innovation, or competes when the innovation was shared). Recall that the full information on outcomes is revealed at the end of every stage game, so if either player defects the punishment state will commence in the next period and will continue for the rest of the game. However, within collaborative trigger strategies there is some flexibility during the stage game in how the manufacturer should handle the ambiguous decision node where no innovation was shared, since this could arise either from the supplier defecting by not sharing an innovation that occurred or from the supplier not having an innovation. We describe below two strategies for the manufacturer: the “skeptical” and the “non-skeptical” strategy. The non-skeptical strategy chooses not to compete during the stage game, and waits for the end of the period for confirmation of the supplier’s possible defection. By contrast, the skeptical strategy chooses to compete during the stage game, but tries to return to the collaborative state if the supplier was innocent (i.e. had no innovation). We also describe two corresponding strategies for the supplier: “forgiving” and “non-forgiving”. A non-forgiving supplier enters the punishment state if the manufacturer chooses to compete when there was no innovation, while a forgiving supplier is willing to stay in the collaborative state.

In preparation for Proposition 1, let us define the following:

[1.] Manufacturer’s Collaborative-Skeptical Trigger Strategy: The manufacturer begins in the collaborative state. The manufacturer enters the punishment state if in the previous round (a) the manufacturer chose to compete if the supplier shared an innovation, (b) the supplier chose not to share an innovation, or (c) the supplier entered the punishment state for any other reason. Once the manufacturer enters the punishment state he stays there for the rest of the game.

- Collaborative state: The manufacturer chooses not to compete if the supplier shares the innovation, and chooses to compete if there was no shared innovation.

- Punishment state: The manufacturer chooses to compete at any decision node.

[2.] Manufacturer’s Collaborative-Non-Skeptical Trigger Strategy: The manufacturer begins in the collaborative state. The manufacturer enters the punishment state if in the previous round (a) the manufacturer chose to compete if the supplier shared an innovation, (b) the supplier chose not to share an innovation, or (c) the supplier entered the punishment state for any other reason. Once the manufacturer enters the punishment state he stays there for the rest of the game.

- Collaborative state: The manufacturer chooses not to compete at any decision node.

- Punishment state: The manufacturer chooses to compete at any decision node.

[3.] Supplier’s Collaborative-Forgiving Trigger Strategy: The supplier begins in the collaborative state. The supplier enters the punishment state if in the previous round (a) the manufacturer chose to compete if the supplier shared an innovation, (b) the supplier chose not to share an innovation, or (c) the manufacturer entered the punishment state for any other reason. Once the supplier enters the punishment state she stays there for the rest of the game.

- Collaborative state: The supplier shares the innovation.
- Punishment state: The supplier does not share the innovation.

[4.] Supplier's Collaborative-Non-Forgiving Trigger Strategy: The supplier begins in the collaborative state. The supplier enters the punishment state if in the previous round (a) the manufacturer chose to compete, (b) the supplier chose not to share an innovation, or (c) the manufacturer entered the punishment state for any other reason. Once the supplier enters the punishment state she stays there for the rest of the game.

- Collaborative state: The supplier shares the innovation.
- Punishment state: The supplier does not share the innovation.

[5.] The thresholds:

$$\hat{\delta}_1 = \frac{ISC_m - ISN_m}{ISC_m - \pi INC_m - (1-\pi)ISN_m}, \text{ and}$$

$$\hat{\delta}_2 = \frac{ISC_m - ISN_m}{ISC_m - (1-\pi)ISN_m + (1-\pi)NN_m - \pi INC_m - (1-\pi)NC_m}.$$

PROPOSITION 1. *Collaborative Equilibria - Firm-as-Single-Employee case:*

[Equilibrium E1:] *If $\delta \geq \hat{\delta}_1$, there exists a subgame perfect Nash Equilibrium where the supplier plays the Collaborative-Forgiving strategy, and the manufacturer plays the Collaborative-Skeptical strategy.*

[Equilibrium E2:] *If $\delta \geq \hat{\delta}_2$, there exists a subgame perfect Nash Equilibrium where the supplier plays the Collaborative-Non-Forgiving strategy, and the manufacturer plays the Collaborative-Non-Skeptical strategy.*

The detailed proofs for equilibria *E1* and *E2* are presented in sections 9.4.1 and 9.4.2 in the Appendix. The proofs show that, with the payoffs in Figure 1 and $\pi = 0.75$ (as will later be used in the experiment), $\delta_1 = 0.69$ and $\delta_2 = 0.70$ guarantee that each equilibrium can be sustained. Based on the two interpretations of δ , the conditions mean 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 for both equilibria to arise¹³. Note that the Folk theorem for infinitely repeated games implies that many strategies can support equilibria with collaborative outcomes¹⁴. 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 least risky for suppliers when matched with manufacturers playing always compete, which is a very common strategy based on previous experimental evidence (Dal Bó and Fréchette 2013). Also, we focus only on pure strategies that lead to an equilibrium with high sharing rates. In mixed strategies, the buyer

¹³ In the experiment we will use $\delta = 0.75$ (a continuation probability of 0.75) to allow for both equilibria to arise

¹⁴ See 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).

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 buyer's actions across several periods.

We focus now on the setting where the firms have a long term relationship and assume that the manufacturer's decision is made by a procurement manager and an engineer. The procurement manager works for the firm for only one period (or equivalently is assigned to this supplier for one period, and rotates to another position in the firm in the next period). The engineer works for the firm (and is assigned to this supplier) throughout the infinite game. We further assume that both employees make recommendations for what the manufacturer should do and that their compensation is the manufacturer's profit. The procurement manager, being a short-run player, only cares about the current period's profits. The engineer, however, is a long-run player that cares about total profits during the whole buyer-supplier relationship.

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 always recommends to 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 supplier and the engineer resembles an infinitely repeated game. Thus, when the engineer's recommendation is always implemented, Proposition 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. In the ambiguous node, the engineer can choose either "compete" (as in $E1$) or "not compete" (as in $E2$).

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' recommendations¹⁵. In this case, trigger strategies analogous to those

¹⁵ We make the assumption that suppliers can observe both employees' recommendations for simplicity. This could represent either the supplier directly observing the buyer's decision-making, or the engineer being able to credibly verify his recommendation. 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 the engineer's strategy probabilistically after observing several rounds of play.

in Proposition 1 can sustain the collaborative outcome, Share - Not compete. The result is presented in the next proposition, for which we define $\hat{\delta}_3 = \frac{ISC_m - ISN_m}{(1 + \frac{\pi}{2})ISC_m - \pi INC_m - (1 - \frac{\pi}{2})ISN_m}$ and $\hat{\delta}_4 = \frac{ISC_m - ISN_m}{[(\frac{\pi}{2} - 1)ISN_m - (\frac{\pi}{2} + 1)ISC_m - \pi INC_m + \frac{(1 - \pi)}{2}NN_m + \frac{(\pi - 1)}{2}NC_m]}$.

PROPOSITION 2. *Collaborative Equilibria - Buyer-as-Two-Employees case:*

[Equilibrium E1':] *If $\delta \geq \hat{\delta}_3$, there exists a subgame perfect Nash Equilibrium where the supplier plays the Collaborative-Forgiving strategy conditional on the engineer's recommendation. The engineer plays the Collaborative-Skeptical strategy conditional on the supplier's action, and the procurement manager plays the non-collaborative strategy of the single shot game.*

[Equilibrium E2':] *If $\delta \geq \hat{\delta}_4$, there exists a subgame perfect Nash Equilibrium where the supplier plays the Collaborative-Non-Forgiving strategy conditional on the engineer's recommendation. The engineer plays the Collaborative-Non-Skeptical strategy conditional on the supplier's action, and the procurement manager plays the non-collaborative strategy of the single shot game.*

The proofs of equilibria E1' and E2' are analogous to those of E1 and E2 in Proposition 1 and are relegated to the Appendix (sections 9.5.1 and 9.5.2). The proofs show that, with the payoffs in Figure 1 and $\pi = 0.75$, $\delta_3 = 0.55$ and $\delta_4 = 0.55$ guarantee that each equilibrium can be sustained.

4. Experimental Design

The sequence of events and payoffs in each round of the experiment follow the stage game presented in Figure 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. First, suppliers are asked whether, if the innovation has occurred, they want to share it with the buyer. Second, buyers are asked whether they want to compete or not in case the supplier shared the innovation, and in case the innovation did not happen or the supplier did not share it¹⁶. After suppliers and buyers have made their decisions, the computer randomly determines whether the innovation occurs (it occurs with probability $\pi = 0.75$) and implements the chosen actions. At the end of the stage game all the subjects in the group are informed whether the innovation occurred, the supplier's and buyer's decision (as well as the individual recommendations of the procurement manager and engineer in the corresponding treatments), and their payoffs.

We design two different experimental settings. The first set-up involves two subjects – one acting as a buyer and the other acting as a supplier. The second set-up is similar except that the buying

¹⁶ We use the strategy method so we can fully understand what subjects' strategies are, even in treatments where the manipulation makes a certain scenario unlikely to happen. For example, we are interested in whether the buyer would compete or not in case the supplier shared the innovation, even in a treatment where suppliers seldom share. The strategy method has been extensively used in the experimental economics literature to elicit full strategies. In a literature survey, Brandts and Charness (2011) analyze twenty-nine comparisons between the strategy method and the direct-response method and find that in no case a treatment effect found with the strategy method was not observed with the direct-response method.

firm consists of two subjects – the first subject acting as a procurement manager and the other subject acting as an engineer. In each set-up, we conduct experiments under several treatments.

In the first set-up, denoted Buyer-as-Single-Employee, subjects are assigned to a role (either a supplier or a buyer), which they keep throughout the experiment. In the first treatment, “Short Run” (SR), the buyer-supplier relationship lasts only one round (a single transaction). We induce this by randomly re-matching buyers and suppliers after each round. In the second treatment, the buyer-supplier relationship is long term (LR). To capture this, the buyer and the supplier will play a repeated game with a random stopping time. Previous experiments (Fréchette and Yuksel 2013) have verified that this is an effective implementation of infinite games, as random termination yields equivalent behavior to payoff discounting. After each round, with probability δ , the relationship continues to the next round and the buyer and the supplier engage in another stage game. On the other hand, with probability $1 - \delta$, the relationship terminates. We use a random number generator to simulate the random stopping time. Once the relationships end, buyers and suppliers will be randomly re-matched again to begin a new relationship. The length of a relationship is equal to the number of *rounds* where the same buyer and supplier engage in stage games. To analyze how the behavior of a subject changes during a session, we use the term *period*, to represent the total number of rounds that a subject has played until now. Thus, $\text{period} = 10$ means that a subject has played a stage game 10 times. In both treatments, the subjects know the continuation and re-matching rules.

In the second set-up, denoted Buyer-as-Two-Employees, there are two subjects working for the buyer – one procurement manager and one engineer. As before, one subject will be assigned to play a role of a supplier. At the firm level the “supplier” and “buyer” have a long term relationship. Subjects in the supplier and engineer roles play together as long as the relationship between the two firms lasts. However, subjects in the procurement manager role are rotated between buyer-supplier pairings each round (representing the procurement manager rotating across departments). We implement this as follows: After each round, a random number is drawn to determine if the relationship between the firms continues. If the relationship continues, suppliers and engineers remain matched for the following round and procurement managers are 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. The stage game in the Buyer-as-Two-Employees set-up is as in the Buyer-as-Single-Employee set-up, except for the second stage (buyer’s decision). In the second 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 the innovation, and if the innovation did

not happen or the supplier did not share it. 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, procurement managers are informed of the last round history in the relationship they have joined before they make their recommendations. All subjects know that this information is provided to procurement managers.

We conduct three different treatments in the Buyer-as-Two-Employees set-up to examine different allocations of decision rights between the engineer and procurement manager: In the procurement manager treatment (denoted PM treatment), the procurement manager's recommendation is always implemented. The opposite happens in the engineer treatment (denoted Eng treatment), where the engineer's recommendation is always implemented. In the joint decision treatment (denoted 50 – 50 treatment), both employees' jointly determine the buyer's action, with the computer randomly picking one recommendation (with equal probability) to implement if they disagree. All subjects are informed that the allocation of decision rights is the one corresponding to the treatment they are in. 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. 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 of 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 with decision authority at some point during the session. As before, the subjects know the grouping, continuation, and re-matching rules.

The experiment consists of five treatments in total, SR, LR, PM, Eng, and 50 – 50, and follows a between-subjects design (each subject is exposed to one treatment). To ensure the subjects' understanding of the game, three examples are presented in the instructions, and the table with payoffs (Figure 3 in the appendix) is shown to participants throughout the experiment. In particular, to avoid biases relative to the continuation probabilities, in the PM, 50 – 50, Eng, and LR treatments it was made explicit that, after each round, the probability that the relationship will continue for another round remains exactly the same. To avoid reputation effects, participants only learn the outcomes and payoffs of their own relationships. In addition, since there is a minimum of four

relationships playing simultaneously in any given session, 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). For the probability of innovation we set $\pi = 0.75$, which allows us to get a high frequency of the interesting outcome where the innovation happens. This relatively high frequency of innovations captures for example the occurrence of small process improvements, rather than big events such as disruptive new technologies (which in reality happen less frequently). For the continuation probability, we used $\delta = 0.75$. A 0.75 continuation probability implies average relationship lengths of four years¹⁷, which is consistent for example with the automotive industry (McMillan 1990). This value guarantees that the collaborative outcome is an equilibrium of the game in the LR, Eng, and 50-50 treatments. Propositions 1 and 2 show that with the payoffs in Figure 1, cooperation can be supported as part of a subgame perfect equilibrium for values of continuation probability greater than 0.69 (for the LR and Eng settings) and 0.55 (for the 50 – 50 setting). Choosing the continuation probability, $\delta = 0.75$ provides an additional slack to ensure that equilibrium outcomes emerge.

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 University of Michigan. A total of 372 undergraduates participated in four sessions of each of the Buyer-as-Single-Employee treatments and six sessions of each of the Buyer-as-Two-Employees treatments. The maximum number of subjects per session was 18 and the minimum was 10 for the Buyer-as-Single-Employee treatments and 12 for the Buyer-as-Two-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 payment¹⁸. 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 earned¹⁹.

¹⁷ This is under the assumption that firms make supplier selection decisions on an annual basis, as is commonly the case in the automotive industry.

¹⁸ We use all observations up to period 30, which is the latest period that was reached in every session. Our main results do not change significantly if we use the observations from all periods.

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.

5. Hypotheses

We derive the following experimental hypotheses from our analysis (mainly propositions 1 and 2). The first hypothesis is derived from the equilibrium outcomes of the Buyer-as-Single-Employee setup. While the non-collaborative outcome is the only equilibrium that can be supported in the short-term relationship, Proposition 1 implies that a collaborative outcome emerges as an equilibrium in the long-term relationship. 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. [Buyer-as-Single-Employee Treatments] *Firms collaborate less in the SR treatment than in the LR treatment. Specifically, compared to the LR treatment, the SR treatment results should show that*

1.a - the supplier chooses to share less frequently,

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

1.c - the frequency of collaborative outcomes (both firms collaborate simultaneously) is lower.

The next hypothesis is for the Buyer-as-Two-Employees treatments. Since the procurement manager is part of the buyer-supplier pairing 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 results of the Eng treatment should be similar to the results of the LR treatment. Finally, in the 50-50 treatment, since the final decision is randomly picked, the buyer will follow the procurement manager's decision and the engineer's decision 50% of the time, respectively. The theory (Proposition 2) prescribes an equilibrium where the supplier always shares, the engineer recommends not to 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. [Buyer-as-Two-Employees Treatments] *In the Firms-as-Two-Employees treatments, collaboration is in between the SR and LR benchmarks:*

2.a - the PM treatment obtains the same outcomes as the SR treatment,

2.b - the Eng treatment obtains the same outcomes as the 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.

We expect that the two treatments under the Buyer-as-Single-Employee setting will serve as benchmarks: We expect (1) the SR treatment has the lowest level of collaboration and (5) the LR treatment the highest level. In the Buyer-as-Two-Employees treatments, our theoretical results

stipulate that the consideration for a long-term relationship will increase when the treatment changes from (2) the PM to (3) the 50-50 to (4) the Eng treatment. Thus, we expect to see that the more collaborative outcomes will emerge as the treatment changes from (1) to (5). 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.

Notice that the procurement manager engages in a relationship for just one round before being rotated 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.*

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 Buyer-as-Two-Employees treatments and analyze the interplay between engineers and procurement managers.

6.1. Descriptive Results

Our hypotheses in the previous section imply that both suppliers and buyers are more likely to choose collaborative actions (hence the stage-game results in a collaborative outcome) as the prospect of a long-term relationship becomes more explicit. Specifically, we expect that the supplier will be least collaborative in (1) SR and (2) PM treatments and most collaborative in (4) Eng and (5) LR treatments. Hypothesis 2.c predicts that in (3) the 50-50 treatment the supplier shares as in (4) and (5). Likewise, we expect to see that buyers are most likely to choose "compete" in (1) SR and (2) PM treatments, and least likely to choose "compete" in (4) Eng and (5) LR treatments. For the buyer's decision, we expect that the result in (3) the 50-50 treatment falls in between (1-2) and

(4-5). From Hypothesis 3, we expect to see an increasing trend of collaboration as we examine the results from treatment (1) to treatment (5) in an increasing order.

Table 1 shows the frequency of the suppliers' decision to "share" and the buyers' decision to "compete if the supplier shared" and "compete if the supplier did not share"²⁰ in each treatment. It also shows the frequency with which the buyer's implemented decision was "compete", and the frequency of the collaborative outcome. For the pairwise comparison of the data across treatments and the trend tests, we consider subject level data (each subject's average decision across all rounds played is considered as one observation for the test). Table 2 presents probit regression results estimating for each treatment the probabilities of suppliers choosing "share", buyers choosing "compete" in cases the supplier chose "share" and "not share", buyers implementing the "compete" decision, and both firms choosing the collaborative action. In all cases we control for round within a relationship, period of play in the session, and subject fixed effects. Recall that our experiment uses the strategy method, which asks suppliers and buyers to choose an action for each contingency. Thus, we are able to collect the data on all the dependent variables in every period.

Table 1 General Results - Frequency of Collaborative Outcomes

	(1)	(2)	(3)	(4)	(5)
Treatment	Supplier's Decision (Share)	Buyer's Decision (Compete if Shared)	Buyer's Decision (Compete if Not Shared)	Buyer's Implemented Decision (Compete)	Collaborative Outcome (Share/Not Compete)
	(%)	(%)	(%)	(%)	(%)
SR	18.2	67.6	78.5	77.6	5.3
PM	29.2	64.3	73.2	71	9.5
50 – 50	33.5	71.2	75.1	74.6	10.8
Eng	40.2	50.7	68.6	62.7	24.6
LR	38.5	59.0	56.6	58.5	21.3

Table 1 shows 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 ($p < 0.001$)²¹. However, pair-wise comparisons across treatments show that the only significant difference is between all treatments and the SR treatment (Wilcoxon rank-sum test $p < 0.05$ for all comparisons to the SR treatment). Average sharing is not significantly different across the PM, 50 – 50, Eng, and LR treatments. We observe similar results in the regression presented

²⁰ In the 50 – 50 treatment, "compete if shared" and "compete if not shared" are the decisions of the player whose recommendation was actually implemented.

²¹ The non-parametric test for trends across ordered groups is an extension of the Wilcoxon rank-sum test (Cuzick 1985).

Table 2 General Results

	(1)	(2)	(3)	(4)	(5)
	Supplier's Decision (Share)	Buyer's Decision (Compete if Shared)	Buyer's Decision (Compete if Not Shared)	Buyer's Implemented Decision (Compete)	Collaborative Outcomes (Share/Not Compete)
Coefficients					
PM	0.624*** (0.206)	-0.173 (0.145)	-0.208 (0.134)	-0.247** (0.126)	0.389* (0.233)
(50 – 50)	0.778*** (0.201)	0.034 (0.142)	-0.122 (0.132)	-0.102 (0.124)	0.462** (0.227)
Eng	0.964*** (0.207)	-0.551*** (0.147)	-0.324** (0.136)	-0.465*** (0.127)	1.055*** (0.230)
LR	0.897*** (0.201)	-0.298** (0.142)	-0.654*** (0.131)	-0.574*** (0.123)	0.849*** (0.225)
Period	-0.008*** (0.003)	-0.018*** (0.002)	-0.014*** (0.003)	-0.017*** (0.003)	0.013*** (0.003)
Round	-0.064*** (0.012)	0.032*** (0.011)	0.015 (0.011)	0.013 (0.011)	-0.009 (0.014)
Constant	-0.938*** (0.154)	0.745*** (0.110)	1.030*** (0.103)	1.045*** (0.098)	-2.085*** (0.181)
Observations	4286	4286	4286	4286	4286
Nr. of Subjects	143	143	143	143	143

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

on Table 2: 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. Therefore, it appears that the supplier's decision to share the innovation depends primarily on the length of the firm-level relationship, and does not differ significantly based on the allocation of decision rights within the buyer firm.

The frequency of the buyer's choice to "compete if the supplier shared" presents a decreasing trend from SR to LR, as predicted in Hypothesis 3.b (non-parametric trend test p-value = 0.028). As shown in column 2 of Table 2, only the Eng and LR treatments present a significant decrease relative to the SR benchmark (marginal effects Eng: -0.213 , LR: -0.114). We also observe that in the 50 – 50 treatment, "compete if the supplier shared" is chosen more often than expected (recall that it was expected to be significantly lower than in the SR benchmark). This deviation from our predictions is analyzed further in the following sections. Similarly, "compete if not shared" has a significant decreasing trend from SR to LR (p-value < 0.001), and it is only significantly lower than in the SR benchmark in the Eng and LR treatments (marginal effects Eng: -0.115 , LR: -0.238). We then look at the frequency of outcomes where the buyer competes (column 4 in Tables 1 and 2). We observe a significant trend of reduced competition (increased collaboration) from SR to LR (test for trends: $p < 0.001$), which supports Hypothesis 3.b. Table 1 shows that, while buyers compete in 77.6% of the outcomes in the SR treatment, they do so 71% of the times in the PM treatment

($p = 0.009$). As in the case of the suppliers' sharing decisions, this suggests that the PM treatment presents increased collaboration relative to the SR benchmark. Nonetheless, the largest difference relative to the SR benchmark is in the Eng and LR treatments (as predicted by Hypotheses 1.b and 2.b). Table 1 shows that the probability that a buyer will "compete" drops to 62.7% and 58.5% in the Eng and LR treatments respectively. These results are supported in Table 2, which shows that the frequency of outcomes where the buyer chose "compete" is lower in the PM treatment than in the SR benchmark (marginal effects: -0.087), and even lower in the Eng and LR treatments (marginal effects: -0.171 and -0.212 for Eng and LR respectively). Therefore, the buyer's decision depends not just on the firm-level relationship length, but also on the decision rights within the firm.

The results described above indicate that there is a significant trend of increased collaboration in both suppliers' and buyers' actions across treatments –when ordered (SR)-(PM)-(50 – 50)-(Eng)-(LR). However, suppliers' and buyers' actions present some deviations from the predictions and these deviations go in opposite directions. Hypothesis 2 predicts that the frequency with which suppliers choose "share" in the PM treatment will be as low as in the SR benchmark, and that in the 50 – 50 and Eng treatments it will be as high as in the LR treatment. We observe that the probability that suppliers choose "share" actually increases even *sooner* than expected. Even in the PM treatment suppliers share significantly more frequently than in the SR benchmark. On the other hand, the frequency with which buyers choose "compete" decreases *later* than predicted. Hypothesis 2 predicts that in the 50 – 50 treatment competition should be lower than in the SR benchmark and higher than in the LR treatment. We find that in the 50 – 50 treatment, the frequency with which buyers choose "compete if the supplier shared" is not significantly lower than in the PM treatment. It is only in the Eng and LR treatments that buyers choose "compete if the supplier shared" significantly less often than in the SR benchmark. We will examine next the impact this has on 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 1 shows that the frequency with which collaborative outcomes occur is not statistically different across the SR, PM, and 50 – 50 treatments. In the Eng treatment, it is significantly higher than in the previous three (Eng vs. 50 – 50: $p = 0.007$) and not significantly different from the LR treatment. The regression presented in Table 2 shows that all treatments present a higher frequency of collaborative outcomes than the SR baseline (PM marginal effects: 0.072). In particular, the Eng and LR treatments have an even higher frequency of collaborative outcomes (Eng and LR marginal effects: 0.249 and 0.183 respectively). As before, these results depart from the hypotheses in two ways: First, the frequency of collaborative outcomes in the PM treatment is (marginally) higher

than in the SR treatment, while they should be equivalent. Second, the frequency of collaborative outcomes in the 50 – 50 treatment is not significantly higher than in the PM treatment, while both buyers and suppliers are expected to be more collaborative. We explore these results in section 6.3. by analyzing each of the Buyer-as-Two-Employees treatments in more detail.

6.2. Profits

In line with the previous results, we find that suppliers', buyers', and total profits present an increasing trend from SR to LR ($p = 0.001, 0.009$, and < 0.001 respectively). Average profits for suppliers, buyers, and both players combined are presented in Table 9 in the Appendix. The results show that suppliers' profits in the PM treatment are slightly higher than the theoretical expected profits from non-collaborative strategies (16.14 vs. 15.25, one sided t-test $p = 0.077$). However, suppliers only earn significantly higher profits than in the SR benchmark in the Eng and LR treatments. On the other hand, buyers benefit from all Buyer-as-Two-Employees treatments. Buyers' profits are significantly higher than in the SR benchmark, in all the other treatments – PM, 50 – 50, Eng, and LR. This is consistent with the previous findings about suppliers' and buyer's actions. While suppliers share more frequently in all the Buyer-as-Two-Employees treatments than in the SR benchmark, the frequency with which buyers choose “compete” only decreases significantly relative to the SR benchmark in the Eng treatment (where the engineer's recommendation is the one that is always implemented).

Table 10 in the Appendix 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 Buyer-as-Two-Employees treatments, 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, it increases 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 difference in total profits across treatments reflects the pattern of increased frequency of suppliers' sharing from SR to LR.

6.3. Employee Decisions in the Buyer-as-Two-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 have also found that the Buyer-as-Two-Employees treatments depart in some ways from the theoretical predictions. In this section, we analyze the results of the Buyer-as-Two-Employees treatments in detail. We observe 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 some differences with the LR treatment. In the last subsection, we show (4) that there exists an interplay between the employees beyond what the theory predicts.

6.3.1. PM and SR Treatments Although the theory predicts that the SR and PM treatments should be identical, the results on Table 3 show important differences. The first and second columns in Table 3 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. Column 5 presents the difference between columns 3 and 4. We observe that this difference 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 11 in the Appendix confirms this result. In the PM treatment, a regression of the average profit per round within a relationship on the average frequency of sharing in that relationship 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, an increase to sharing 60% of the times would be associated with an increase in expected profit of 2.52 points per round. Since the average supplier profit in the PM treatment is 16.59 points per round, this implies a 15% increase in profits. In the experiment, suppliers seem to (at least partially) acknowledge this difference: they share 18.2% of the times in the SR treatment and 29.2% of the times in the PM treatment.

Table 3 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	75.1	14.16	15.93	-1.78	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 other columns present the suppliers’ expected profits given the buyer’s choices in the experiment.

6.3.2. 50-50 Treatment As shown in Proposition 2, in the 50 – 50 treatment an equilibrium where the supplier chooses “share” in every round, the engineer recommends “not compete if the supplier shared”, and the procurement manager recommends “compete” in every round, can be supported with the continuation probability of 0.75 used in the experiment. Since one of the two

recommendations 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 engineers and procurement managers compete more often than in the PM and Eng treatments. Figure 4 in the Appendix shows that the percentage of times an engineer chooses “compete when the supplier shared” is higher in the 50 – 50 treatment than in the PM and Eng treatments ($p = 0.013$ for PM, $p = 0.033$ for Eng). Similarly, a procurement manager chooses to “compete when the supplier shared” significantly more often in the 50 – 50 treatment than in the PM and Eng treatments ($p = 0.055$ and $p = 0.071$ for PM and Eng respectively). 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%). This result suggests that a more ambiguous allocation of decision rights may have a negative effect on employees’ propensity to collaborate.

We also study how collaboration evolves throughout a relationship in the 50 – 50 treatment relative to the PM and Eng treatments. Figure 2 shows how a collaborative outcome in the first round is sustained throughout the relationship in the different treatments. The figure indicates the frequency of a collaborative outcome in rounds 2 to 6 of the relationship conditional on collaboration in round 1 (solid line), and conditional on non-collaboration in round 1 (dashed line). We observe that, 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 rounds, in the PM and 50 – 50 treatments this is less likely to occur. We also observe that in all treatments, if the collaborative outcome is not reached in the first round, it is very unlikely that it will be reached in a subsequent round. To confirm these results, Table 4 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 a collaborative outcome 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 when there was a collaborative outcome in the previous round (0.31 and 0.43 in the PM and 50 – 50 treatments vs. 0.77 and 0.79 in the Eng and LR treatments). As a result, the 50 – 50 treatment is

more similar to the PM treatment than the Eng treatment in terms of how collaboration is sustained in a relationship.

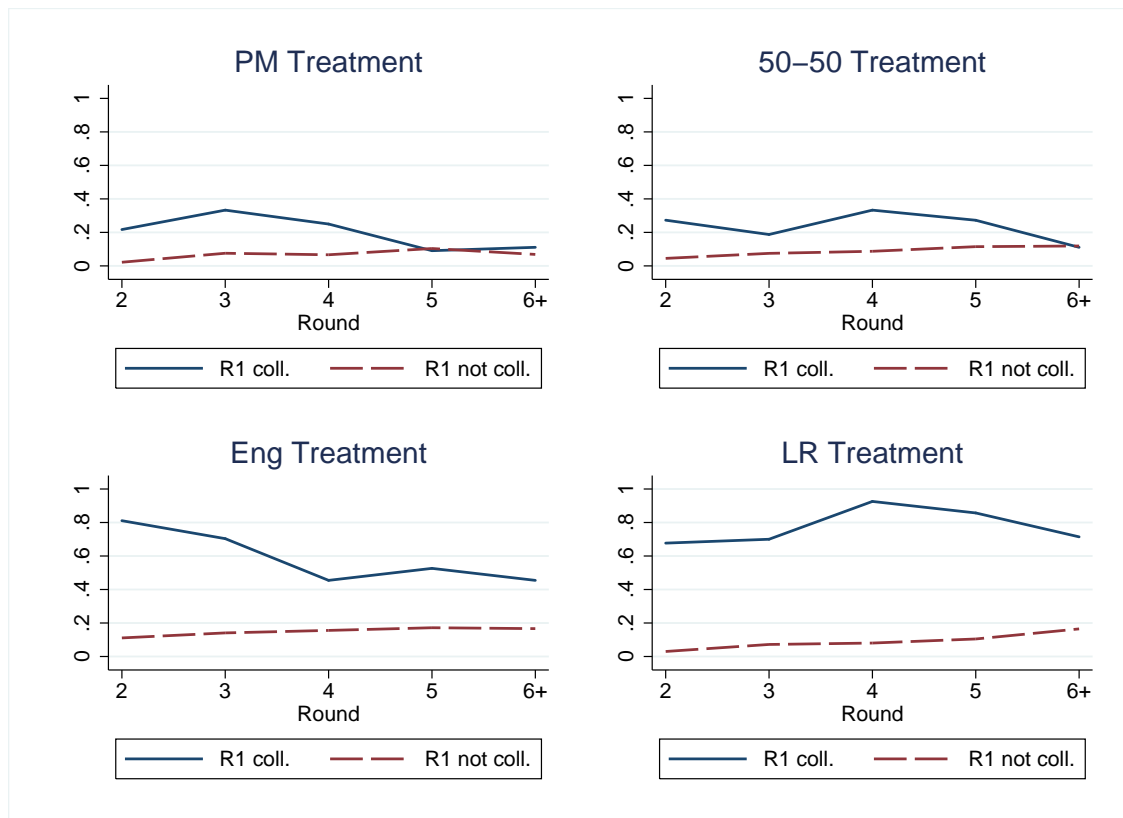
The previous result, that collaboration is harder to sustain in the 50 – 50 treatment than in the Eng treatment, is not entirely surprising. In fact, half the times the procurement manager’s recommendation is implemented and procurement managers are expected (according to the theory) not to collaborate. However, beyond the procurement managers’ non-collaborative actions, when we look at the suppliers’ and engineers’ strategies in both treatments we observe differences which are not explained by the theory. First, while in the 50 – 50 the supplier’s decision should only be correlated with the engineer’s previous decision, we observe that when the procurement manager’s decision was implemented in the previous round, the supplier’s decision is also correlated with the previous procurement manager’s decision (as shown in table 12 in the Appendix). Since procurement managers choose “compete if shared” more often than the engineers, this has a negative impact in sustaining cooperation in the 50 – 50 treatment. Second, engineers punish the supplier more for not sharing in the Eng treatment than in the 50 – 50 treatment. When we compare the engineer’s average “compete if shared” in all the rounds before and after the first time the supplier did not share in a relationship, we observe that the increase in the decision to compete after the first time the supplier did not share is higher in the Eng treatment than in the 50 – 50 treatment (0.63 before vs. 0.72 after in the 50 – 50 treatment, and 0.36 before vs. 0.60 after in the Eng treatment). This also holds for the decision to “compete if not shared” (0.72 before vs. 0.77 after in the 50 – 50 treatment, and 0.62 before vs. 0.72 after in the Eng treatment). In addition, the engineers’ response immediately after a round where the supplier did not share the innovation is also different between the 50 – 50 and Eng treatments. We look at the difference in engineers’ frequency of choosing “compete if shared” and “compete if not shared” in the first round following a round where the supplier shared the innovation, and the first round following a round where the supplier did not share the innovation. In the 50 – 50 treatment, the frequency of “compete if not shared” following a round where the supplier shared is 0.67, and it is 0.78 following a round where the supplier did not share. In the Eng treatment it is 0.58 following a round where the supplier shared vs. 0.76 following a round where the supplier did not share. Similarly, the frequency of “compete if shared” increases from 0.53 following a round a round where the supplier shared to 0.75 following a round where the supplier did not share in the 50 – 50 treatment, while it increases from 0.36 to 0.62 in the Eng treatment²². These results suggest that the punishment that engineers impose on suppliers for not sharing is stronger in the Eng treatment than in the 50 – 50 treatment, which reinforces the difficulty of sustaining collaboration in the 50 – 50 treatment.

²² We note that the engineers’ frequency of choosing “compete if shared” and “compete if not shared” is not significantly different if in the previous round of the relationship the innovation did not happen than if it happened and the supplier did not share. Procurement managers’ choices are not significantly different if the supplier shared in the previous round than if he did not share or if the innovation did not occur.

Table 4 Collaboration throughout Relationships

		SR	PM	50–50	Eng	LR
Pr (CO)*		0.053	0.095	0.108	0.246	0.213
(A)	Pr (CO $R = 1$)	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. “Pr (CO | $R = 1$)” indicates the probability of collaboration in the first round of a relationship. “Pr (CO | CO prev round, $R > 1$)” indicates the probability of a collaborative outcome conditional on a collaborative outcome in the previous period, for periods 2 onwards. “Pr (CO | no CO prev round, $R > 1$)” indicates the probability of a collaborative outcome conditional on not having a collaborative outcome in the previous period, for periods 2 onwards. (%) represents the percentage of all collaborative outcomes that occur in a particular treatment corresponding to cases (A), (B), and (C).

**Figure 2 Effect of first round collaboration of subsequent rounds**

6.3.3. Eng and LR Treatments In the Eng treatment the engineer’s 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). Our experimental results show that there are no significant differences between the

two treatments in terms of either the suppliers' average sharing decisions or the buyers' average competing decisions (columns 1 and 4 in Tables 1 and 2).

We have established in the theory section the existence of two equilibria where collaboration – the supplier shares the innovation and the buyer does not compete if the supplier shared – can be sustained (Proposition 1). For these equilibria to arise, both the buyer and the supplier need to be “non-defectors”: that is, while in the collaborative state, the supplier always shares and the buyer never competes if an innovation was shared. The difference between the two equilibria resides in the actions taken when an innovation is not shared. Equilibrium *E1* arises if buyers are skeptical and suppliers are forgiving. Equilibrium *E2* arises if buyers non-skeptical and suppliers are non-forgiving. In order to identify these types of players in our data, we define a “pre-betrayal round” to be every round in a relationship up until the first time that either (1) the innovation happens and the supplier does not share, or (2) the innovation happens, the supplier shares, and the buyer chooses to compete. Note that a “betrayal” can only happen if the innovation occurs. We define how much of a “defector” a supplier is based on the fraction of times she chose not to share in all the pre-betrayal rounds she played. Similarly, we define how much of a “defector” a buyer is depending on the fraction of times he chose “compete” when an innovation was shared in all the pre-betrayal rounds he played. We next define how “skeptical” a buyer is as the fraction of times he chose “compete” in the ambiguous node in all the pre-betrayal rounds he played. Finally, we define the supplier’s “forgiveness” as the fraction of times a supplier chose “share” in pre-betrayal rounds where in the previous round the innovation did not happen and the buyer chose to compete. The histograms in Figure 8 show the distributions of each type of player in our experiment. In this section, we look at the data from the Eng and LR treatments combined and we refer to the buyers in the LR treatment and the engineers in the Eng treatment combined as “buyers” (we later comment on the differences between these two treatments). The individual measures of “skeptical” for buyers and “forgiving” for suppliers were only defined for non-defector buyers and non-defector suppliers respectively. We define continuous and dichotomous variables for these four measures to use in our analysis. The cutoff to define whether a subject is a defector is at 0.9 for buyers (with this cutoff 37.9% of the buyers are defectors) and at 0.9 for suppliers (so that 15.5% are defectors)²³. The cutoff for the discrete forgiving dummy variable is at strictly greater than zero, which makes 40% of the non-defector suppliers forgiving (note that 60% had a value of 0, therefore no cutoff would have made them forgiving). The cutoff for the discrete skeptical dummy variable is at greater or equal than 0.5 (the median of the distribution)²⁴.

²³ The cutoffs were set based on the histograms in Figure 8 and so that only subjects with extreme behavior (100%) are considered defectors.

²⁴ The frequency of matchings between the different types of suppliers and buyers given the cutoff points that we chose, is presented in Tables 13 and 14. Note that these frequencies do not capture the frequency with which each

First, we observe by looking at the histograms in Figure 8 that a fraction of suppliers and buyers show a persistent defector behavior. In particular, over a third of the buyers chose to compete after the supplier shared the innovation in every pre-betrayal round they played. We also note that the suppliers' measure looks more uniformly distributed, while the buyers' behavior seems to be more extreme. Column 7 in Table 5 shows that being a defector is profitable for a buyer ($\beta = 7.223$, $p = 0.010$), which explains why car manufacturers may have incentives to re-share their suppliers' innovations. In addition, Table 15 in the Appendix shows that while suppliers' profits are higher when their type matches the buyer's type (both are defectors or both are non-defectors), buyers are indifferent between being defectors or not regardless of the supplier's type. From an efficiency perspective, our data shows that subjects being defectors has a significant negative effect on total surplus. Both buyers and suppliers are worse off when matched with a defector (Table 5 shows this with a regression of profits on continuous type variables controlling for subject random effects, and Table 15 shows the differences in average profits across types using the dichotomous type variables). Column 5 in Table 6 shows that total surplus increases significantly when the supplier is not a defector (relative to a baseline case where both the buyer and supplier are defectors). In addition, if the supplier is not a defector, total surplus increases even further when the buyer is not a defector either ($\beta = 32.462$, $p < 0.001$).

While the buyer and supplier being non-defectors is generally desirable in terms of efficiency – as it favors sustained collaboration– we have seen that, in addition, two different collaborative equilibria may arise depending on whether the buyer is skeptical and the supplier is forgiving ($E1$) or whether the buyer is non-skeptical and the supplier is non-forgiving ($E2$). Consistent with these two equilibria, we observe in our experiment a high heterogeneity in the distributions of skepticism and forgiveness. The two bottom histograms in Figure 8 show suppliers' forgiveness and buyers' skepticism (conditional on suppliers and buyers not being defectors respectively). In the “forgiveness” histogram, we observe that 60% of the suppliers were non-forgiving in every occasion where they faced the decision to forgive or not. These suppliers considered buyers choosing “compete” if an innovation was not shared in a pre-betrayal round as a reason to trigger punishment, as equilibrium $E2$ (but not $E1$) prescribes. Despite the high fraction of non-forgiving suppliers, 40% of the buyers were highly skeptical and chose “compete” if an innovation was not shared in at least seventy percent of the pre-betrayal rounds they played.

Given the high heterogeneity in skeptical and forgiving types, a natural question that follows is what to expect in the cases where the matching between these types is different than those resulting in $E1$ and $E2$. When the buyer is skeptical and the supplier is non-forgiving, a new equilibrium type of equilibria is reached in the data, but rather the frequency with which different types of players are matched with each other in the experiment.

(E3) arises. In this equilibrium, the first round where the innovation does not happen triggers a transition from the collaborative state to the punishment state. This is because the buyer will choose “compete” and the supplier will not forgive this action. Therefore, this equilibrium is only partially collaborative. The proof is presented in section 9.4.3. in the Appendix. A fourth equilibrium where the buyer is non-skeptical and the supplier is forgiving cannot arise. It can be easily shown that in this case the buyer has an incentive to deviate and compete when an innovation is not shared since the supplier would forgive this action and will not trigger punishment²⁵.

Table 5 Effect of Supplier and Buyer types (continuous measures) on supplier and buyer profit

Coefficients	Supplier profit in a relationship				Buyer profit in a relationship			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sup Defector (Cont.)	-12.805*** (2.445)				-28.437*** (3.116)			
Sup Forgiving (Cont.)		0.652 (3.518)				6.676* (3.863)		
Buyer Defector (Cont.)			-20.047*** (1.410)				7.228*** (2.788)	
Buyer Skeptical (Cont.)				-9.301*** (2.513)				-1.573 (3.286)
Period	0.094 (0.074)	0.077 (0.107)	0.108* (0.063)	0.167* (0.100)	-0.045 (0.123)	0.068 (0.177)	-0.060 (0.133)	-0.186 (0.143)
Constant	25.250*** (1.750)	20.534*** (2.406)	29.363*** (1.314)	27.678*** (2.270)	54.070*** (2.464)	40.154*** (3.319)	36.481*** (2.534)	41.066*** (2.793)
Observations	452	265	452	275	452	265	452	275
Nr. of Subjects	58	84	58	57	84	79	84	51

OLS regression with subject random effects. Each relationship between a buyer and a supplier is considered as one observation. Columns 2, 4, 6, and 8 consider only relationships where the supplier and the buyer are non-defectors. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

The differences in buyer’s play at the ambiguous node (skeptical or not), and the supplier’s reaction to that action (forgiving or not), has a significant impact on profits. Table 5 shows that being matched with a skeptical buyer leads to significant lower supplier profits ($\beta = -9.301$, $p < 0.001$). Table 15 in the Appendix shows that this effect is stronger when the supplier is non-defector (defector supplier’s profit: 16.77 with skeptical vs. 20.83 with non-skeptical, non-defector supplier’s profit: 24.32 vs. 30.26). On the buyer’s side, being matched with a forgiving supplier leads to (marginally) significantly higher profits ($\beta = 6.676$, $p = 0.084$), in particular if the buyer is skeptical (Table 15). As a result, column 6 in Table 6 shows that total surplus is the lowest when the buyer is skeptical and the supplier non-forgiving and increases significantly if either the buyer is non-skeptical ($\beta = 10.420$, $p = 0.046$) or the supplier is forgiving ($\beta = 15.157$, $p = 0.008$). This is because being skeptical does not help the buyer but it hurts the supplier, and being non-forgiving does not help the supplier while

²⁵ The buyer’s incentive compatibility requires that $NN_m + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1-\pi)NN_m] \geq NC_m + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1-\pi)NN_m]$, which does not hold for any δ with the payoffs in Figure 1 since $NC_m = 17$ and $NC_m = 12$.

Table 6 Effect of Supplier and Buyer types (continuous measures) on total profits

Coefficients	Total Supply Chain Profit					
	(1)	(2)	(3)	(4)	(5)	(6)
Supplier Defector (Cont.)	-41.344*** (3.438)					
Supplier Forgiving (Cont.)		4.765 (6.457)				
Buyer Defector (Cont.)			-13.392*** (3.147)			
Buyer Skeptical (Cont.)				-9.869** (4.646)		
Sup Def – Buyer NoDef (Dum.)					6.524 (6.432)	
Sup NoDef – Buyer Def (Dum.)					20.161*** (5.686)	
Sup NoDef – Buyer NoDef (Dum.)					32.462*** (5.549)	
Skept – Forg (Dum.)						15.157*** (5.701)
Non Skept – Non Forg (Dum.)						10.420** (5.217)
Non Skept – Forg (Dum.)						17.201*** (5.888)
Period	0.048 (0.132)	0.151 (0.186)	0.050 (0.150)	-0.054 (0.202)	0.046 (0.143)	-0.041 (0.225)
Constant	79.486*** (2.684)	61.159*** (4.334)	66.148*** (2.861)	69.308*** (3.949)	35.680*** (5.559)	62.122*** (3.962)
Observations	452	265	452	275	452	223
Nr. of Subjects	58	35	84	51	84	51

OLS regression with subject random effects. Each relationship between a buyer and a supplier is considered as one observation. Columns 2 to 6 consider only relationships where the supplier and the buyer are non-defectors. Robust standard errors reported in parentheses. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

it hurts the buyer (Table 5). This indicates that actions matter not only when there is an innovation, but also in the periods where the innovation does not occur. While some buyers are skeptical and choose to compete, others are non-skeptical and choose not to compete. This skepticism is a loss in efficiency which suppliers punish by being non-forgiving (which hinders surplus even further). In addition, following the previous definition of a “betrayal”, we find that a betrayal is more likely to happen in a relationship between skeptical and non-forgiving types (probability of a relationship having a betrayal = 0.813) than if the buyer is instead non-skeptical (0.630, $p = 0.05$), or if the supplier is instead forgiving (0.444, $p = 0.0005$) or both (0.364, $p < 0.0001$).

Previously, we pooled the data from the Eng (subjects playing as engineers) and LR (subjects playing as buyers) treatments since, according to our model, they have equivalent incentives. We now want to see whether there are any significant differences between these two treatments. When we compare the average in buyers’ continuous measures of “Defector” and “Skeptical” across the Eng and LR treatments we find that, while the difference in “defector” is not significant (0.508 vs.

0.556, p-value = 0.804), the difference in “skeptical” is (0.733 vs. 0.552, p-value= 0.045). This result indicates that the presence of the procurement managers (which is the only difference across the two treatments) makes the buyer’s decision less collaborative specifically by making engineers more skeptical. This explains some of our previous results. First, in Table 1 we observe that “compete if not shared” is higher in the Eng treatment than in the LR treatment. We know now that this is due to individual differences in subject’s “skeptical” measure. Second, Figure 2 shows that a collaborative outcome in the first round of a relationship is less likely to be sustained in the Eng treatment than in the LR treatment. This is consistent with more skeptical buyers in the Eng treatment, which increases the likelihood that the semi-collaborative equilibrium $E3$ will arise. There is no significant difference in suppliers’ measures of “Defector” or “Forgiving” across the two treatments.

6.3.4. Interplay Between Employees Figure 9 shows how the engineer’s play in the Eng treatment correlates with his own previous play and with the play of the previous procurement manager he interacted with. The bar chart on the left shows the engineer’s choice to “compete if the supplier shared” in every round of the relationship (except the first one) with two columns, one for the case where he chose to compete in the previous round and another for the case where he chose not to compete in the previous round. 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 relationship²⁶. This result is not surprising: it is consistent with hypothesis 2.b, which was derived assuming trigger strategies, but it 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 9 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 7. Column 6 shows that, the engineer’s decision is correlated with the previous procurement manager’s decision in the round immediate before, even after controlling for the engineer’s own decision in the previous round. 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’s recommendation 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 of play in the relationship and are only consistent with their own previous actions (note that the regression only

²⁶ 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.

considers the cases where in the previous round the innovation did occur, so that the procurement manager is always 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 Buyer-as-Two-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 – but only do so when the procurement manager has a say in the final decision (PM and 50 – 50 treatments). We therefore see in the 50 – 50 and Eng treatments that both suppliers and engineers deviate from the collaborative equilibria described in Propositions 1 and 2 which prescribe that the supplier's and engineer's choices should depend only on each other's.

Table 7 Strategy Analysis

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

7. Discussion

Our experimental results show that, when firms have a long-term relationship, the allocation of decision rights matters. One of the most interesting observations is that the results of the 50 – 50 treatment substantially deviate from the theoretical predictions. The theory predicts that in the PM and 50 – 50 treatments the procurement manager always recommends “compete”, and that in

the Eng and 50 – 50 treatments the engineer always recommends “not compete”. The experimental results show that procurement managers choose “compete” more often in the 50 – 50 treatment than in the PM treatment, and engineers choose “compete” more often in the 50 – 50 treatment than in the Eng treatment. That is, both procurement managers and engineers choose “compete” more often in a setup where the probability that their recommendation will be implemented is 0.5, than when it will be implemented for sure. The results show that the average frequency with which suppliers choose “share” is not significantly different across the PM, 50 – 50, and Eng treatments (0.29, 0.34, and 0.40 respectively). One could expect that suppliers would punish the procurement managers and engineers for choosing “compete” more often in the 50 – 50 treatment. However, suppliers do not share less often in the 50 – 50 treatment than in the PM or Eng treatments (subject level average differences $p = 0.301$ and $p = 0.391$ respectively). This explains why the suppliers’ profits and share of the total surplus are the low in the 50 – 50 treatment (Table 9).

We also observe that, among the rounds where collaborative outcomes arise in the 50 – 50 treatment, a relatively high proportion (37%) occurs when the procurement manager’s recommendation is implemented. This result seems to be driven by a number of procurement managers who often choose “not compete” when the supplier shared. While 50% of the procurement managers in the 50 – 50 treatment choose to compete in more than 90% of the rounds they played, another 23% choose to compete only 50% of the times or less. This is surprising since procurement managers change partners in every round, even if the relationship between the firms continues. A potential explanation for this can be found in Kandori (1992b) and Kandori (1992a), who have shown that collaboration can be supported in a sub-game perfect equilibrium in a setting where subjects change partners over time. They show that 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 chose not to “compete if shared” in the 50 – 50 treatment. The subjects in a session may constitute a “community” where, if all procurement managers choose “compete”, they ultimately get punished in the future relationships they join. This can provide enough incentives for procurement managers to collaborate²⁷.

Another surprising result is that actions of short-term agents (procurement managers) influence the team’s future actions even *after* they have left the teams. On the surface, the result seems to be in line with the existing literature on group decision making. Using the setting of sequential gift

²⁷ The result in Kandori (1992b) and Kandori (1992a) can also explain why procurement managers’ decision to “compete if the supplier shared” is lower than 100% in the PM and SR treatments (they are 64% and 68% respectively).

exchange game²⁸, Ambrus et al. (2015) analyze how the preferences of individuals get aggregated when they make decisions as a group²⁹. They show that when subjects are allowed to discuss freely and deliberate before making a group decision, group members can influence each other's decisions³⁰. In our setting, social influence is less likely to occur. Our setting does not allow for discussion and deliberation. In addition, unlike their setting where everyone plays the same game with the same objective function, procurement managers and engineers have different matching rules and different monetary incentives in our setting. In spite of such limitations, our results still show that a short-term agent influences group decisions even after they left the team. We conjecture that this impact will be much stronger if subjects deliberate together (rather than submit their individual decisions). Our result implies that, even if a firm rotates and reassigns procurement managers often, they may have a long-lasting influence on the relationship with a supplier.

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 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 employee 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. We observe that, in addition to the length of the relationship between the firms, the allocation of decision rights to employees also matters.

²⁸ The gift exchange game (Fehr et al. 1993, Brandts and Charness 2004) 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.

²⁹ 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.

³⁰ Ambrus et al. (2015) reference two social psychology mechanisms which explain why subjects may behave differently 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.

Having both short- and long-run focused employees involved in the decision 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.

Our theory and experiments also show that, in long-term relationships, what happens in the ambiguous situation where the buyer receives no innovations from the supplier (and is unsure of whether the supplier had no innovation or if she had the innovation and chose not to share it) also matters. Many buyers in our experiment are skeptical and choose to compete in this case. These buyers are better off when matched with a forgiving supplier who does not punish this action. However, our results show that 60% of the suppliers are completely non-forgiving and switch to a punishment state following a skeptical action. Skepticism on the buyer's side has a significant detrimental effect on suppliers' profits, while it does not make buyers better off. Therefore, buyers being skeptical is a direct loss in efficiency (total surplus) which is further accentuated when the supplier is non-forgiving.

Finally, 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.

References

- Ambrus, Attila, Ben Greiner, Parag A Pathak. 2015. How individual preferences are aggregated in groups: An experimental study. *Journal of Public Economics* **129** 1–13.
- Aoyagi, Masaki, Guillaume Fréchette. 2009. Collusion as public monitoring becomes noisy: Experimental evidence. *Journal of Economic theory* **144**(3) 1135–1165.

-
- Baker, George, Robert Gibbons, Kevin J Murphy. 2002. Relational contracts and the theory of the firm. *Quarterly Journal of economics* 39–84.
- Becker-Peth, Michael, Elena Katok, Ulrich W Thonemann. 2013. Designing buyback contracts for irrational but predictable newsvendors. *Management Science* **59**(8) 1800–1816.
- Beer, Ruth, Hyun-Soo Ahn, Stephen Leider. 2017. Can trustworthiness in a supply chain be signaled? *Forthcoming in Management Science* .
- Bernstein, Fernando, A Gürhan Kök. 2009. Dynamic cost reduction through process improvement in assembly networks. *Management Science* **55**(4) 552–567.
- Bó, Pedro Dal. 2005. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *American Economic Review* 1591–1604.
- Bó, Pedro Dal, Guillaume R Fréchette. 2011. The evolution of cooperation in infinitely repeated games: Experimental evidence. *The American Economic Review* **101**(1) 411–429.
- Bolton, G.E., 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.
- Brandts, Jordi, Gary Charness. 2004. Do labour market conditions affect gift exchange? some experimental evidence*. *The Economic Journal* **114**(497) 684–708.
- Brandts, Jordi, Gary Charness. 2011. The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics* **14**(3) 375–398.
- Brinkhoff, Andreas, Özalp Özer, Gökçe Sargut. 2015. All you need is trust? an examination of inter-organizational supply chain projects. *Production and Operations Management* **24**(2) 181–200.
- Burt, David N. 1989. Managing suppliers up to speed. *Harvard Business Review* **67**(4) 127–135.
- Cabral, Luis, Erkut Y Ozbay, Andrew Schotter. 2014. Intrinsic and instrumental reciprocity: An experimental study. *Games and Economic Behavior* **87** 100–121.
- Camera, Gabriele, Marco Casari. 2009. Cooperation among strangers under the shadow of the future. *The American Economic Review* **99**(3) 979–1005.
- Cooper, David J, Kai-Uwe Kuhn. 2009. Communication, renegotiation, and the scope for collusion .
- Cox, James C. 2002. Trust, reciprocity, and other-regarding preferences: Groups vs. individuals and males vs. females. *Experimental Business Research*. Springer, 331–350.
- Cui, T.H., J.S. Raju, Z.J. Zhang. 2007. Fairness and channel coordination. *Management Science* **53**(8) 1303–1314.
- Cuzick, Jack. 1985. A wilcoxon-type test for trend. *Statistics in medicine* **4**(4) 543–547.
- Dal Bó, Pedro, Guillaume R Fréchette. 2013. Strategy choice in the infinitely repeated prisonersâ€™ dilemma. *Available at SSRN 2292390* .

- Davis, Andrew M, Kyle Hyndman. 2016. An experimental investigation of managing quality through monetary and relational incentives. *Forthcoming in Management Science* .
- Dyer, Jeffrey H. 1996. How chrysler created an american keiretsu. *Harvard Business Review* .
- Dyer, Jeffrey H, Nile W Hatch. 2006. Relation-specific capabilities and barriers to knowledge transfers: creating advantage through network relationships. *Strategic management journal* **27**(8) 701–719.
- Engle-Warnick, Jim. 2007. *Five indefinitely repeated games in the laboratory*. Citeseer.
- Engle-Warnick, Jim, Robert L Slonim. 2004. The evolution of strategies in a repeated trust game. *Journal of Economic Behavior & Organization* **55**(4) 553–573.
- Engle-Warnick, Jim, Robert L Slonim. 2006a. Inferring repeated-game strategies from actions: evidence from trust game experiments. *Economic theory* **28**(3) 603–632.
- Engle-Warnick, Jim, Robert L Slonim. 2006b. Learning to trust in indefinitely repeated games. *Games and Economic Behavior* **54**(1) 95–114.
- Fehr, Ernst, Georg Kirchsteiger, Arno Riedl. 1993. Does fairness prevent market clearing? an experimental investigation. *The Quarterly Journal of Economics* 437–459.
- Fischbacher, U. 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* **10**(2) 171–178.
- Fréchette, Guillaume R, Sevgi Yuksel. 2013. Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Available at SSRN 2225331* .
- Fudenberg, Drew, Eric Maskin. 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* **54**(3) 533–554.
- Gibbons, Robert. 1998. Incentives in organizations. Tech. rep., National Bureau of Economic Research.
- Gibbons, Robert. 2001. Trust in social structures: Hobbes and coase meet repeated games. *Trust in society* 332–353.
- Gibbons, Robert. 2005. Incentives between firms (and within). *Management science* **51**(1) 2–17.
- Gibbons, Robert, Rebecca Henderson. 2012. Relational contracts and organizational capabilities. *Organization Science* **23**(5) 1350–1364.
- Gray, John V, Aleda V Roth, Brian Tomlin. 2009. The influence of cost and quality priorities on the propensity to outsource production*. *Decision Sciences* **40**(4) 697–726.
- Han, Sang-Lin, David T Wilson, Shirish P Dant. 1993. Buyer-supplier relationships today. *Industrial Marketing Management* **22**(4) 331–338.
- Helper, Susan, Rebecca Henderson. 2014. Management practices, relational contracts, and the decline of general motors. Tech. rep., National Bureau of Economic Research.
- Ho, T.H., 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, Dorothee, Kyle B Hyndman. 2015. Flexibility and reputation in repeated prisoners' dilemma games. *Available at SSRN 2562727* .
- Hyndman, Kyle, Dorothee Honhon. 2014. Flexibility in long-term relationships: An experimental study. *Working paper* .
- Hyndman, Kyle, Santiago Kraiselburd, Noel Watson. 2014. Coordination in games with strategic complementarities: An experiment on fixed vs. random matching. *Production and Operations Management* **23**(2) 221–238.
- Iyer, Ananth V, Leroy B Schwarz, Stefanos A Zenios. 2005. A principal-agent model for product specification and production. *Management Science* **51**(1) 106–119.
- Kalwani, Manohar U, Narakesari Narayandas. 1995. Long-term manufacturer-supplier relationships: do they pay off for supplier firms? *The Journal of Marketing* 1–16.
- Kandori, Michihiro. 1992a. Repeated games played by overlapping generations of players. *The Review of Economic Studies* **59**(1) 81–92.
- Kandori, Michihiro. 1992b. Social norms and community enforcement. *The Review of Economic Studies* **59**(1) 63–80.
- Katok, E., D.Y. Wu. 2009. Contracting in supply chains: A laboratory investigation. *Management Science* **55**(12) 1953–1968.
- Katok, Elena, Valery Pavlov. 2013. Fairness in supply chain contracts: A laboratory study. *Journal of Operations Management* **31**(3) 129–137.
- Kim, Sang-Hyun, Serguei Netessine. 2013. Collaborative cost reduction and component procurement under information asymmetry. *Management Science* **59**(1) 189–206.
- Klier, Thomas H. 2006. The supplier industry in transition—the new geography of auto production. *Chicago Fed Letter* **229**.
- Kreps, David M. 1990. Corporate culture and economic theory. *In: Alt, J., Shepsle, K. (Eds.), Perspectives on Positive Political Economy. Cambridge Univ. Press, Cambridge* .
- Kugler, Tamar, Gary Bornstein, Martin G Kocher, Matthias 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.
- Li, Cuihong, Laurens G Debo. 2009. Strategic dynamic sourcing from competing suppliers with transferable capacity investment. *Naval Research Logistics (NRL)* **56**(6) 540–562.
- Liker, Jeffrey K, Thomas Y Choi. 2004. Building deep supplier relationships. *Harvard business review* **82**(12) 104–113.
- Loch, C.H., Y. Wu. 2008. Social preferences and supply chain performance: An experimental study. *Management Science* **54**(11) 1835–1849.

- Lovejoy, William S. 2010. Bargaining chains (long version). *Ross School of Business Paper* (1146).
- McMillan, John. 1990. Managing suppliers: Incentive systems in Japanese and US industry. *California Management Review* **32**(4) 38.
- Miller, Gary. 2001. Why is trust necessary in organizations? the moral. *Trust in society* 307.
- Miller, Gary J, Robert William Smith. 1993. *Managerial dilemmas: The political economy of hierarchy*. Cambridge University Press.
- Monczka, Robert M. 2000. *New product development: strategies for supplier integration*. ASQ Quality Press.
- Murnighan, J Keith, Alvin E Roth. 1983. Expecting continued play in prisoner's dilemma games a test of several models. *Journal of Conflict Resolution* **27**(2) 279–300.
- Nash, John F. 1950. The bargaining problem. *Econometrica: Journal of the Econometric Society* 155–162.
- Newman, Richard G. 1988. Single source qualification. *Journal of Supply Chain Management* **24**(2) 10.
- Özer, Ö., Y. Zheng, K.Y. Chen. 2011. Trust in forecast information sharing. *Management Science* **57**(6) 1111–1137.
- Özer, Özalp, Upender Subramanian, Yu Wang. 2016. Information sharing, advice provision, or delegation: what leads to higher trust and trustworthiness. *Manag Sci* .
- Özer, Özalp, Yanchong Zheng. 2016. Establishing trust and trustworthiness for supply chain information sharing. *Information Exchange in Supply Chain Management*, A. Ha, C. Tang, eds, Springer .
- Özer, Özalp, Yanchong Zheng, Yufei Ren. 2014. Trust, trustworthiness, and information sharing in supply chains bridging China and the United States. *Management Science* **60**(10) 2435–2460.
- Radner, Roy. 1985. Repeated principal-agent games with discounting. *Econometrica: Journal of the Econometric Society* 1173–1198.
- Roth, Alvin E, J Keith Murnighan. 1978. Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology* **17**(2) 189–198.
- Rubinstein, Ariel. 1979. Equilibrium in supergames with the overtaking criterion. *Journal of Economic Theory* **21**(1) 1–9.
- Rudzki, Robert A. 2004. The advantages of partnering well. *SUPPLY CHAIN MANAGEMENT REVIEW*, V. 8, NO. 2 (MAR. 2004), P. 44-51: ILL **8**(2).
- Sherstyuk, Katerina, Nori Tarui, Tatsuyoshi Saijo. 2013. Payment schemes in infinite-horizon experimental games. *Experimental Economics* **16**(1) 125–153.
- Spiliotopoulou, Eirini, Karen Donohue, Mustafa Çağrı Gürbüz. 2015. Information reliability in supply chains: the case of multiple retailers. *Production and Operations Management* .
- Stallkamp, Thomas T. 2005. *SCORE!: A Better Way to Do Business: Moving from Conflict to Collaboration*. Pearson Education.

- Swinney, Robert, Serguei Netessine. 2009. Long-term contracts under the threat of supplier default. *Manufacturing & Service Operations Management* **11**(1) 109–127.
- Taylor, Terry A, Erica L Plambeck. 2007. Supply chain relationships and contracts: The impact of repeated interaction on capacity investment and procurement. *Management science* **53**(10) 1577–1593.
- Treleven, Mark. 1987. Single sourcing: a management tool for the quality supplier. *Journal of Purchasing and Materials Management* **23**(1) 19–24.

9. Appendix

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

$$\begin{aligned} \max_{p,w} \quad & Q(p - C_m - C_s) \\ \text{s.t.} \quad & \Pi_m(p, w) = \alpha Q(p - C_m - C_s) \\ & \text{and } \Pi_s(p, w) = (1 - \alpha)Q(p - C_m - C_s) \end{aligned}$$

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 $w^* = (1 - \alpha)(p - C_m) + \alpha C_s$

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

$$\Pi_m = \Pi_s = \frac{(a - b(C_s + C_m))^2}{8b} \quad (4)$$

³¹ For more on the surplus split in case of monopolies with exogenous bargaining power, see Lovejoy (2010).

9.2. 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

$$\arg \max_{\alpha} [(u_s - t_s)(u_m - t_m)]$$

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 to the α that solves:

$$\arg \max_{\alpha} [(1 - \alpha)Q(p - C_s - C_m)][\alpha Q(p - C_s - C_m) - Q(p - C_a - C_m)]$$

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:

$$\begin{aligned} \max_{p, w} \quad & Q(p - C_m - C_s) \\ \text{s.t.} \quad & \Pi_m(p, w) = \alpha Q(p - C_m - C_s) \\ & \text{and } \Pi_s(p, w) = (1 - \alpha)Q(p - C_m - C_s) \end{aligned}$$

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^* , 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:

$$\Pi_m = \frac{[a - b(C_m + C_s)][a - b(C_m + C_a)]}{4b}, \quad (5)$$

and the supplier's profit is:

$$\Pi_s = \frac{[C_a - C_s][a - b(C_s + C_m)]}{4}. \quad (6)$$

9.3. 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) as in Table 8 and calculate the payoffs under these conditions. Figure 1 shows the extensive form of the game with these payoffs.

Suppose that supplier 1 started off with production cost $C_{s1} = 7$ and the manufacturer 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 nothing³². 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 56.25 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

³² We assume that if supplier 1 and the alternative supplier have the same costs, supplier 1 wins the deal.

optimal retailer 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).

9.4. Proofs for Equilibria, $E1$, $E2$, $E3$ - Buyer-as-Single-Employee setup:

We show that $E1$, $E2$, and $E3$ can be sustained in equilibrium with the parameters used in our setup (see Table 8), and the resulting payoffs of the stage game depicted in Figure 1.

9.4.1. Equilibrium $E1$: In equilibrium $E1$ the buyer is skeptical and the supplier is forgiving. On equilibrium, they stay in the collaborative state where the supplier shares and the buyer chooses “not compete if the supplier shared” and “compete if the supplier did not share”. In particular in this equilibrium, if the innovation does not happen and the buyer chooses “compete”, they stay in the collaborative state. If anyone deviates, they switch to a punishment state (off-equilibrium) where the supplier never shares and the buyer always chooses “compete”.

We check that no player wants to deviate in the **collaborative state**:

E1.a) The supplier does not want to deviate from “share” as long as $ISN_s + \frac{\delta}{(1-\delta)}[\pi ISN_s + (1-\pi)NC_s] \geq INC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. Or equivalently, $ISN_s \geq INC_s$, which holds with our payoffs.

E1.b) The buyer does not want to deviate from “not compete” if the supplier shared as long as $ISN_m + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1-\pi)NC_m] \geq ISC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$, or equivalently, as long as $\delta \geq \frac{ISC_m - ISN_m}{[ISC_m - \pi INC_m - (1-\pi)ISN_m]}$. With the values of the payoffs of our stage game, this holds for $\delta \geq 0.69$.

E1.c) The buyer does not want to deviate from “compete” if the supplier did not share as long as $NC_m + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1-\pi)NC_m] \geq NN_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$. The inequality holds with our payoffs since $NC_m > NN_m$ and $ISN_m > INC_m$.

We now check that no player wants to deviate in the **punishment state**:

E1.d) The supplier does not want to deviate from “not share” as long as $INC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s] \geq ISC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. The inequality always holds with our payoffs since $INC_s > ISC_s$.

E1.e) The buyer does not want to deviate from “compete” as long as $[\pi INC_m + (1 - \pi)NC_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1 - \pi)NC_m] \geq [\pi INN_m + (1 - \pi)NN_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1 - \pi)NC_m]$. The inequality holds with our payoffs since $INC_m > INN_m$ and $NC_m > NN_m$.

9.4.2. Equilibrium E2: In equilibrium *E2*, the buyer is non-skeptical and the supplier is non-forgiving. On equilibrium, they stay in the collaborative state where the supplier shares and the buyer chooses “not compete if the supplier shared” and “not compete if the supplier did not share”. In particular in this equilibrium, they stay in the collaborative state unless the supplier does not share the innovation or the buyer chooses to compete. If anyone deviates, they switch to a punishment state (off-equilibrium) where the supplier never shares and the buyer always chooses “compete”.

We first check that no player wants to deviate in the **collaborative state**:

E2.a) The supplier does not want to deviate from “share” as long as $ISN_s + \frac{\delta}{(1-\delta)}[\pi ISN_s + (1 - \pi)NN_s] \geq INN_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1 - \pi)NC_s]$. The inequality always holds with our payoffs since $ISN_s > INN_s$, $ISN_s > INC_s$, and $NN_s > NC_s$.

E2.b) The buyer does not want to deviate from “not compete” if the supplier shared as long as $ISN_m + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1 - \pi)NN_m] \geq ISC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1 - \pi)NC_m]$ or equivalently, as long as $\delta \geq \frac{ISC_m - ISN_m}{[ISC_m - (1-\pi)ISN_m + (1-\pi)NN_m - \pi INC_m - (1-\pi)NC_m]}$. With our payoffs, this is equivalent to $\delta \geq 0.70$.

E2.c) The buyer does not want to deviate from “not compete” if the supplier did not share as long as $NN_m + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1 - \pi)NN_m] \geq NC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1 - \pi)NC_m]$ or equivalently, as long as $\delta \geq \frac{NC_m - NN_m}{\pi[ISN_m - INC_m - NN_m + NC_m]}$. With our payoffs, this is $\delta \geq 0.171$.

We now check that no player wants to deviate in the **punishment state**:

E2.d) The supplier does not want to deviate from “not share” as long as $INC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1 - \pi)NC_s] \geq ISC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1 - \pi)NC_s]$. With our payoffs, the inequality holds since $INC_s > ISC_s$.

E2.e) The buyer does not want to deviate from “compete” as long as $[\pi INC_m + (1 - \pi)NC_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1 - \pi)NC_m] \geq [\pi INN_m + (1 - \pi)NN_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1 - \pi)NC_m]$. With our payoffs, the inequality holds since $INC_m > INN_m$ and $NC_m > NN_m$.

9.4.3. Equilibrium E3: In equilibrium *E3*, the buyer is skeptical and the supplier is non-forgiving. On equilibrium, there is a collaborative state where the supplier shares and the buyer chooses “not compete if the supplier shared” and “compete if the supplier did not share”. They stay in the collaborative state unless the buyer chooses “compete” or the supplier does not share. In the punishment state, the supplier never shares and the buyer always competes. Note that the buyer choosing “compete” happens with positive probability in the collaborative state (with probability π , the innovation will not happen and the buyer will choose “compete”). Thus, this is

not a fully collaborative equilibrium, as there is a positive probability that they will transition from the collaborative state to the punishment state.

First, we define V_i^j as the expected value of staying in state i , $i \in \{c = \text{collaborative}, p = \text{punishment}\}$, as role j , $j \in \{s = \text{supplier}, m = \text{manufacturer}\}$. We compute for each role, the **expected value of staying in the punishment state**:

$$\text{Supplier: } V_p^s = [\pi INC_s + (1 - \pi)NC_s] + \delta V_p^s, \text{ or equivalently } V_p^s = \frac{1}{(1-\delta)}[\pi INC_s + (1 - \pi)NC_s].$$

$$\text{Manufacturer: } V_p^m = [\pi INC_m + (1 - \pi)NC_m] + \delta V_p^m \text{ or equivalently, } V_p^m = \frac{1}{(1-\delta)}[\pi INC_m + (1 - \pi)NC_m].$$

We now compute for each role, the **expected value of staying in the collaborative state**:

$$\text{Supplier: } V_c^s = \pi(ISN_s + \delta V_c^s) + (1 - \pi)(NC_s + \delta V_p^s) \text{ or equivalently, } V_c^s = \frac{\pi ISN_s + (1-\pi)NC_s + \frac{\delta}{(1-\delta)}(1-\pi)[\pi INC_s + (1-\pi)NC_s]}{(1-\delta\pi)}.$$

$$\text{Manufacturer: } V_c^m = \pi(ISN_m + \delta V_c^m) + (1 - \pi)(NC_m + \delta V_p^m) \text{ or equivalently, } V_c^m = \frac{\pi ISN_m + (1-\pi)NC_m + \frac{\delta}{(1-\delta)}(1-\pi)[\pi INC_m + (1-\pi)NC_m]}{(1-\delta\pi)}.$$

For our values of payoffs, the results above imply that $V_c^s > V_p^s$ for all $\delta \in [0, 1)$, and $V_c^m > V_p^m$ for all $\delta \in [0, 1)$.

We check that no player wants to deviate in the **collaborative state**:

E3.a) The supplier does not want to deviate from “share” as long as $ISN_s + \frac{\delta}{(1-\delta)}V_c^s \geq INC_s + \frac{\delta}{(1-\delta)}V_p^s$. With our payoffs, the inequality always holds.

E3.b) The buyer does not want to deviate from “not compete” if the supplier shared as long as $ISN_m + \frac{\delta}{(1-\delta)}V_c^m \geq ISC_m + \frac{\delta}{(1-\delta)}V_p^m$. With our payoffs, the inequality holds for $\delta \geq 0.56$.

E3.c) The buyer does not want to deviate from “compete” if the supplier did not share as long as $NC_m + \frac{\delta}{(1-\delta)}V_p^m \geq NN_m + \frac{\delta}{(1-\delta)}V_p^m$. With our payoffs, the inequality always holds.

We now check that no player wants to deviate in the **punishment state**:

E3.d) The supplier does not want to deviate from “not share” as long as $INC_s + \frac{\delta}{(1-\delta)}V_p^s \geq ISC_s + \frac{\delta}{(1-\delta)}V_p^s$. With our payoffs, the inequality always holds.

E3.e) The buyer does not want to deviate from “not compete” as long as $[\pi INC_m + (1 - \pi)NC_m] + \frac{\delta}{(1-\delta)}V_p^m \geq [\pi INN_m + (1 - \pi)NN_m] + \frac{\delta}{(1-\delta)}V_p^m$. With our payoffs, the inequality always holds.

9.5. Proofs for Equilibria, $E1'$, $E2'$, $E3'$ - Buyer-as-Two-Employees setup (50-50 case):

In the Buyer-as-Two-Employees setups, the supplier and manufacturer have a long-term relationship, and the manufacturer has two employees (a long term engineer and a short term procurement manager). We assume that both employees make recommendations for what the manufacturer should do. In the 50 – 50 case, if both employee’s recommendations agree, their recommendation is implemented. If they disagree, one of the two is implemented at random, both with equal probability.

We show that three equilibria $E1'$, $E2'$, and $E3'$, analogous to the equilibria in the Buyer-as-Single-Employee setup, $E1$, $E2$, and $E3$, can be supported in the Buyer-as-Two-Employees 50 – 50 setup under the values of parameters used in the experiment.

In these equilibria, the supplier plays the same strategies as in the Buyer-as-a-Single-Employee long-run case. The engineer plays the same strategies as the manufacturer in the Buyer-as-a-Single-Employee long-run case, and the procurement manager always plays the non-collaborative strategy of the single-period game. The procurement manager's incentive to always play the non-collaborative strategies -compete if the supplier shared and compete if the supplier did not share or the innovation did not happen- are straight forward: since the procurement manager works for the manufacturer for only one round, his incentives are those of a single period game. In each round of play of the stage game, the procurement manager's recommendation is implemented with probability 0.5, however, the procurement manager's recommendation never triggers abandonment of the collaborative state. We consider equilibria where the supplier's and engineer's decisions to stay in the collaborative state depend only on the supplier's and engineer's previous actions.

The proof shows that the condition for the supplier not to leave the collaborative state is tighter than in the Buyer-as-Single-Employee case. This is because in the collaborative state, the procurement manager's non-collaborative action will be implemented with probability 0.5. Thus, the supplier's expected payoff from collaboration is lower than in the Firms-as-a-Single-Employee case. On the contrary, the condition for the engineer not to leave the collaborative state is less tight than in the Buyer-as-a-Single-Employee state. This is because the engineer benefits from the increased payoff derived from the procurement manager's recommendation to compete without facing the supplier's punishment. Thus, any δ that guarantees that the buyer does not deviate from the collaborative state in the Firms-as-a-Single-Employee case, will also guarantee that the engineer does not want to deviate from the collaborative state in the Firms-as-Two-Employees case.

9.5.1. Equilibrium $E1'$: We check that no player wants to deviate in the **collaborative state**:

E1'.a) The supplier does not want to deviate from sharing in the collaborative state as long as $\frac{ISN_s+ISC_s}{2} + \frac{\delta}{(1-\delta)}[\pi\frac{ISN_s+ISC_s}{2} + (1-\pi)NC_s] \geq INC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. With our payoffs, the inequality holds for every δ since $\frac{ISN_s+ISC_s}{2} > INC_s$.

E1'.b) The engineer does not want to deviate from “not compete” if the supplier shared as long as $ISN_m + \frac{\delta}{(1-\delta)}[\pi(\frac{ISN_m+ISC_m}{2}) + (1-\pi)NC_m] \geq ISC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$, or equivalently $\delta \geq \frac{ISC_m-ISN_m}{(\frac{\pi}{2}-1)ISN_m+(\frac{\pi}{2}+1)ISC_m-\pi INC_m} = 0.55$.

E1'.c) The engineer does not want to deviate from “compete” if the supplier did not share as long as $NC_m + \frac{\delta}{(1-\delta)}[\pi(\frac{ISN_m+ISC_m}{2}) + (1-\pi)NC_m] \geq NN_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$. With our payoffs, the inequality holds for all δ since $NC_m > NN_m$ and $\frac{ISN_m+ISC_m}{2} > INC_m$.

As explained above, the conditions for the engineer to remain in the collaborative state are less tight than in the Firms-as-a-Single-Employee case. Thus, any δ that guarantees that the manufacturer does not want to deviate from the collaborative state in the Firms-as-a-Single-Employee case, also guarantees that the engineer does not want to deviate from the collaborative state in the Firms-as-Two-Employees 50 – 50 case. Also note that the conditions for the supplier and the engineer not to deviate from the **punishment state** are as in the Firms-as-a-Single-Employee case and thus will be omitted.

9.5.2. Equilibrium $E2'$: We check that no player wants to deviate in the **collaborative state**:

E2'.a) The supplier shares in the collaborative state as long as $\frac{ISN_s+ISC_s}{2} + \frac{\delta}{(1-\delta)}[\pi(\frac{ISN_s+ISC_s}{2}) + (1-\pi)(\frac{NN_s+NC_s}{2})] \geq \frac{INN_s+INC_s}{2} + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. The inequality holds for all δ since $\frac{ISN_s+ISC_s}{2} > \frac{INN_s+INC_s}{2} > INC_s$ and $\frac{NN_s+NC_s}{2} > NC_s$.

E2'.b) The engineer does not want to deviate from “not compete” if the innovation was shared as long as $ISN_m + \frac{\delta}{(1-\delta)}[\pi(\frac{ISN_m+ISC_m}{2}) + (1-\pi)(\frac{NN_m+NC_m}{2})] \geq ISC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$ or equivalently, as long as $\delta \geq \frac{ISC_m - ISN_m}{[(\frac{\pi}{2}-1)ISN_m - (\frac{\pi}{2}+1)ISC_m - \pi INC_m + \frac{(1-\pi)}{2}NN_m + \frac{(\pi-1)}{2}NC_m]}$. With our payoffs, this is equivalent to $\delta \geq 0.55$.

E2'.c) The engineer does not want to deviate from “not compete” if the innovation was not shared as long as $NN_m + \frac{\delta}{(1-\delta)}[\pi(\frac{ISN_m+ISC_m}{2}) + (1-\pi)(\frac{NN_m+NC_m}{2})] \geq NC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$ or equivalently, as long as $\delta \geq \frac{NC_m - NN_m}{[\frac{\pi}{2}ISN_m + \frac{\pi}{2}ISC_m - \pi INC_m + \frac{(\pi-1)}{2}NN_m + \frac{(\pi+1)}{2}NC_m]}$. With our payoffs, this is $\delta \geq 0.08$.

9.5.3. Equilibrium $E3'$: We first compute the **supplier's expected value of staying in the punishment and collaborative states**;

$$\text{Punishment: } V_p^s = \frac{1}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s].$$

$$\text{Collaboration: } V_c^s = \frac{\pi(\frac{ISN_s+ISC_s}{2}) + (1-\pi)NC_s + \frac{\delta}{(1-\delta)}(1-\pi)[\pi INC_s + (1-\pi)NC_s]}{(1-\delta\pi)}.$$

We can see that $V_c^s > V_p^s$ for all $\delta \in [0, 1)$. We now check that the supplier does not want to deviate from the collaborative state:

E3'.a) The supplier chooses share in the collaborative state as long as $\frac{ISN_s+ISC_s}{2} + \frac{\delta}{(1-\delta)}V_c^s \geq INC_s + \frac{\delta}{(1-\delta)}V_p^s$, which holds for all δ .

We next compute the **engineer's expected value of staying in the punishment and collaborative states**;

$$\text{Punishment: } V_p^m = \frac{1}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m].$$

$$\text{Collaboration: } V_c^m = \frac{\pi(\frac{ISN_m+ISC_m}{2}) + (1-\pi)NC_m + \frac{\delta}{(1-\delta)}(1-\pi)[\pi INC_m + (1-\pi)NC_m]}{(1-\delta\pi)}.$$

We can see that $V_c^m > V_p^m$ for all $\delta \in [0, 1)$. We now check that the engineer does not want to deviate from the collaborative state:

E3'.b) The engineer chooses “not compete” if the innovation was shared in the collaborative state as long as $ISN_m + \frac{\delta}{(1-\delta)}V_c^m \geq ISC_m + \frac{\delta}{(1-\delta)}V_p^m$. With our payoffs, the inequality holds for $\delta \geq 0.45$.

E3'.c) The engineer does to want to deviate from “compete” if the innovation was not shared in the collaborative state as long as $NC_m + \frac{\delta}{(1-\delta)}V_p^m \geq NN_m + \frac{\delta}{(1-\delta)}V_p^m$. With our payoffs, the inequality holds for all δ .

9.6. Additional Tables and Figures

Table 8 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 3 Payoffs table shown in the experiment

Table 9 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
Expected profit from non-collaborative strategies (*)	15.25	20.75	36	42.4

(*) Refers to the outcome of the strategies where suppliers never choose “share” and buyers always choose “compete”.

Table 10 Profits

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

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 11 Supplier's profit from sharing - PM treatment

Coefficients	Average profit per round (Supplier)
Average sharing in a relationship	5.043*** (1.950)
Relationship number	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: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 12 Supplier's sharing (50 – 50 Treatment)

Coefficients	Supplier's Share (50 – 50 Treatment)
Previous round share	0.485*** (0.140)
Previous round compete (Eng)	-1.125*** (0.177)
Prev_Compete_Eng x Prev_Eng_Impl	-0.895*** (0.218)
Previous round compete (PM)	0.012 (0.178)
Prev_Compete_PM x Prev_PM_Impl	-1.021*** (0.203)
Period	-0.019** (0.008)
Round	0.015 (0.033)
Constant	0.870*** (0.251)
Observations	687
Nr. of Subjects	31

Probit regression with subject random effects. Standard errors reported in parentheses. The first round of each relationship was omitted. The variable “Prev_Compete_Eng x Prev_Eng_Impl” captures the interaction effect of the engineer’s decision in the previous round and whether his recommendation was implemented. “Prev_Compete_PM x Prev_PM_Impl” is the analogous for the procurement manager. Significance is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 13 Frequency of matching between suppliers' and buyers' Defector/Non-Defector types

		Buyer	
		Defector	Non-Defector
Supplier	Defector	5.8%	11.5%
	Non-Defector	33.4%	49.3%

Frequency of matching between a supplier and a buyer type as a fraction of all the relationships in the Eng and LR treatments (in total 452 relationships). Note that this indicates the frequency with which the different types are matched, not the frequency with which each equilibria occurred.

Table 14 Frequency of matching between suppliers' and buyers' types conditional on both players being Non-Defectors


		Buyer	
		Skeptical	Non-Skeptical
Supplier	Forgiving N/A*	6.3%	20.6%
	Non-forgiving	21.5%	20.6%
	Forgiving	16.2%	14.8%

Frequency of matching between a supplier and a buyer type in the Eng and LR treatments conditional on buyer and supplier being non-defectors (in total 223 relationships). (*) “Forgiving N/A” refers to non-defector suppliers who where not faced with a situation where they had to be either forgiving or non-forgiving.

Table 15 Effect of types on other player's profit

	Supplier's Average Profit in the Relationship			
	Buyer Defector	Buyer Non-Defector	Buyer Skeptical	Buyer Non-skeptical
All suppliers	11.49	25.95***	22.64	28.74***
Supplier Defector	14.38	18.64**	16.77	20.83**
Supplier Non-defector	10.99	27.65***	24.32	30.26***
Supplier Forgiving	10.23	32.03***	27.93	36.52
Supplier Non-Forgiving	11.73	26.53***	25.09	28.04
	Buyer's Average Profit in the Relationship			
	Supplier Defector	Supplier Non-Defector	Supplier Forgiving	Supplier Non-Forgiving
All buyers	23.41	42.85***	46.27	40.65*
Buyer Defector	21.89	45.45***	47.15	42.94
Buyer Non-defector	24.17	41.09***	45.67	39.30
Buyer Skeptical	24.52	41.45***	48.80	34.85**
Buyer Non-skeptical	23.76	40.80***	42.26	43.94

Note: Each relationship is one observation. Columns 3 and 4 include only relationships where the supplier and the buyer are not defectors. Significance in the difference between columns (1) and (2) and between columns (3) and (4) is indicated in columns (2) and (4) respectively. It is denoted: * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

**Figure 4****Figure 5** Compete if Shared by Treatment

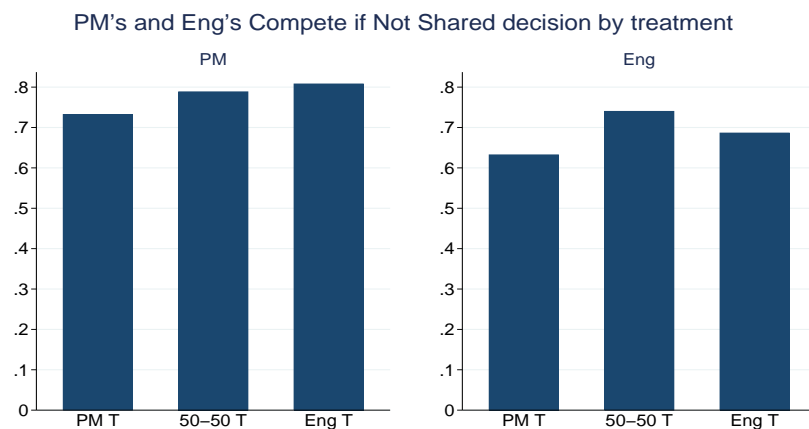


Figure 6

Figure 7 Compete if Not Shared by Treatment

Supplier's and buyer's types distributions

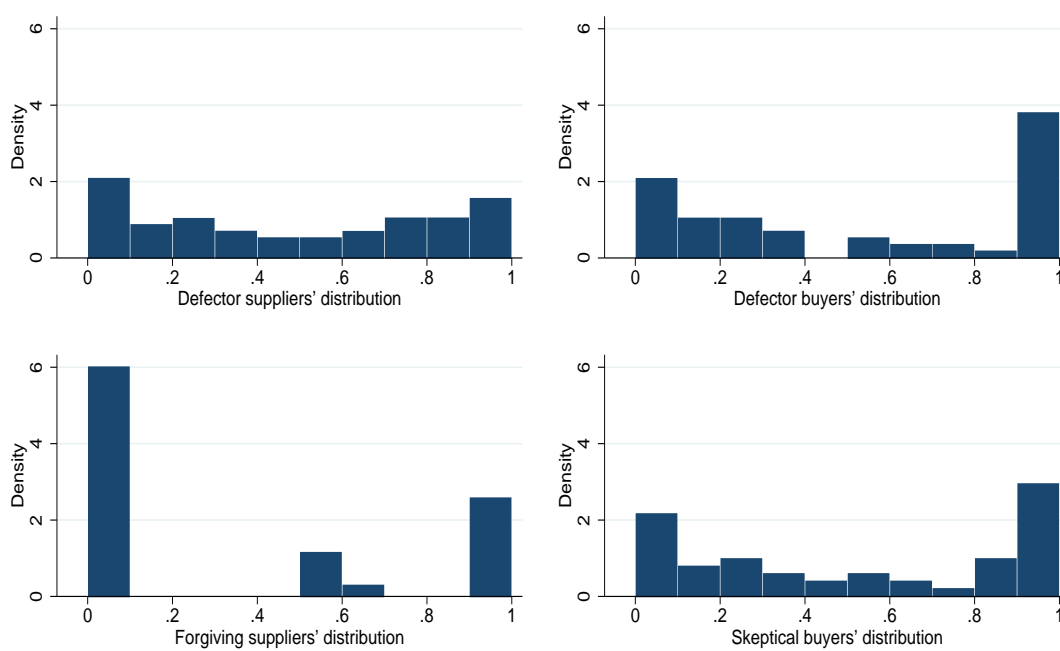


Figure 8 Buyers' and Suppliers' Types (Eng and LR treatments)




Figure 9 Engineer's Decision