M MICHIGAN ROSS

Working Paper

The Impact of Decision Rights on Innovation Sharing

Ruth Beer

Kelley School of Business Indiana University

Hyun-Soo Ahn

Stephen M. Ross School of Business University of Michigan

Stephen Leider

Stephen M. Ross School of Business University of Michigan

Ross School of Business Working Paper Working Paper No. 1369 December 2018

This paper can be downloaded without charge from the Social Sciences Research Network Electronic Paper Collection: http://ssrn.com/abstract=2967525

UNIVERSITY OF MICHIGAN

The Impact of Decision Rights on Innovation Sharing

Ruth Beer

Kelley School of Business, Indiana University, Bloomington, IN 47404, 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—a practice that is common in the automotive industry—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). Anecdotal evidence from automotive suppliers tells that in some occasions the buyer's decision is in the hands of a long-run focused employee (an "engineer"), while in other occasions it is a short-run focused employee (a "procurement manager") who has more control. In this paper, we examine how the allocation of decision rights to short-run and long-run focused employees affects collaboration between the two firms. To do so, we model a relationship between a supplier and a buyer where the buyer is a dual decision maker, consisting of long-run and/or short-run focused employees. We characterize the equilibrium of this model and show that the frequency of collaborative outcomes increases from a case where the decision is made by an employee with a short-term objective, to a case where the decision is made jointly, to a case where a decision is made by an employee with a long-run objective. Our experimental results verify this prediction, for the most part. One result not predicted by the theory is that when decision rights are equally split between the two employees, both employees become significantly less trustworthy. With an additional treatment which allows for free-form communication between the employees, we identify social interaction effects in the form of a "bias to agreement" as a plausible driver of the more competitive behavior. We find that 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.

1. Introduction

Manufacturers often benefit from innovations and process improvements discovered by their suppliers, 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 reduction comes from

tier-1 or tier-2 suppliers.¹ For example, General Motors' suppliers developed hinges that did not need welding, which resulted in significant cost savings for GM (Klier 2006). Similarly, Chrysler achieved 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 frame.²

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 need 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 many 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). 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). To reverse this deeply rooted mistrust, recent initiatives, such as Ford's "Aligned Business Framework" and GM's "Strategic Supplier Engagement" programs, have focused on building longer-run strategic partnerships with suppliers and encouraging innovation sharing. While this has led to recent improvements in measures of the

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.

quality of the working relationship between American automakers and their suppliers, their ratings are still quite poor.⁵

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 gains. 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. For example, for some divisions within GM, the responsibility for managing the relationships with suppliers is primarily with the procurement managers, while in others engineers have extensive control. Conversations with GM's suppliers suggest that the nature of a buyer-supplier relationship heavily depends on which employees within the firm manages the relationship.⁷ Procurement managers are often evaluated by performance metrics that focus on short-run (immediate) cost savings. Thus, their incentives and promotions are driven by these cost measures. This is further exacerbated by the fact that, in many organizations, procurement managers have a short tenure as they 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: Does the allocation of decision-making rights (e.g., to a short-sighted procurement manager or a long run focused engineer) affect the likelihood of a collaborative relationship? What are the behavioral mechanisms driving the buyer's and supplier's strategies under different allocations of decision rights?

To answer these questions, we focus on three focal cases: a case where the procurement manager has full control, a case where the engineer has full control, and a joint decision making case where decision rights are equally split between the two employees. Our theoretical results predict that there is a trend of increasing collaboration from the procurement manager control case, to the joint control case, to the engineer control case. The experimental results confirm this prediction, for the most part. One result not predicted by the theory is that when there is uncertainty about which employee's recommendation will be implemented (the joint decision rights case), both employees are significantly less trustworthy. We propose two potential explanations for this behavior and test them with additional experimental treatments. The first one is diffusion of responsibility, by which an engineer may feel less responsible to take the collaborative action when the procurement manager is part of the decision making process (and thus, also able to help the supplier). To test this, we designed an additional treatment which favors diffusion of responsibility: if at least one employee recommends to take the collaborative action, then that decision is implemented. If neither chooses to be collaborative, the non-collaborative action to open up competition is implemented. Since each employee can unilaterally implement the collaborative action, this modified decision making mechanism should lead to less collaborative behavior by the engineer as it alleviates the engineer's own responsibility for the final decision. On the contrary, we find that this treatment does not change an engineer's behavior significantly, suggesting that diffusion of responsibility is not a main driver of the less collaborative behavior in the joint control case. The second potential explanation are social interaction effects by which in the joint control case the employees derive utility from making their recommendations match. We test this behavioral mechanism with an additional treatment

where employees are given the option to chat before making their recommendations. This treatment keeps the same decision making mechanism as the join decision case but it increases the employees' ability to agree by allowing for free-form communication. We find that most employees choose to engage in the chats and try to reach an agreement with the other employee (despite having no strategic incentives to do so). In addition, chat analysis suggests that the concerns expressed by the employees make agreeing to compete easier than agreeing not to compete, and that an agreement to compete is more likely to be honored. This suggests that social interaction effects in the form of a bias to agreement are a more plausible explanation for the increased competitive behavior in the joint control case.

Finally, we compare the procurement manager and engineer full control cases with the theoretically-equivalent benchmarks where the buyer is a single decision maker and either has a short-run or a long-run relationship with the supplier. We find that there are no significant differences between the case where the engineer has full control and the case where the buyer is a single decision maker and has a long-term relationship with the supplier. Contrary to the theoretical prediction, the supplier shares the innovation significantly more often when the procurement manager has full control than when the buyer is a single decision maker and has a short-term relationship. Our analysis finds that this can be explained by strategic motives given the buyer's more collaborative behavior. Overall, we find that 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. 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.

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 the 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 relations, relational contracts provide

enough incentive for the supplier to invest. Similarly, Li and Debo (2009) 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 long-term supplier relations experimentally (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 2017, 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. Their results suggest that a 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) find that repeated interactions strengthen the impact of signaling investments, leading to higher profits and efficiency. Davis and Hyndman (2016) show that a long term relationship with a threat of punishment leads to increased quality and supply chain efficiency. Hyndman et al. (2014) find that fixed pairs reach a better alignment to coordinate on an optimal total investment level than pairs that are randomly re-matched after every round, while Hyndman and Honhon (2014) find 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 in the context of innovation sharing. Second, we introduce joint decision making within the buyer. Our focus is on the firms' allocation of decision rights to a short-term focused procurement manager and a long-run focused engineer. 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 discover an innovation (the innovation happens randomly with exogenous probability π) and, if so, she needs to decide whether to share it with a buyer or not. The buyer then chooses whether to elicit bids from new suppliers to lower the price (we call this decision "to compete") or to single source and offer the supplier a generous price ("not to compete").

As Figure 1 shows, the stage game has six possible outcomes: If the innovation occurs (I), 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 (N), the possible outcomes are Compete (NC) and Do not Compete (NN). The payoffs in Figure 1 (which are later used in the experiment) are derived from a bargaining model presented in section OA-1 of the Online Appendix.⁸

⁸ The model applies the Nash bargaining solution concept to split the supply chain surplus between the supplier and the buyer (Lovejoy 2010). We analyze separately the case where the buyer chooses not to compete (referred in the literature as bilateral bargaining) and the case where the buyer chooses to compete (bargaining with supplier competition). We then apply a plausible set of parameters to the bargaining model, and derive the values of the stage game payoffs presented in Figure 1.

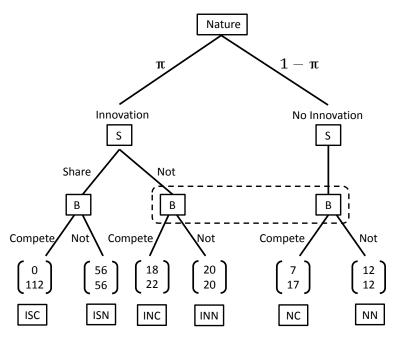


Figure 1 Stage Game

With these payoffs, if the innovation occurs, the total surplus (the size of a pie that can be shared between the two parties) increases relative to the case where the innovation does not occur. In addition, if the supplier shares the innovation with the buyer, the total surplus increases even further. However, sharing the innovation makes the supplier more vulnerable to competition. Specifically, the minimum payoff from not sharing is seven and from sharing it is zero. The buyer'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 buyer choosing not to compete as "collaborative" actions since both firms benefit from their counterpart's action. Since the innovation occurs randomly in each period, when choosing his action the buyer cannot distinguish between the case where the innovation has occurred but the supplier decided not to share it and the case where the innovation has not occurred in the first place. This is captured in the information set of the buyer (from now onwards, we will denote this decision node as the buyer's "ambiguous node"). When the outcomes are realized, the buyer learns whether an innovation had occurred, and therefore whether the supplier had shared (capturing, for example, a case where the buyer learns this information while bargaining with the supplier). We make this simplification so that a buyer can condition his action directly on the supplier's action in the previous round. Without this simplification, the buyer 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. The laboratory experiment reproduces the same information structure.⁹

It is straightforward to see that the only subgame perfect Nash equilibrium of the single-period play of the stage game in Figure 1 is the non-collaborative outcome. To see why, we solve by backward induction, starting with the buyer's strategy. In all decision nodes, the buyer's payoff is higher when he chooses to compete than when he chooses not to compete (112 vs. 56 if the supplier shared, 22 vs. 20 if the supplier did not share, and 17 vs. 12 if the innovation did not occur). We then roll back to the supplier's strategy given that the buyer chooses "compete". If the supplier chooses to share the innovation, she earns zero profit. If she does not share, then she earns positive profits. As a result, the supplier does not share and the only Nash equilibrium in a one-period play of the game are INC if the innovation occurs, and NC if the innovation does not occur.

3.1. Firms' Long-Term Relationship

To capture the long-term relationship between the two firms we 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 δ today and a dollar to be received n periods from today is worth δ^n today. 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). 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, $i \in \{s = \text{supplier}, b = \text{buyer}\}$, from each possible outcome. Additionally, recall that at the end of the stage game, the buyer learns whether the innovation occurred and whether the supplier shared. Therefore, the buyer can condition his strategies on the past sharing decisions of the supplier. Similarly, the supplier can condition her strategies on the buyer's past decisions to compete.

We characterize collaborative equilibria of the infinitely repeated game—that is equilibria where the supplier shares the innovation, and the buyer does not compete when the innovation is shared.

⁹ 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 buyer is preceded by a random innovation. Second, the buyer 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.

¹⁰ Recent experimental work by Fréchette and Yuksel (2017) 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.

Since we are especially interested in those equilibria with the highest frequency of the collaborative "Share-Not Compete" outcome, we focus on trigger strategies which provide the highest disincentive to deviate from collaboration.¹¹ While the Folk theorem for infinitely repeated games implies that many strategies can support equilibria with collaborative outcomes, trigger strategies provide the largest set of parameters under which collaboration can be sustained in equilibrium.¹² In addition, we focus only on pure strategies that lead to an equilibrium with high collaboration rates by both firms. In mixed strategies, the buyer could induce the supplier to share by using, for example, a strategy where he chooses not to compete 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. From all these plethora of equilibria, we are interested in those with the highest collaboration rates from both the supplier and the buyer.

Buyers-as-Two Employees Setup. Our main focus is a "Buyer-as-Two-Employees" setup, where the buyer's decision is made by a procurement manager and an engineer in a long-term relationship. The procurement manager works for the team for only one period (he is assigned to manage the relationship with a supplier for one period, and rotates to another team in the next period). The engineer works for the firm (and is assigned to this supplier) throughout the infinitely repeated game. We further assume that both employees make recommendations for what the buyer should do, and that their compensation is equal to the buyer'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. We consider an allocation of decision rights where, if both recommendations agree, that decision is implemented. If they disagree, the buyer implements one of the two recommendations at random: the procurement manager's recommendation with probability p and the engineer's recommendation with probability (1-p), where p is known to all players. We are interested in three focal cases: when p=1 and the procurement manager has full control; when p=0 and the engineer has full control; and when p = 0.5 and both employees have joint control with the decision rights being equally split between the two—if their recommendations disagree, both recommendations are equally likely to be implemented. This last case captures a setting where the employees are uncertain about which of the two recommendations a senior manager will find most persuasive.

¹¹ While the analysis of the strategies played is not the main focus in this paper, we observe trigger-like behavior by many subjects in the experiment. A summarized strategy analysis is presented in section OA-3 of the Online Appendix.

¹² 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).

3.1.1. Procurement Manager's Full Control Case To analyze the case where the procurement manager has full control (p=1), consider 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. Again, by backward induction, the procurement manager always recommends to compete (and the buyer always implements this recommendation) and the supplier always chooses not to share. Thus, 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.

3.1.2. Engineer's Full Control and Joint Control Cases With the payoffs in Figure 1 and $\pi = 0.75$, which we later use in the experiment, collaborative equilibria can be supported when p = 0 and the engineer has full control, and when p = 0.5 and the decision rights equally split between the two employees.

Consider first the case where the engineer has full control. Since the engineer works for the buyer indefinitely and the firms have a long-term relationship, the game between the supplier and the engineer resembles an infinitely repeated game. In Proposition 1, we show that when the engineer's recommendation is always implemented, trigger strategies can sustain a repetition of the collaborative outcome Share-Not Compete (ISN) in every period where there is an innovation.

Consider now the case where if both employees' recommendations disagree, one of the two recommendations is implemented at random, both with equal probability. For simplicity, we assume that the supplier can perfectly observe both employees' recommendations. This could represent either the supplier directly observing the buyer's decision-making, or the engineer being able to credibly verify his recommendation.¹³ In Proposition 1 we show that in this case, trigger strategies—where the supplier and the engineer condition their actions on the supplier's and the engineer's previous actions—can sustain a collaborative equilibrium where the supplier always shares the innovation and where the collaborative outcome Share-Not Compete (ISN) occurs whenever the engineer's recommendation is implemented.

The equilibrium with high collaboration rates characterized in Proposition 1 is supported with trigger strategies, where the supplier and the engineer choose the collaborative action until either one 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. Therefore, if either player defects, the punishment state will commence in the next period and will continue for the rest of the game. Within collaborative trigger strategies, there is some flexibility during

¹³ 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.

the stage game in how the buyer should handle the ambiguous decision node where no innovation was shared—as this could arise either from the supplier defecting by not sharing an innovation that occurred or from the supplier not having an innovation. Proposition 1 describes the sufficient condition for two different collaborative equilibria to arise. The first is equilibrium E1, where the engineer chooses to compete in the ambiguous node during the stage game, but tries to return to the collaborative state if the supplier was innocent (i.e. had no innovation), and the supplier is willing to stay in the collaborative state. The second is equilibrium E2, where the engineer chooses not to compete in the ambiguous node during the stage game (and waits for the end of the period for confirmation of the supplier's possible defection), and the supplier enters the punishment state if the buyer chooses to compete when there was no innovation. Note that, even though these equilibria differ in the strategies played in the ambiguous node, both equilibria can sustain the most efficient outcome where the supplier always shares whenever an innovation occurs. At the same time, both equilibria have the highest possible collaboration rate from the buyer when an innovation is shared.

In preparation for Proposition 1, let us define:

$$\hat{\delta} = \frac{ISC_b - ISN_b}{[(\pi(1-p)-1)ISN_b + (\pi p+1)ISC_b - \pi INC_b + (1-\pi)(1-p)NN_b + (1-\pi)(p-1)NC_b]}$$

PROPOSITION 1. [Collaborative Equilibria - Buyer-as-Two-Employees (engineer control and joint control cases):] If $\delta \geq \hat{\delta}$, there exist an equilibrium in the form of E1 and an equilibrium in the form of E2, that support sharing in every round, and "not compete" in every round where the engineer's recommendation is implemented.

The detailed proofs are presented in section 7.1 in the Appendix. The proofs imply that, with the payoffs in Figure 1 and $\pi=0.75$ (which we later use in the experiment), each equilibrium can be sustained in the engineer's full control (p=0) and the joint control (p=0.5) cases. Specifically, $\delta \geq 0.70$ guarantees that E1 and E2 can arise when the engineer's recommendation is always implemented, and $\delta > 0.55$ guarantees that E1 and E2 can arise when both employees' recommendations are equally likely to be implemented upon disagreement. Based on the two interpretations of δ , the conditions mean that the buyer 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 arise.

¹⁴ The proofs in the Appendix show that the condition for the supplier not deviating from collaboration is tighter in the joint control case than in the engineer control case. This is because in the joint control case, the procurement manager's non-collaborative action is implemented with probability 0.5, resulting in a lower expected payoff from collaboration for the supplier than the engineer control case. On the contrary, the condition for the engineer not leaving the collaborative state is less tight in the joint control case than in the engineer control case. 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 engineer does not deviate from the collaborative state in the engineer control case, also guarantees that the engineer does not want to deviate from the collaborative state in the joint control case.

Finally note that in the simpler case where we consider only the relationship between the supplier and the buyer without allocation of decision rights on the buyer side ("Buyer-as-Single-Employee" case), we can identify two benchmark cases. The lowest collaboration benchmark is the case where the firms transact only once and the only Nash equilibrium is the non-collaborative outcome of the stage game. This case is analogous to the case where the procurement manager has full control. The highest collaboration benchmark is the case where the two firms have an infinitely long relationship. In this case, collaborative equilibria analogous to E1 and E2 can be supported under the same conditions as in the engineer's full control case.

4. Experimental Design

The sequence of events and the payoffs in each round of the experiment follow the stage game presented in Figure 1. We set the probability of innovation at $\pi = 0.75$, which allows us to get a high frequency of the interesting outcome where the innovation happens.¹⁵ We design the experiment using the strategy method, by which participants make conditional decisions for each possible scenario that may arise, in order to elicit subjects' complete strategies. 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 it (ambiguous node).¹⁶ After suppliers and buyers have made their decisions, the computer randomly determines whether the innovation occurs 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 the buyer's decision (as well as the procurement manager's and the engineer's individual recommendations in the corresponding treatments), and their payoffs.

We conduct three treatments corresponding to the Buyer-as-Two-Employees set-up, where subjects are randomly matched in groups of three players. One subject is assigned the role of supplier, and two subjects are assigned to be employees working for the buyer—one in the role of procurement manager and the other in the role of engineer. At the firm level, the "supplier" and "buyer" have a long term relationship. Subjects in the roles of supplier and engineer play together as long as the relationship between the two firms lasts. Subjects in the role of procurement manager are rotated

¹⁵ 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).

¹⁶ We use the strategy method so we can fully understand subjects' strategies, 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.

among buyer-supplier pairings in each round. We capture the long-term relationship between the two firms as a repeated game with a random stopping time. That is, after each round, the relationship continues to the next round with probability δ . On the other hand, with probability $1-\delta$, the relationship ends. In the experiment, we use a continuation probability of $\delta = 0.75$ to allow for the collaborative equilibrium in Proposition 1 to arise, both in the engineer's control and the joint control cases. The length of a relationship is equal to the number of rounds where the same buyer and supplier engage in stage games (the variable round is re-set at the beginning of each new relationship). A continuation probability of $\delta = 0.75$, implies an average relationship length of four rounds. This value was set so that the relationship length is relatively short (and, therefore, collaboration is rather hard to sustain as in our motivating examples), while it still guarantees that the collaborative equilibrium in Proposition 1 can arise. ¹⁷ 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 (for example, period = 10 means that a subject has played a stage game 10 times). We implement this as follows: After each round, we us a random number generator to determine if the relationship between the firms continues or ends. If the relationship continues, the supplier and the engineer 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. These matching rules are common knowledge to all players.

Each round consists of two stages. In the first stage, the supplier decides whether, if the innovation has occurred, they want to share it with the buyer. 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 (ambiguous node). Since the engineer has been matched with the same supplier starting from the first round of the relationship between the firms, the engineer 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 full history of play in the relationship. 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. Subjects playing as suppliers get the payoff of the supplier

¹⁷ This relationship length is consistent with the automotive industry for example, where firms usually make supplier selection decisions on an annual basis and relationships of four years are not uncommon (McMillan 1990).

firm and subjects playing as procurement managers and engineers each get the payoff of the buying

The three treatments in the Buyer-as-Two-Employees setup capture the three different allocations of decision rights between the engineer and the procurement manager that we analyze in the theory section. In the procurement manager decision treatment (denoted PM), the procurement manager's recommendation is always implemented. The opposite happens in the engineer decision treatment (denoted ENG), where the engineer's recommendation is always implemented. We then consider a joint control treatment where decision rights are equally split (denoted 50-50), where the computer randomly picks one recommendation to implement (both with equal probability) if they disagree. All subjects are informed that the allocation of decision rights is the one corresponding to the treatment they are in.¹⁸

Lastly, we conduct two benchmark treatments corresponding to the Buyer-as-Single-Employee setup, where each firm is a single decision maker. Subjects are randomly and anonymously matched into pairs and are assigned a role (supplier or buyer), which they keep throughout the session. In the first treatment, the buyer and the supplier have a short-term relationship which lasts only one round (we denote it SR treatment). We induce this by randomly re-matching buyers and suppliers after each round. In the second treatment, the buyer and the supplier have a long-term relationship (we denote it LR treatment). We implement this with the same random end procedure used in the Buyer-as-Two-Employees setup, also with a continuation probability of 0.75.

Overall, the experiment consists of three main treatments (PM, 50 - 50, and ENG), and two additional benchmark treatments, (SR and LR). We adopt a between-subjects design, i.e. each subject is exposed to only one treatment. To ensure subjects' understanding of the game, three examples are presented in the instructions, and the table with payoffs (Figure 4 in the Online 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 is made explicit that the probability that the relationship will continue for another round remains exactly the same after every round. Finally, participants only learn the outcomes and payoffs of their own relationships to avoid reputation effects.

4.1. Hypotheses

We derive the following experimental hypotheses from our theoretical analysis. The first hypothesis is derived from the equilibrium outcomes of the Buyer-as-Two-Employees setup. Consider first the

¹⁸ Note that in the PM (ENG) treatment, the engineer's (procurement manager's) recommendation is not monetarily incentivized, as it has no direct impact on the payoffs. Nevertheless, since both recommendations are revealed at the end of each round, the recommendation allows the employee with no decision rights to express his preference to the other employee and to the supplier. By randomly re-assigning a new role to engineers and procurement managers after each relationship, most subjects get to play the role with decision authority at some point during the session.

PM treatment, where the procurement manager's recommendation is always implemented. Since the procurement manager is part of the buyer-supplier pairing for only one period, the only equilibrium outcome in the PM treatment is that the supplier does not share and the procurement manager's recommendation is to compete in every round. Consider now the ENG treatment, where the engineer's recommendation is always implemented. Proposition 1 prescribes that an equilibrium where the supplier always shares and the engineer always recommends not to compete when the supplier shared can be sustained with the values of parameters used in the experiment. Finally, in the 50-50 treatment, Proposition 1 prescribes an equilibrium where the supplier always shares, the engineer recommends not to compete and the procurement manager recommends to compete. Since the final decision is randomly determined, the buyer follows the procurement manager's recommendation and the engineer's recommendation each 50% of the times. Thus, the frequency of rounds with collaborative outcomes in the 50-50 treatment should be higher than in the PM treatment but lower than in the ENG treatment. Overall, the theoretical results stipulate that we should see a gradient of increased collaboration from (1) the PM to (2) the 50-50 to (3) the ENG treatment.

HYPOTHESIS 1. [Trends across Buyer-as-Two-Employees treatments] There is a trend of increasing collaboration from PM to 50-50 to ENG:

- 1.a the frequency of sharing is higher in the 50-50 and ENG treatments relative to PM,
- 1.b the frequency of competing (if the supplier shared) decreases from PM to 50-50 to ENG, and
- 1.c the frequency of collaborative outcomes increases from PM to 50-50 to ENG.

In order to establish a reference for the difference between the two extreme cases of allocation of decision rights (PM and ENG), we compare these two treatments to the theoretically-equivalent benchmark cases where the buyer is a single decision maker (SR and LR respectively). Since the relationship between the procurement manager and the supplier resembles a one-shot game, 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.

HYPOTHESIS 2. [Comparison with Buyer-as-Single-Employee Benchmarks] The treatments where the procurement manager or the engineer have a full control obtain the same outcomes as the SR and LR benchmarks respectively:

- 2.a the PM treatment obtains the same outcomes as the SR treatment, and
- 2.b the ENG treatment obtains the same outcomes as the LR treatment.

5. Experimental Results

The experiments were conducted in z-Tree (Fischbacher 2007) between March and September of 2014 at the behavioral laboratory of a large public university in the Midwest of the United States. A total of 372 undergraduates participated in four sessions of each of the Buyer-as-Single-Employee treatments (84 in PM, 93 in 50-50, and 81 in ENG) and in six sessions of each of the Buyer-as-Two-Employees treatments (52 in SR, and 62 in LR). 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 while all the other treatments ended after 50 minutes (including the time to read the instructions). The average number of rounds per relationship was 3.9, with a minimum of 1 and a maximum of 11. Average payoffs were \$11, consisting of a \$5 show up fee plus the payoffs of two randomly selected rounds at a conversion rate of \$0.10 per point earned. The selected rounds are conversion rate of \$0.10 per point earned.

In the next sub-section we test Hypothesis 1 by comparing the three treatments in the Buyer-as-Two-Employees setup. We then test Hypothesis 2 by comparing the PM and ENG treatments with their corresponding Buyer-as-Single-Employee benchmarks.

5.1. Employees' Decisions in the Buyer-as-Two-Employees Treatments

Hypothesis 1 predicts that the frequency of the collaborative outcome—the supplier shares and the buyer does not compete—shows an increasing trend from (1) the PM treatment, to (2) the 50-50 treatment, to (3) the ENG treatment. This results from the predictions for the supplier's and the buyer's decisions in each treatment. The supplier is expected to be least collaborative in the (1) PM treatment and most collaborative in the (2) 50-50 and (3) ENG treatments. Likewise, we expect that buyers are most likely to choose "compete" in the (1) PM treatment, least likely to choose "compete" in (3) ENG treatment, while the result in (2) the 50-50 treatment falls in between (1) and (3).

Table 1 shows the frequency of the supplier's decision to "share", the buyer's decision to "compete if the supplier shared" and "compete if the supplier did not share", and the frequency of the collaborative outcome (share/not compete), in each treatment.²¹ Note that we use the denomination "if

¹⁹ We use all observations up to period 30, which is the latest period that was reached in every session (except one session, which lasted 29 periods). Our main results do not change significantly if we use the observations from all periods.

²⁰ 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.

 $^{^{21}}$ In the 50-50 treatment, "compete if shared" and "compete if not shared" are the decisions of the player whose recommendation was actually implemented.

the supplier did not share" for the buyer's decision in the ambiguous node, where the buyer does not know whether the innovation occurred and the supplier did not share it or whether the innovation did not occur. The table presents data that is aggregated at the subject level for each subject in the role of supplier. That is, one observation is equal to the fraction of times a decision or an outcome happened in all the rounds played by a supplier in a session. Recall that, since we used the strategy method, we were able to elicit the buyer's decision whether to compete if the supplier shared and if the supplier did not share, in each round and for each supplier.

Providing support for Hypothesis 1, we find that there is a trend of increased collaboration from (1) PM, to (2) 50-50, to (3) ENG (a non-parametric test for trends is presented in the last row of Table 1).²² As the trend is only marginally significant for the supplier's decision to share, the significant increasing trend in the frequency of collaborative outcomes is mostly driven by buyers' decreasing trend in the decision to compete if the supplier shared. Pairwise comparisons of the results in Table 1 show that the supplier's decision to share is not significantly different across treatments (Wilcoxon rank-sum test p-values: PM vs. 50-50 = 0.301, PM vs. ENG = 0.121, 50-50vs. ENG = 0.391). For the buyer's decision to compete if the supplier shared, pairwise comparisons show that it is significantly lower in the ENG treatment than in the PM and 50-50 treatments (ENG vs. PM: p = 0.017, ENG vs. 50-50: p < 0.001). When we compare the PM and 50-50 treatments, we find that "compete if the supplier shared" is chosen significantly more often than expected in the 50-50 treatment (recall that it was expected to be significantly lower than in the PM treatment and it is actually marginally higher, p = 0.078). This deviation from our predictions is analyzed further in the following sections. The buyer's decision to compete if the supplier did not share shows a similar pattern but the pairwise differences are not statistically significant (p-values: PM vs. 50-50 = 0.303, PM vs. ENG = 0.488, 50-50 vs. ENG = 0.138). Finally, 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. Table 1 shows that the frequency with which collaborative outcomes occur is significantly higher in the ENG treatment than in the other two treatments (ENG vs. PM: p = 0.004, ENG vs. 50-50: p = 0.007), however it is not statistically different across the PM, and 50 - 50 treatments (p = 0.656).

Table 2 presents panel probit regressions with subject random effects for suppliers in the PM, 50-50, and ENG treatments. In all cases we control for round within a relationship and period of play in the session. The first column shows that the supplier's decision to share is not significantly higher in the 50-50 treatment than in the PM baseline and that in the ENG treatment it is only marginally higher. The second column shows that the frequency of the collaborative outcome is not higher in

²² The non-parametric test for trends across ordered groups is an extension of the Wilcoxon rank-sum test (Cuzick 1985).

Trend test (p-value)

0.097

	Table 1	cheral results - Buyer-	us- Two-Employees Treatmen	
	Supplier's	Buyer's	Buyer's	Collaborative
	Decision	Decision	Decision	Outcome
	(Share)	(Compete if Shared)	(Compete if Not Shared)	(Share/Not Compete)
Treatment	(%)	(%)	(%)	(%)
PM	29.2	64.3	73.2	9.5
50 - 50	33.5	71.2	75.1	10.8
ENG	40.2	50.8	68.7	24.6

Table 1 General Results - Buyer-as-Two-Employees Treatments

The number of observation is equal to the number of subjects who played in the role of supplier in each treatment: 28 in PM, 31 in 50-50, and 27 in ENG.

0.553

0.003

0.033

the 50-50 treatment than in PM, but it is significantly higher in the ENG treatment. These findings confirm the results in columns (1) and (4) of Table 1. First, it appears that the supplier's decision to share the innovation does not differ significantly based on the allocation of decision rights within the buyer firm. Second, the frequency of the collaborative outcome 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 5.1.1. by analyzing the 50-50 treatment in more detail.

Table 2 Share and Collaborative Outcomes - Buyer-as-Two-Employees Treatments

	Supplier's	Collaborative
	Decision	Outcomes
Coefficients	(Share)	(Share/Not Compete)
50 - 50	0.154	0.072
	(0.181)	(0.192)
ENG	0.340*	0.660***
	(0.188)	(0.194)
Constant	-0.290**	-1.622***
	(0.142)	(0.157)
Controls	Yes	Yes
Observations	2576	2576
Nr. of Subjects	86	86
$\overline{\text{Test } 50\text{-}50} = \overline{\text{ENG}}$	0.308	0.002

Probit regression with subject random effects. Standard errors reported in parentheses. Significance is denoted: * p < 0.10 ** p < 0.05 *** p < 0.01. Data for suppliers in the PM, 50-50, and ENG treatments. Controls include period and round.

A surprising result in Table 1 is that the frequency of "compete if shared" in the 50-50 treatment is marginally higher than in the PM treatment (while it was expected to be lower). To study why, we analyze the recommendations of engineers and procurement managers separately. We aggregate the data at the subject level for each subject who played in the role of employee working for the buyer by computing the fraction of times the "compete" decision was chosen in all the rounds where the subject played in the role of procurement manager (engineer). Figure 2 shows that the percentage of

times an engineer chooses "compete if the supplier shared" is higher in the 50-50 treatment than in the PM and ENG treatments (50-50 vs. PM: 67% vs. 52%, p=0.013; 50-50 vs. ENG: 67% vs. 52%, p=0.033). Similarly, procurement managers choose to "compete if the supplier shared" more often in the 50-50 treatment than in the PM and ENG treatments (76% vs. 65%, p=0.055 and 76% vs. 65%, p=0.071 for 50-50 vs. PM and 50-50 vs. ENG respectively). Table 3 confirms the previous result with panel probit regressions of the decision to "compete if share" on treatment dummies. For these regressions we use the decisions of all the subjects who played in the role of employee working for the buyer in the 50-50 and ENG treatments (columns 1 and 2) and in the 50-50 and PM treatments (3 and 4). We control by period, round, and subject random effects. Columns 1 and 3 use only the decisions in the rounds where subjects played in the role of engineers and columns 2 and 4 use the decisions in the rounds where subjects played in the role of procurement managers. Taken together, Figure 2 and Table 3 confirm that the higher frequency in the decision to compete in the 50-50 treatment is driven by both engineers and procurement managers being more likely to compete in the 50-50 treatment than in the PM and ENG treatments.

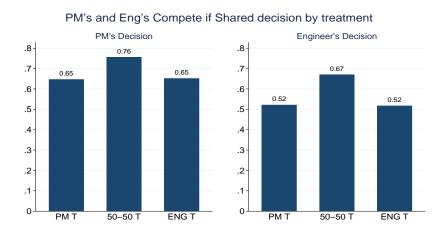


Figure 2 Compete if Shared by Treatment

5.1.1. The 50-50 Treatment In this section we propose two potential explanations for why the frequency of the decision to compete when the supplier shared is higher in the 50-50 treatment for both the procurement manager and the engineer, and we test them with two additional experimental treatments.²³ The first mechanism that can explain why both employees are less collaborative in the joint control treatment is diffusion of responsibility, by which an individual is less likely to help others in presence of another individual who can help. In our setting, this could translate into an

.

²³ The new treatments were conducted between February and April of 2018 in the same laboratory and with the same subject pool and recruitment procedure as the original treatments.

50-50 vs. ENG treatments Coefficients Engineer 's Decision PM's Decision Engineer's Decision PM's Decision ENG treatment **888.0 -0.571^* (0.383)(0.321)PM treatment -0.930*** -0.620** (0.334)(0.306)1.413*** Constant 1.584*** 1.557*** 1.561*** (0.284)(0.239)(0.247)(0.256)Controls Yes Yes Yes Yes

Table 3 Compete if Shared Decision - Buyer-as-Two-Employees Treatments

Probit regression with subject random effects. Standard errors reported in parentheses. Significance is denoted: * p < 0.10 ** p < 0.05 *** p < 0.01. Data of employees working for the buyer. Columns (1) and (2) pool data from the 50 - 50 and ENG treatments exclusively. Columns (3) and (4) pool data from the 50 - 50 and PM treatments exclusively. Controls include period and round.

1736

115

1770

118

1770

118

1736

115

Observations

Nr. of Subjects

employee being less willing to "help" the supplier (i.e. choose not to compete) when another employee is also present and able to do so. The second mechanism that we study is *social interaction effects*, by which the engineer and the procurement manager may be biased to make their recommendations match when they are both involved in the decision making process.

Diffusion of Responsibility. A potential mechanism that can be at play in the joint control treatment is a social inhibition of helping. In an early review of the literature on diffusion of responsibility, Latané and Nida (1981) define this phenomenon as follows: "The knowledge that others are present and available to respond, even if the individual cannot see or be seen by them, allows the shifting of some responsibility for helping". In our setting this could imply that in the 50-50 treatment (where both employees are involved in the decision making process) the engineer feels less responsible for taking the collaborative action. Concrete examples of diffusion of responsibility have been documented as the bystander effect, where individuals are inhibited from intervening in an emergency when others are present. Experimental evidence of this has been found in student bystanders to another student having a seizure similar to epilepsy (Darley and Latané 1968), in students reacting to a smoke filling room (Latane and Darley 1968), bystanders to a stranded motorist (Hurley and Allen 1974), in providing assistance by email (Barron and Yechiam 2002), and in the form of moral "wiggle room" (Dana et al. 2007).

Dana et al. (2007) find experimental evidence of diffusion of responsibility in a binary dictator game, where the dictator chooses between an equal (fair) or an unequal (unfair) allocation between himself and a recipient. They compare this baseline game with a modified dictator game where two dictators are matched with one recipient, and each dictator can unilaterally implement the fair outcome. By eliminating a dictator's sole responsibility for the unfair outcome, this treatment allows

them to test whether diffusion of responsibility favors higher choices of the unfair outcome. Indeed, they find that the frequency of the unfair outcome is significantly higher in the multiple dictator game than in the baseline (74% vs. 35% respectively). We design a new Buyer-as-Two-Employees treatment (where an engineer and a procurement manager make recommendations for what the buyer should do), that builds on the manipulation in Dana et al. (2007). Specifically, if at least one employee recommends not to compete, then that decision is implemented. Otherwise, the decision to compete is implemented. Since both recommendations need to be "compete" for this decision to be implemented, we denote this treatment two-to-compete (TWO) to COMP). In this setting, the procurement manager has the same incentive to compete as in the ENG treatment. Given that choosing to compete is a dominant strategy for the procurement manager, the engineer's incentives are also the same as in the ENG treatment (note that the engineer can unilaterally ensure the same outcome as in the ENG treatment). Consistently with diffusion of responsibility, engineers could potentially "justify" that the decision to compete is only implemented if both employees allow it, reducing their own personal responsibility for the decision. Therefore, if diffusion of responsibility is at play, engineers should be more prone to compete in TWO to COMP than in the ENG treatment. In addition to providing a setting where diffusion of responsibility can be tested, the TWO to COMP treatment allows us to capture a realistic setting where the firm has a high-level policy (to compete) but the employees still have discretion to unilaterally enforce a collaborative action towards a supplier.

A total of 48 subjects participated in five sessions of the TWO_to_COMP treatment. Since our focus is on the comparison between the ENG and TWO_to_COMP treatments, we pool the data of these two treatments and exclude all other treatments for analysis. The seeds for the random relationship length and the total number of periods used was kept the same as in the ENG treatment for comparison. To compare the two treatments, we conduct a panel probit regression of the frequency with which engineers choose to compete if the supplier shared and if the supplier did not share on a dummy variable which takes value one in the ENG treatment and zero in the TWO_to_COMP treatment. We find that the coefficients are not significant ($\beta = -0.096$, p = 0.394; and $\beta = 0.025$, p = 0.359 if the supplier shared and did not share respectively) suggesting that, in the setting where diffusion of responsibility is favored and the responsibility for a "compete" outcome cannot be attributed to either employee exclusively, engineers do not choose to compete significantly more often (the full regression is presented in Table 10 in the Online Appendix).

Social Interaction Effects. According to a definition given by Cooper et al. (2011), social interaction effects are present when "an individual's utility from an action is enhanced by others taking the same action". In our setting this would imply that the engineer and the procurement manager

get an additional psychological utility if their recommendations match. Cooper et al. (2011) find evidence of this effect in lab experiments where subjects make decisions under risk and uncertainty and learn other group members' choices before making their own decision. While in their setting there are no direct strategic interactions between subjects (i.e. each subject's payoff depends only on their own decision), we conjecture similar effects could be present when subjects make decisions jointly. The experimental economics literature has explored whether groups behave differently than individuals in various settings. Most of the results in this literature indicate that groups behave more in line with standard game-theoretical predictions under the assumption of rationality and selfishness than individuals (Kocher and Sutter 2007). Subjects acting on their own, are more likely to be influenced by biases, cognitive limitations, and social preferences (Charness and Sutter 2012).²⁴ For example, Bornstein and Yaniv (1998) find that groups offer less and are also willing to accept less than individuals in the ultimatum game. Cox (2002) finds that, groups in the role of responders in trust games send back smaller amounts than individuals. Kugler et al. (2007) find that groups are just as trustworthy as individuals, but in the role of senders they are less trusting than individuals. Similarly, Ambrus et al. (2015) find that groups on average reciprocate less than individuals in the gift-exchange setting.²⁵ The fact that both the procurement manager and the engineer are less trustworthy in the 50-50 treatment (which most closely resembles group-decision making) is in line with this literature. The main differential feature of our group decision setting is the misalignment of monetary and strategic incentives between the two employees, which should (in theory) hinder employees' propensity to reach an agreement.

To explore this, we conduct a joint control treatment with communication (we denote it 50-50+CHAT treatment). In this treatment, the decision mechanism is as in the 50-50 treatment—the engineer and the procurement manager make their recommendations, if they agree that decision is implemented and, if they disagree, one of the two recommendations is selected at random—but, in addition, the engineer and the procurement manager can communicate with each other before making their recommendations. This treatment aims to provide the employees with a better ability to reach an agreement, should they wish to do so. To implement this as part of the computerized experiment, we follow the principles in Cooper and Kagel (2005) who introduced chat analysis to obtain insights on subject's learning in a signaling game. In our design, the two employees enter

²⁴ This susceptibility to outside influence has been studied in social psychology, and it has been attributed to three goals fundamental to human behavior, people's desire to form accurate perceptions of reality and react accordingly, a motivation to develop and preserve meaningful relationships, and a desire maintain a favorable self-concept (see Cialdini and Goldstein (2004) for an extensive literature review).

²⁵ The related literature on group identity has shown that individuals tend to be more altruistic towards an in-group match (Chen and Li 2009) and that individuals behave more aggressively towards an out-of-group individual when they have their group members in the audience (Charness et al. 2007).

a private chatroom for 90 seconds where they can exchange text messages before their individual decision screens are displayed. Subjects receive no direct instructions about what messages to send to each other (they are only instructed to be civil to each other and to not identify themselves) and each employee is able to leave the chatroom at any time. By incorporating free-form communication, this treatment captures a more realistic setting where employees may have meetings or informal conversations before making their recommendations to the firm. Note that, as in the 50-50 treatment, the two employees have divergent strategic interests and therefore there is little reason for them to try to reach an agreement. If, despite this, subjects choose to engage in communication their messages may allow us to better understand the underlying behavioral mechanisms driving their behavior.

In total 75 subjects participated in five sessions of the 50-50+CHAT treatment. To make the data of the two treatments comparable, we matched the seeds in each session of the 50-50+CHAT treatment with a session of the 50-50 treatment (so that the relationship lengths of the two treatments are equivalent). In addition, since the 50-50+CHAT treatment had fewer periods per session (13 on average), we only use the corresponding number of periods in the 50-50 treatment when we compare the two treatments. We pool the data from the 50-50 and the 50-50+CHAT treatments and we conduct a panel probit regression of the engineer's and the procurement manager's frequency of choosing "compete" on a dummy variable that takes value one in the 50-50 treatment and zero in the 50-50+CHAT treatment. For engineers and procurement managers we find that the frequency of choosing to compete is not significantly different across the two treatments (the full regression is presented in Table 11 in the Online Appendix: $\beta = 0.471$, p = 0.377 for the engineer; $\beta = 0.201$, p = 0.312 for the procurement manager). Given that adding the chatroom stage to the 50-50 treatment does not significantly affect the overall frequency with which employees recommend to compete, we can use the conversations in the 50-50+CHAT treatment to better understand the employees' behavior.

We follow the procedures in Cooper and Kagel (2005) to analyze the conversations between the two employees. First, we developed a coding scheme based on sample chats with the goal of capturing the main relevant statements related to reaching agreement. Specifically, the codes capture for each chat:

1) whether subjects showed intention to come to an agreement, 2) whether the subjects agreed to take the same action, agreed to take different actions, or did not reach an agreement, 3) in case they reached an agreement to take the same action, whether they agreed to compete or not to compete, 4) whether the subjects expressed motivations or concerns about the supplier's past sharing behavior, about building supplier's trust, or about maximizing their own period payoff (the full coding scheme and instructions for coding are presented in the Online Appendix). The next step was to recruit two graduate students as research assistants, who were trained in the experimental background and

coding scheme. For the total 330 chatrooms generated in the five sessions, the coders assigned 1711 codes and agreed in 95.7% of the codes.²⁶ The last step was to analyze the frequency with which each code arose and how the codes correlate with the employees' actual recommendations.

Table 4 Chats Analysis - Agreement

		Frequency	If shared		If did not share	
		(# of chats)	Comp Not Comp		Comp	Not Comp
No Agreement		26				
Agree different action		12				
Agree same action	Eng proposed	121	84 (69%)	37 (31%)	88 (73%)	33~(27%)
	PM proposed	102	75 (74%)	26~(26%)	79 (78%)	22(22%)
	Both proposed same	60	55 (92%)	5 (8%)	53 (88%)	7 (12%)

Note: Six chats were excluded from the classification No agreement/Agree different action/Agree same action, due to ambiguous coding. Within the category Agree same action-PM proposed, one chat was excluded from the agreement to Compete/Not compete classification, due to ambiguous coding.

Table 5 Frequency of Deviation from Agreement

	I	Engineer	Procurement Manager		
	If shared	If did not share	If shared	If did not share	
Agreement to compete	4.2%	2.7%	1.9%	2.3%	
Agreement not to compete	7.4%	11.3%	16.2%	11.3%	
Difference* (p-value)	0.291	0.004	< 0.001	0.002	

(*) Wilocoxon rank-sum test. Number of observations: Total number of chats coded as compete if shared 216, not compete if shared 68, compete if not shared 222, not compete if not shared 62.

Table 6 Chats Analysis - Drivers of Behavior

	[Supplier_past_share]		[Desire_B	uild_Trust]	[Earn_Points]	
	Frequency	% of	Frequency	% of	Frequency	% of
	(# chats)	total chats	(# chats)	total chats	(# chats)	total chats
Engineer	45	14%	37	11%	23	7%
PM	24	7%	17	5%	29	9%
Difference Test 1 (p-value)	0.007		0.004		0.405	
Difference Test 2 (p-value)	0.001		0.003		0.330	

<u>Codes</u>: Supplier_past_share = Employee discussed the supplier's past sharing behavior. Desire_Build_Trust = Employee expresses desire to build trust/build a good relationship with the supplier. [Earn_Points] = Employee expresses desire to get a high payoff/earn many points in present round. Test 1 is a Wilcoxon rank-sum test, Test 2 is a Wilcoxon signed-rank test.

The chats analysis shows several important insights: The first result is that engineers and procurement managers choose to engage in conversation, even when—in theory—they have no incentive

²⁶ A third coder was recruited and trained for tie-breaking in the reminder 4.3% of the codes. After tie-breaking, only 0.8% of the codes remained ambiguous and were excluded from the analysis.

to do so. From the total 330 chatrooms only three were dropped because the subjects left the chatroom immediately. The codes indicate that in all of the reminder 327 chatrooms, the procurement manager expressed desire to discuss strategies. The engineers did so in 323 of the 327 chatrooms. The second result is that engineers and procurement managers usually seek to reach an agreement. Table 4 shows that an agreement is not reached in only 26 chats, while in 295 chats an agreement is reached (in 12 chats the employees agree to take different actions and in 283 chats they agree to take the same action). Table 5 shows that, conditional on having reached an agreement, deviations from the agreement are rare (particularly when the agreement was to compete in which case deviations are less than 5%) indicating a positive correlation between the chat agreement and the employee's subsequent recommendation. As a result, Figure 3 shows that the frequency with which the recommendations of the two employees match in the 50-50+CHAT treatment is significantly higher than in the 50-50 treatment (p-value for difference across treatments < 0.001 in both charts). The third result is that reaching an agreement to compete is easier than reaching an agreement not to compete. Table 4 shows that in the majority of the chats where the employees agreed to take the same action, the agreement was to compete (by approximately 70% vs. 30% against not compete). Note that the frequency of agreement to compete is particularly higher (p-value < 0.01) in the chats where both employees spontaneously proposed the same action—and thus where agreement arose most easily. In addition to being easier to reach, an agreement to compete is significantly more likely to be honored, particularly by procurement managers (see Table 5 for frequencies of deviation from agreement).

Finally, we examined whether subjects expressed any motivations or concerns to support a particular action. While in most chats either both employees propose the same action or an agreement is reached fairly soon (one employee proposes and the other immediately agrees), there is a minority of chats where strategies and concerns are discussed in depth. Table 6 shows that the main concerns leading to "not compete" (i.e. suppliers' past sharing and future relationship building) are more prevalent among engineers, whereas concerns leading to "compete" (i.e. maximizing the current period payoff) are equally present for both employees. The fact that the concerns leading to competing are shared by both employees also contributes to making an agreement to compete easier to reach.²⁷

Taken together, our experimental results in the 50-50 and 50-50+CHAT treatments suggest that social interaction effects arise in the form of a bias to agreement. To account for this, we modify our

²⁷ Test 1 in Table 6 presents a Wilcoxon rank-sum test where the decision of each procurement manager (engineer) in each chat is one observation. Test 2 presents a Wilcoxon sign-rank test, where the decisions of the procurement manager and the engineer are compared in pairs within a chat. The results of the two tests are consistent.

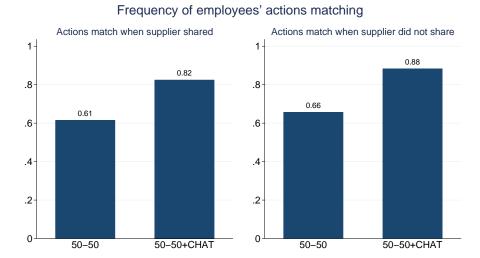


Figure 3 Frequency with which the employees' actions match

theoretical model and assume that, beyond their monetary payoffs, the engineer and the procurement manager get an additional utility if their recommendations match. We capture this utility with a behavioral parameter, $\beta_i \in [0, \bar{\beta}]$ for employee $i, i \in \{E = \text{engineer}, P = \text{procurement manager}\}$, which is common knowledge to all players. Specifically, in each round where their recommendations match, the engineer and the procurement manager earn their monetary payoff plus an extra utility, β_i , each. The experimental results show that this bias favors the outcome where the procurement manager and the engineer agree to compete, over the theoretical prediction that the engineer would be more inclined to "not compete" while the procurement manager has incentive to compete. Therefore, we are interested in how the bias towards agreement affects the conditions under which the following equilibria can arise: 1) agreement to compete: a non-collaborative equilibrium where the procurement manager and the engineer compete in every round and the supplier never shares the innovation, 2) agreement not to compete: a collaborative equilibrium where the procurement manager and the engineer do not to compete if the supplier shared the innovation and the supplier shares the innovation whenever there is one, and 3) disagreement despite bias: a collaborative equilibrium where the engineer does not compete if the supplier shared and the supplier shares whenever there is an innovation, and the procurement manager plays the non-collaborative strategy of the single-shot game. Note that this last case corresponds to the equilibrium characterized in Proposition 1, which we have already shown can be supported when firms' do not have bias towards agreement.

In Appendix 7.2 we show the sets of parameters β_i under which each of the equilibria can be supported for the values of the parameters used in the experiment. The results confirm that the bias to agreement significantly favors the equilibrium with an agreement to compete. In fact, while

this equilibrium can be supported even without employees' bias to agreement, the bias makes both employees' incentive compatibility constraints easier to satisfy. An equilibrium with agreement not to compete if the supplier shared the innovation can be supported, however, it requires the procurement manager's bias to agreement, β_P , to be quite high. Specifically, in order for the result to hold, incentive compatibility for the procurement manager requires $\beta_P \geq 56$, which is equivalent to giving up half the total monetary surplus from a shared innovation. Finally, an equilibrium where employees disagree despite the bias to agreement can also be supported only in a small subset of β_i 's. Specifically, the engineer's incentive compatibility constraint requires $\beta_E \in [0, 20.40]$, and the procurement manager's $\beta_P \leq 28$. The experimental results of the 50-50+CHAT treatment are consistent with these findings. An agreement to compete (recall that the incentive compatibility constraints of this case become easier to satisfy with the bias to agreement) is the most prominent experimental outcome as it occurs in 214 of the 321 coded chats. An agreement not to compete if the supplier shared and a disagreement despite the bias to agreement (which can also be supported in equilibrium although the incentive compatibility constraints become harder to satisfy with the bias to agreement) only occur in 68 and 12 of the 321 coded chats respectively.

5.2. Comparison with Buyer-as-Single-Employee Benchmarks

Hypothesis 2 predicts that the PM and ENG treatments obtain the same outcomes as the Buyeras-Single-Employee benchmarks (SR and LR respectively). In this section we compare the PM and SR treatments, and the ENG and LR treatments. Table 7 pools data of the PM and SR treatments (columns 1 to 3) and the ENG and LR treatments (columns 4 to 6) and presents panel probit regressions with subject random effects of the decision variable on treatment indicator variables. Columns 1 and 4 include the decision to share by subjects in the role of suppliers. Columns 2 and 3 include the decisions to compete by buyers in the SR treatment, and by subjects in the role of employee working for the buyer in relationships where they were procurement managers in the PM treatment. Similarly, columns 5 and 6 include buyers' decisions in the LR treatment and the decision of subjects in the role of employee working for the buyer in the relationships where they played as engineers in the ENG treatment. We observe that, as predicted by the theory, the supplier's and the buyer's decisions are not significantly different in the ENG treatment and the LR benchmark. In the PM treatment the supplier shares significantly more often than in the SR benchmark. The buyer competes directionally less in the PM treatment relative to SR, but the difference is only marginally significant for the decision to "compete if not shared" (which captures the buyer's decision in the ambiguous node). As a result, the frequency of collaborative outcomes for a supplier is significantly higher in the PM treatment than in the SR treatment. Table 12 in the Online Appendix shows that a treatment dummy has a significantly positive effect on the frequency of collaborative outcomes $(\beta = 0.400, p = 0.036).$

Table 7

			6 1 3		
	PM vs. SR	,		ENG vs. L	R
Supplier's	Buyer's	Buyer's	Supplier's	Buyer's	Buy
Decision	Decision	Decision	Decision	Decision	Decis
(63	/ C	/ C	(01	/ 🛪 .	/~

Comparison with Buyer-as-Sinlge-Employee Benchmarks

		PM vs. S.	R	ENG vs. LR			
	Supplier's	Buyer's	Buyer's	Supplier's	Buyer's	Buyer's	
	Decision	Decision	Decision	Decision	Decision	Decision	
	(Share)	(Compete	(Compete	(Share)	(Compete	(Compete	
		if shared)	if not shared)		if shared)	if not shared)	
PM	0.580***	-0.355	-0.713*				
	(0.206)	(0.407)	(0.415)				
ENG				0.068	-0.404	0.684	
				(0.218)	(0.438)	(0.457)	
Constant	-0.710***	1.221***	1.852***	-0.247	0.776**	1.031***	
	(0.159)	(0.342)	(0.355)	(0.163)	(0.357)	(0.374)	
Controls	Yes	Yes	Yes	Yes	Yes	Yes	
Observations	1620	1620	1620	1736	1736	1736	
Nr. of Subjects	54	82	82	58	84	84	

Probit regression with subject random effects. Standard errors reported in parentheses. Significance is denoted: * p < 0.10 ** p < 0.05 *** p < 0.01. Columns 1-3 pool data of the PM and SR treatments, columns 4-6 pool data of the ENG and LR treatments. Controls include period and round.

Overall, the main difference with the Buyer-as-Single-Employee benchmarks is that suppliers share significantly more often in the PM treatment than in the SR benchmark. One potential driver of this behavior is social preferences: in the PM treatment the supplier interacts with the same engineer repeated times and, as a result, fairness concerns towards the buyer may become more salient (Fehr and Schmidt 1999, Bolton and Ockenfels 2000). Given that the frequency of collaborative outcomes for suppliers is significantly higher in PM than in SR, an interesting question is whether the buyer's marginally more collaborative behavior in the PM treatment relative to SR justifies the significant increase in the frequency of supplier's decision to share. The third and fourth columns in Table 13 in the Online Appendix show the supplier's expected profit from sharing and from not sharing, and the difference between the two, given how the buyers respond to these two actions in the experiment. We observe that the difference is positive in the PM treatment (17.08 vs. 15.99, diff = 1.09) and negative in the SR treatment (15.63 from sharing vs. 15.84 from not sharing, diff = -0.22). This implies that, in expectation, sharing is profitable in the PM treatment and not in the SR treatment. In the experiment, suppliers seem to acknowledge this difference: they share 18.2% of the times in the SR treatment and 29.2% of the times in the PM treatment. Table 14 in the Online Appendix confirms that sharing is profitable 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. These results suggest that, while behavioral factors (such as social preferences) may lead suppliers to share significantly more often in the PM treatment relative to SR, suppliers also have strategic incentives to do so as the buyer becomes marginally more collaborative in PM relative to SR.

5.3. Profits

The results described above indicate that, as predicted by the theory, there is a significant trend of increased collaboration across the Buyer-as-Two-Employees treatments when ordered (PM)-(50 – 50)-(ENG), and the PM and ENG treatments are fairly similar to their benchmarks. However, we also find that suppliers' and buyers' actions present some deviations from the predictions. Interestingly, these deviations go in opposite directions for the supplier and for the buyer. We observe that the frequency with which suppliers share actually increases 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. We find that the frequency with which buyers choose "compete if the supplier shared" in the 50-50treatment 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. These results have a direct impact on profits. Table 8 presents subject-level average profits for suppliers, buyers, and both players combined. The results show that suppliers only earn significantly higher profits than in the SR benchmark in the ENG and LR treatments (SR vs. ENG, p < 0.001; SR vs. LR, p = 0.031; K-Wallis test, p < 0.001).²⁸ On the other hand, buyers benefit from all Buyer-as-Two-Employees treatments. Buyers' profits are higher than in the SR benchmark in all the other treatments (directionally in SR vs. PM, p = 0.104; and statistically significant in: SR vs. 50-50, p = 0.014; SR vs. ENG, p = 0.014; SR vs. LR p = 0.013), and are not significantly different across these four treatments (K-Wallis test p = 0.728). This is consistent with our 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 and LR treatments.

Table 15 in the Online Appendix confirms the previous results with a regression of suppliers', buyers', and total profits on treatment dummies controlling for period, round, and subject random effects. Suppliers' profits only increase relative to the SR benchmark in the ENG and LR treatments (ENG: $\beta = 5.843$, p < 0.001; LR: $\beta = 3.962$, p = 0.003), while buyers' profits increase in all the Buyer-as-Two-Employees treatments, as well as in the LR treatment (PM: $\beta = 7.455$, p = 0.005; 50-50: $\beta = 9.132$, p < 0.001; ENG: $\beta = 9.801$, p = 0.001; LR: $\beta = 8.993$, p = 0.001).

Taken together, the results imply that total profits are higher in all the treatments where the firms have a long term relationship—for any allocation of decision rights—capturing the efficiency

²⁸ Kruskal-Wallis equality-of-populations rank test across multiple samples.

Table 8 Average Profits

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

(*)Refers to the expected profits computed analytically for the benchmark case of lowest collaboration, where the supplier never shares and the buyer always chooses to compete. Number of observations for each role: 26 in SR, 28 in PM, 31 in 50-50, 27 in ENG, 31 in LR.

gains of long-term supplier relations. In particular, total profits are higher in the PM and 50-50 treatments than in the SR benchmark, and are even higher in the ENG and LR treatments (K-Wallis test p=0.005). For the supplier to be able to benefit from this increased surplus, the engineer needs to be in charge of managing the relationship. In this case, the increased supplier profit does not come at a cost for the buyer.

6. Discussion and Concluding Remarks

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 the allocation of decision rights to employees matters. Overall, we confirm that there is a trend of increased collaboration from procurement manager control, to joint control, to engineer control, as predicted by the theory. However, we observe two main deviations from this result. First, suppliers do not share more in the joint control case than in the procurement manager control case, and in the engineer control case they share only marginally more than in the procurement manager control case. Second, we find that both procurement managers and engineers behave significantly less collaboratively in the joint control case than when the procurement manager

or when the engineer have full control. We explore two potential explanations for this result. The first one is diffusion of responsibility, by which an individual is inhibited from helping others when another individual is also available to help. To explore this potential explanation, we designed a Buyer-as-Two-Employees treatment motivated by a manipulation in Dana et al. (2007), where if at least one of the two employees recommends not to compete, that decision is implemented. Otherwise the buyer competes. Since both employees can unilaterally ensure the decision not to compete, this setup allows engineers to reduce their own psychological personal responsibility for the decision to compete. Therefore, behavior that is consistent with diffusion of responsibility would show that the engineer competes more in the new treatment (where the other employee is also able to help the supplier) than when his own decision is always implemented. On the contrary, we find that this treatment does not significantly change the engineer's behavior, suggesting the diffusion of responsibility is not the main behavioral driver at play. The second mechanism that we consider is social interaction effects, by which the procurement manager and the engineer might be biased to make their recommendations match. To explore this possibility, we conduct a second additional treatment which allows for richer interactions between the two employees to arise. The new treatment is equivalent to the treatment with equal split of decision rights but, in addition, the engineer and the procurement manager are allowed to chat before making their recommendations. Since the employees have divergent incentives in their relationship with the supplier, they have no strategic incentives to interact or to try to reach an agreement. Nevertheless, we find that employees actively engage in the chats and seek to reach an agreement. Furthermore, an analysis of the chats suggests that an agreement to compete is easier to reach and more likely to be honored than an agreement not to compete. These results suggest that when both employees have equal decision rights, social interaction effects in the form of a bias to agreement arise which hamper the employees' collaborative behavior.

We then compare the cases of procurement manager and engineer full control with the corresponding benchmarks in a setup where the buyer is a single decision maker. As predicted by the theory, we find that the case where the engineer has full control is not significantly different from the benchmark case where the buyer is a single employee and the two firms have a long-term relationship. When we compare the case where the procurement manager has full control with the case where the buyer is a single employee and the two firms have a sort-term relationship, we find that suppliers share more in the procurement manager control case than in the short-term relationship benchmark. One potential driver of this behavior is social preferences—suppliers' fairness concerns towards the buyer may become more salient in the procurement manager control case as the supplier repeatedly interacts with the same engineer. However, since the buyer is also marginally more

collaborative in the procurement manager control case, the supplier also has strategic incentives to be more collaborative.

Finally, we analyze suppliers' and buyers' profits and total supply chain profits. We find that suppliers benefit only from long-run focused employees, while buyers benefit from any of the joint decision cases. As a result, total supply chain profits are higher in the PM and 50-50 treatments (which benefit the buyer) than in the SR benchmark, and are even higher in the ENG and LR treatments (where both the supplier and the buyer earn higher profits).

Our results suggest that, beyond the efficiency gains associated with a long-term relationship between supply chain partners, the allocation of decision rights to employees within the buying firm matters. Suppliers are willing to share the innovation even when a short-run focused employee is in charge of making the decision on the buyer side—and, interestingly, this decision pays off for the supplier. Relative to this case, shifting the decision rights to a long-run focused employee significantly helps the supplier without making the buyer worse off, resulting in significant efficiency gains. Our results also show that uncertainty about which employee will be the final decision-maker leads to less collaboration by both types of employees. When given the opportunity to communicate with each other, employees try to reach an agreement and the non-collaborative action is more likely to prevail. This more competitive behavior does not benefit the buyer and it is particularly detrimental to suppliers' profits. Understanding the interplay between employees is, therefore, key for a buyer deciding whether (and how) to build teams to manage his supplier relations.

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.
- Barron, Greg, Eldad Yechiam. 2002. Private e-mail requests and the diffusion of responsibility. *Computers in Human Behavior* **18**(5) 507–520.
- 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, Gary E, Axel Ockenfels. 2000. Erc: A theory of equity, reciprocity, and competition. *American economic review* **90**(1) 166–193.
- 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.
- Bornstein, Gary, Ilan Yaniv. 1998. Individual and group behavior in the ultimatum game: are groups more ?rational? players? Experimental Economics 1(1) 101–108.
- 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 interorganizational 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.
- Charness, Gary, Luca Rigotti, Aldo Rustichini. 2007. Individual behavior and group membership. *American Economic Review* 97(4) 1340–1352.
- Charness, Gary, Matthias Sutter. 2012. Groups make better self-interested decisions. *Journal of Economic Perspectives* **26**(3) 157–76.
- Chen, Yan, Sherry Xin Li. 2009. Group identity and social preferences. *American Economic Review* **99**(1) 431–57.
- Cialdini, Robert B, Noah J Goldstein. 2004. Social influence: Compliance and conformity. *Annu. Rev. Psychol.* **55** 591–621.
- Cooper, David J., John H. Kagel. 2005. Are two heads better than one? team versus individual play in signaling games. *American Economic Review* **95**(3) 477–509. doi:10.1257/0002828054201431. URL http://www.aeaweb.org/articles?id=10.1257/0002828054201431.
- Cooper, David J, Kai-Uwe Kuhn. 2009. Communication, renegotiation, and the scope for collusion.
- Cooper, David J, Mari Rege, et al. 2011. Misery loves company: Social regret and social interaction effects in choices under risk and uncertainty. *Games and Economic Behavior* **73**(1) 91–110.

- 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 ÅŹ dilemma. Available~at~SSRN~2292390~.
- Dana, Jason, Roberto A Weber, Jason Xi Kuang. 2007. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* **33**(1) 67–80.
- Darley, John M, Bibb Latané. 1968. Bystander intervention in emergencies: diffusion of responsibility. *Journal of personality and social psychology* 8(4p1) 377.
- 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, Klaus M Schmidt. 1999. A theory of fairness, competition, and cooperation. *The quarterly journal of economics* **114**(3) 817–868.
- Fischbacher, U. 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2) 171–178.
- Fréchette, Guillaume R, Sevgi Yuksel. 2017. Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Experimental Economics* **20**(2) 279–308.
- 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~.
- Hurley, Dennis, Bem P Allen. 1974. The effect of the number of people present in a nonemergency situation.

 The Journal of Social Psychology 92(1) 27–29.
- Hyndman, Kyle, Dorothée 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.
- 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 transitionâÅTthe new geography of auto production. Chicago Fed Letter 229.
- Kocher, Martin G, Matthias Sutter. 2007. Individual versus group behavior and the role of the decision making procedure in gift-exchange experiments. *Empirica* **34**(1) 63–88.

- 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.
- Latane, Bibb, John M Darley. 1968. Group inhibition of bystander intervention in emergencies. *Journal of personality and social psychology* **10**(3) 215.
- Latané, Bibb, Steve Nida. 1981. Ten years of research on group size and helping. *Psychological bulletin* **89**(2) 308.
- 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.
- 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 Çagri 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 Busine: 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.

7. Appendix

7.1. Proofs for Equilibria, E1 and E2 - Buyer-as-Two-Employees setup:

We characterize equilibria for the Buyer-as-Two-Employees setup, where 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. If their recommendations agree, that decision is implemented. If they disagree, the buyer implements one of the two recommendations at random, the procurement manager's recommendation with probability p and the engineer's recommendation with probability p and the engineer's recommendation with probability p and the engineer's recommendation with probability p and the engineer has full control, and when p=0.5 decision rights are equally split between the two employees. We show that in the Buyer-as-Two-Employees setup with p=0 and p=0.5, two collaborative equilibria, E1 and E2, can be supported with the payoffs of the stage game depicted in Figure 1 and a probability of innovation p=0.75.

In both equilibria, 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 if 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. While in each round of play of the stage game the procurement manager's recommendation is implemented with probability 0.5, the procurement manager's recommendation never triggers abandonment of the collaborative state. We consider equilibria where the supplier's and the engineer's decisions to stay in the collaborative state depend only on the supplier's and the engineer's previous actions.

The supplier and the engineer play trigger strategies, where they choose the collaborative 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. Therefore, if either player defects, the punishment state will commence in the next period and will continue for the rest of the game. 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—as 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 nonskeptical 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.

Based on this, we 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:

$$\begin{split} \hat{\delta}_1 &= \frac{ISC_m - ISN_m}{(\pi p + 1)ISC_m - (1 - \pi + \pi p)ISN_m - \pi INC_m}, \text{ and } \\ \hat{\delta}_2 &= \frac{ISC_m - ISN_m}{[(\pi (1 - p) - 1)ISN_m + (\pi p + 1)ISC_m - \pi INC_m + (1 - \pi)(1 - p)NN_m + (1 - \pi)(p - 1)NC_m]}. \end{split}$$

We next characterize two collaborative equilibria, E1 and E2, and show that they can be supported under the condition presented in Proposition 1.

If $\delta \geq \hat{\delta}_1$, there exists a subgame perfect Nash Equilibrium, E1, 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. In 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".

If $\delta \geq \hat{\delta}_2$, there exists a subgame perfect Nash Equilibrium, E2, 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. In 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 show below that both equilibria can be supported with the payoffs of the stage game depicted in Figure 1 and $\pi = 0.75$ in the cases of p = 0 and p = 0.5 (note that for the values of the parameters used in the experiment, $\hat{\delta}_2 > \hat{\delta}_1$, and therefore $\delta > \hat{\delta}_2$ in Proposition 1 guarantees that both equilibria, E1 and E2, are supported).

7.1.1. 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 $[pISC_s + (1-p)ISN_s] + \frac{\delta}{(1-\delta)}[\pi(pISC_s + (1-p)ISN_s) + (1-\pi)NC_s] \ge INC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$, or equivalently $p \le \frac{ISN_s INC_s}{ISN_s ISC_s} = 0.68$. Therefore, with our payoffs, the inequality always holds for p = 0 and p = 0.5.
- **E1.b)** The engineer does not want to deviate from "not compete" if the supplier shared as long as $ISN_m + \frac{\delta}{(1-\delta)} [\pi(pISC_m + (1-p)ISN_m) + (1-\pi)NC_m] \ge ISC_m + \frac{\delta}{(1-\delta)} [\pi INC_m + (1-\pi)NC_m],$ or equivalently $\delta \ge \frac{ISC_m ISN_m}{(\pi p + 1)ISC_m (1-\pi + \pi p)ISN_m \pi INC_m}$. With our payoffs, that is $\delta > 0.55$ for p = 0.5, and $\delta > 0.69$ for p = 0.
- **E1.c)** The engineer does not want to deviate from "compete" if the supplier did not share as long as $NC_m + \frac{\delta}{(1-\delta)} [\pi(pISC_m + (1-p)ISN_m) + (1-\pi)NC_m] \ge NN_m + \frac{\delta}{(1-\delta)} [\pi INC_m + (1-\pi)NC_m]$. With our payoffs, the inequality holds when p=0 and p=0.5 for all δ since $NC_m > NN_m$ and $pISC_m + (1-p)ISN_m > INC_m$.

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 $INC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s] \ge ISC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. The inequality always holds with our payoffs since $INC_s > ISC_s$.
- **E1.e)** The engineer does not want to deviate from "compete" as long as $[\pi INC_m + (1-\pi)NC_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m] \ge [\pi INN_m + (1-\pi)NN_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$. The inequality holds with our payoffs since $INC_m > INN_m$ and $NC_m > NN_m$.

- 7.1.2. 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 $[pISC_s + (1-p)ISN_s] + \frac{\delta}{(1-\delta)}[\pi(pISC_s + (1-p)ISN_s) + (1-\pi)(pNC_s + (1-p)NN_s)] \ge [pINC_s + (1-p)INN_s] + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. Since $pNC_s + (1-p)NN_s > NC_s$ and $pINC_s + (1-p)INN_s > INC_s$, a sufficient condition for the equilibrium to be supported for all δ is $pISC_s + (1-p)ISN_s > pINC_s + (1-p)INN_s$, or equivalently $p < \frac{ISN_s INN_s}{ISN_s + INC_s ISC_s INN_s}$. With our payoffs that is p < 0.67, which holds for the focal cases p = 0 and p = 0.5.
- **E2.b)** The engineer does not want to deviate from "not compete" if the innovation was shared as long as $ISN_m + \frac{\delta}{(1-\delta)}[\pi(pISC_m + (1-p)ISN_m) + (1-\pi)(pNC_m + (1-p)NN_m)] \ge ISC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$ or equivalently, as long as $\delta \ge \frac{ISC_m ISN_m}{[(\pi(1-p)-1)ISN_m + (\pi p+1)ISC_m \pi INC_m + (1-\pi)(1-p)NN_m + (1-\pi)(p-1)NC_m]}$. With our payoffs, this is equivalent to $\delta \ge 0.55$ for p = 0.5 and $\delta \ge 0.70$ for p = 0.
- **E2.c)** The engineer does not want to deviate from "not compete" if the innovation was not shared as long as $NN_m + \frac{\delta}{(1-\delta)}[\pi(pISC_m + (1-p)ISN_m) + (1-\pi)(pNC_m + (1-p)NN_m)] \ge NC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$ or equivalently, as long as $\delta \ge \frac{NC_m NN_m}{[\pi(1-p)ISN_m + \pi pISC_m \pi INC_m + [(1-\pi)(1-p)-1]NN_m + [(1-\pi)p + \pi]NC_m]}$. With our payoffs, this is $\delta \ge 0.10$ for p = 0.5 and $\delta \ge 0.17$ for p = 0.

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 $INC_s + \frac{\delta}{(1-\delta)} [\pi INC_s + (1-\pi)NC_s] \ge ISC_s + \frac{\delta}{(1-\delta)} [\pi INC_s + (1-\pi)NC_s]$. With our payoffs, the inequality holds since $INC_s > ISC_s$.
- **E2.e)** The engineer does not want to deviate from "compete" as long as $[\pi INC_m + (1-\pi)NC_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m] \ge [\pi INN_m + (1-\pi)NN_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m]$. With our payoffs, the inequality holds since $INC_m > INN_m$ and $NC_m > NN_m$.

7.2. Model with Bias to Agreement

We consider the case where the engineer and the procurement manager get an additional utility (beyond their monetary payoffs) if their recommendations match. We capture this utility with a behavioral parameter, $\beta_i \in [0, \bar{\beta}]$ for employee $i, i \in \{E = \text{engineer}, P = \text{procurement manager}\}$, which is common knowledge to all players. Specifically, in each round where their recommendations match, the engineer and the procurement manager earn their monetary payoff plus an extra "bonus", β_i , each.

We focus on the 50-50 case (if their recommendations agree, that decision is implemented, and if they disagree one of the two recommendations is chosen at random, both with equal probability), where both employees' decisions are payoff-relevant. We analyze how the bias towards agreement

affects the conditions under which the following equilibria can arise: 1) **agreement to compete**: a non-collaborative equilibrium where the procurement manager and the engineer agree to compete in every round and the supplier never shares the innovation, 2) **agreement not to compete**: a collaborative equilibrium where the procurement manager and the engineer agree not to compete if the supplier shared the innovation (and either compete in the ambiguous node consistent with equilibrium E1, or not compete in the ambiguous node consistent with equilibrium E2) and the supplier shares the innovation whenever there is one, and 3) **disagreement despite bias:** a collaborative equilibrium where the engineer does not compete if the supplier shared and the supplier shares whenever there is an innovation, and the procurement manager plays the non-collaborative strategy of the single-shot game. Note that in this last case the equilibria E1 and E2 characterized in Proposition 1 can be supported despite firms' bias towards agreement.

7.2.1. Agreement to compete:

Proposition 2. Equilibrium with Agreement to Compete:

Under the values of the parameters used in the experiment, there exists a subgame perfect Nash Equilibrium where the procurement manager and the engineer compete in every round, and the supplier never shares the innovation.

Note that this equilibrium resembles the punishment state of the trigger strategies in the collaborative equilibrium, where the collaborative state cannot be reinstated if an employee deviates to "not compete" or if the supplier starts sharing.

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

The engineer does not want to deviate from compete as long as $[\pi INC_m + (1-\pi)NC_m + \beta_E] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m + \beta_E] \ge [\pi INN_m + (1-\pi)NN_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m + \beta_E].$ With our payoffs, the inequality always holds with our parameters since $INC_m > INN_m$ and $NC_m > NN_m$.

The procurement manager does not want to deviate from compete as long as $[\pi INC_m + (1 - \pi)NC_m + \beta_P] \ge [\pi INN_m + (1 - \pi)NN_m]$. With our payoffs, the inequality always holds since $INC_m > INN_m$ and $NC_m > NN_m$.

7.2.2. Agreement not to compete: We consider two equilibria where the procurement manager and engineer agree not to compete if the supplier shared, one where they agree to compete in the ambiguous node (analogous to E1) and another where they agree not to compete in the ambiguous node (analogous to E2).²⁹

²⁹ Note that while the engineer and the supplier play trigger strategies in an infinitely repeated game, the procurement manager joins the relationship for only one round. Nevertheless, we assume that the procurement manager knows

Proposition 3. Equilibrium with Agreement Not to Compete:

[Equilibrium Analogous to E1:] If $\beta_E \in [0,81.5)$ and $\beta_P \geq 56$, then with the values of the parameters used in the experiment, 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. If the engineer and the supplier are in the collaborative state, the procurement manager matches the engineer's Collaborative-Skeptical strategy. If the engineer and the supplier are in the punishment state, the procurement manager chooses to compete in every round.

[Equilibrium Analogous to E2:] If $\beta_E \in [0, 29.25)$ and $\beta_P \geq 56$, then with the values of the parameters used in the experiment, 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. If the engineer and the supplier are in the collaborative state, the procurement manager matches the engineer's Collaborative-Non-Skeptical strategy. If the engineer and the supplier are in the punishment state, the procurement manager chooses to compete in every round.

Proof.

[Equilibrium analogous to E1] 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 $ISN_s + \frac{\delta}{(1-\delta)}[\pi ISN_s + (1-\pi)NC_s] \ge INC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. Or equivalently, $ISN_s \ge INC_s$, which holds with our payoffs.
- **E1.b)** The engineer does not want to deviate from "not compete" if the supplier shared as long as $ISN_m + \beta_E + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1-\pi)NC_m + \beta_E] \ge ISC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m + \beta_E],$ or equivalently, as long as $\delta \ge \frac{ISC_m ISN_m \beta_E}{[ISC_m \pi INC_m (1-\pi)ISN_m \beta_E]}$ —with the values of the payoffs of our stage game, this is $\delta \ge \frac{56 \beta_E}{81.5 \beta_E}$. For $\delta = 0.75$, the inequality holds for $\beta_E \in [0, 81.5)$ and the incentive compatibility constraint is easier to satisfy as β_E increases in this range.
- **E1.c)** The engineer does not want to deviate from "compete" if the supplier did not share as long as $NC_m + \beta_E + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1-\pi)NC_m + \beta_E] \ge NN_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m + \beta_E]$. The inequality holds with our payoffs since $NC_m > NN_m$ and $ISN_m > INC_m$.
- **E1.d)** The procurement manager does not want to deviate from "not compete" if the supplier shared as long as $ISN_m + \beta_P \ge ISC_m$, or equivalently with our payoffs $\beta_P \ge ISC_m ISN_m = 56$.

whether the engineer and procurement are in the "collaborative state" or in the "punishment state" of their trigger strategies before making his recommendation. This captures the realistic setting where the procurement manager is aware of previous history with the supplier. To allow for this in the experiment, whenever a procurement manager joined an existing relationship (rounds two onwards within a relationship), they were informed of all the players' decisions in that relationship in the previous round.

E1.e) The procurement manager does not want to deviate from "compete" in the ambiguous node as long as $NC_m + \beta_P \ge NN_m$, which holds for all β_P since $NC_m > NN_m$.

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

- **E1.f)** The supplier does not want to deviate from "not share" as long as $INC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s] \ge ISC_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. With our payoffs, the inequality holds since $INC_s > ISC_s$.
- **E1.g)** The engineer does not want to deviate from "compete" as long as $[\pi INC_m + (1-\pi)NC_m + \beta_E] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m + \beta_E] \ge [\pi INN_m + (1-\pi)NN_m] + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m + \beta_E]$. With our payoffs, the inequality always holds with our parameters since $INC_m > INN_m$ and $NC_m > NN_m$.
- **E1.h)** The procurement manager does not want to deviate from "compete" as long as $\pi INC_m + (1-\pi)NC_m + \beta_P \ge [\pi INN_m + (1-\pi)NN_m]$. With our payoffs, the inequality holds for all β_P since $INC_m > INN_m$ and $NC_m > NN_m$.

[Equilibrium analogous to E2] We 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 $ISN_s + \frac{\delta}{(1-\delta)}[\pi ISN_s + (1-\pi)NN_s] \ge INN_s + \frac{\delta}{(1-\delta)}[\pi INC_s + (1-\pi)NC_s]$. The inequality always holds with our payoffs since $ISN_s > INN_s$, $ISN_s > INC_s$, and $NN_s > NC_s$.
- **E2.b)** The engineer does not want to deviate from "not compete" if the supplier shared as long as $ISN_m + \beta_E + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1-\pi)NN_m + \beta_E] \geq ISC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m + \beta_E]$ or equivalently, as long as $\delta \geq \frac{ISC_m ISN_m \beta_E}{[ISC_m (1-\pi)ISN_m + (1-\pi)NN_m \pi INC_m (1-\pi)NC_m \beta_E]}$ —with our payoffs, this is $\delta \geq \frac{56 \beta_E}{80.25 \beta_E}$. When $\delta = 0.75$, the inequality holds for $\beta_E \in [0, 80.25)$ and the incentive compatibility constraint is easier to satisfy as β_E increases in this range.
- **E2.c)** The engineer does not want to deviate from "not compete" if the supplier did not share as long as $NN_m + \beta_E + \frac{\delta}{(1-\delta)}[\pi ISN_m + (1-\pi)NN_m + \beta_E] \ge NC_m + \frac{\delta}{(1-\delta)}[\pi INC_m + (1-\pi)NC_m + \beta_E]$ or equivalently, as long as $\delta \ge \frac{NC_m NN_m \beta_E}{\pi[ISN_m INC_m NN_m + NC_m] \beta_E}$ —with our payoffs, this is $\delta \ge \frac{5-\beta_E}{29.25-\beta_E}$. When $\delta = 0.75$, the inequality holds for $\beta_E \in [0, 29.25)$ and the incentive compatibility constraint is easier to satisfy as β_E increases in this range.
- **E2.d)** The procurement manager does not want to deviate from "not compete" if the supplier shared as long as $ISN_m + \beta_P \ge ISC_m$, or equivalently with our payoffs $\beta_P \ge ISC_m ISN_m = 56$.
- **E2.e)** The procurement manager does not want to deviate from "not compete" in the ambiguous node as long as $NN_m + \beta_P \ge NC_m$, or equivalently with our payoffs $\beta_P \ge NC_m NN_m = 5$.

The conditions for the punishment state are as in the previous case and thus omitted.

7.2.3. Disagreement despite bias to agreement: In this section we characterize the conditions under which equilibria analogous to those in Proposition 1 (where the engineer plays collaborative trigger strategies and the procurement manager plays the non-collaborative strategy of the single shot game) can be supported despite the employees' bias to agreement. Note that in this equilibrium the engineer and the procurement manager disagree in the collaborative state of the supplier-engineer relationship. However, the engineer's and the procurement manager's recommendations agree in the punishment state.

Proposition 4. Equilibrium with Disagreement despite Bias to Agreement:

[Equilibrium Analogous to E1:] If $\beta_E \in [0, 25.69]$ and $\beta_P \leq 28$, then with the values of the parameters used in the experiment, 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. The procurement manager plays the non-collaborative strategy of the single shot game.

[Equilibrium Analogous to E2:] If $\beta_E \in [0, 20.40]$ and $\beta_P \leq 2.5$, then with the values of the parameters used in the experiment, 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. The procurement manager plays the non-collaborative strategy of the single shot game.

Proof.

[Equilibrium analogous to 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 $ISN_m + \frac{\delta}{(1-\delta)} [\pi(\frac{ISN_m + ISC_m}{2}) + (1-\pi)(NC_m + \beta_E)] \ge ISC_m + \beta_E + \frac{\delta}{(1-\delta)} [\pi INC_m + (1-\pi)NC_m + \beta_E]$, or equivalently $\delta \ge \frac{ISC_m ISN_m + \beta_E}{(\frac{\pi}{2} 1)ISN_m + (\frac{\pi}{2} + 1)ISC_m \pi INC_m + (1-\pi)\beta_E}$, which with our payoffs is $\frac{56 + \beta_E}{102.5 + (1-\pi)\beta_E}$. For $\delta = 0.75$, the inequality holds for $\beta_E \in [0, 25.69]$ and the incentive compatibility constraint is harder to satisfy as β_E increases in this range.
- **E1'.c)** The engineer does not want to deviate from "compete" if the supplier did not share as long as $NC_m + \beta_E + \frac{\delta}{(1-\delta)} \left[\pi \frac{(ISN_m + ISC_m)}{2} + (1-\pi)(NC_m + \beta_E)\right] \ge NN_m + \frac{\delta}{(1-\delta)} \left[\pi INC_m + (1-\pi)NC_m + \beta_E\right]$, or equivalently with our payoffs, $\delta \ge \frac{NN_m NC_m \beta_E}{NN_m NC_m + \frac{\pi}{2}ISN_m + \frac{\pi}{2}ISC_m \pi INC_m (1-\pi)\beta_E}$ —which with our parameters is $\delta \ge \frac{-5-\beta_E}{41.5-(1-\pi)\beta_E}$. With $\delta = 0.75$, the inequality holds for $\beta_E \in [0, 166)$ and the incentive compatibility constraint becomes easier to satisfy as β_E increases in this range.

- **E1'.d)** The procurement manager does not want to deviate from "compete" if the supplier shared as long as $\frac{ISC_m+ISN_m}{2} \ge ISN_m + \beta_P$, or equivalently with our payoffs, $\beta_P \le \frac{112+56}{2} 56 = 28$.
- **E1'.e)** The procurement manager does not want to deviate from "compete" in the ambiguous node as long as $NC_m + \beta_P \ge \frac{NN_m + NC_m}{2}$, which holds for all β_P since $NC_m > NN_m$.

The conditions for the punishment state are as in the previous case and thus omitted.

[Equilibrium analogous to 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 $\frac{ISN_s+ISC_s}{2} + \frac{\delta}{(1-\delta)} \left[\pi(\frac{ISN_s+ISC_s}{2}) + (1-\pi)(\frac{(NN_s+NC_s)}{2})\right] \ge \frac{INN_s+INC_s}{2} + \frac{\delta}{(1-\delta)} \left[\pi INC_s + (1-\pi)NC_s\right]$. 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 $ISN_m + \frac{\delta}{(1-\delta)} \left[\pi \frac{(ISN_m + ISC_m)}{2} + (1-\pi) \frac{(NN_m + NC_m)}{2}\right] \ge ISC_m + \beta_E + \frac{\delta}{(1-\delta)} \left[\pi INC_m + (1-\pi)NC_m + \beta_E\right]$ or equivalently, as long as $\delta \ge \frac{ISC_m ISN_m + \beta_E}{\left[(\frac{\pi}{2} 1)ISN_m + (\frac{\pi}{2} + 1)ISC_m \pi INC_m + \frac{(1-\pi)}{2}NN_m + \frac{(\pi-1)}{2}NC_m\right]}$ with our payoffs, this is $\delta \ge \frac{56 + \beta_E}{101.875}$. With $\delta = 0.75$, the inequality holds for $\beta_E \in [0, 20.40]$ and the incentive compatibility constraint becomes harder to satisfy as β_E increases in this range.
- **E2'.c)** The engineer does not want to deviate from "not compete" if the innovation was not shared as long as $NN_m + \frac{\delta}{(1-\delta)} \left[\pi \frac{(ISN_m + ISC_m)}{2} + (1-\pi) \frac{(NN_m + NC_m)}{2}\right] \ge NC_m + \beta_E + \frac{\delta}{(1-\delta)} \left[\pi INC_m + (1-\pi)NC_m + \beta_E\right]$ or equivalently, as long as $\delta \ge \frac{NC_m NN_m + \beta_E}{\left[\frac{\pi}{2}ISN_m + \frac{\pi}{2}ISC_m \pi INC_m \frac{(\pi+1)}{2}NN_m + \frac{(\pi+1)}{2}NC_m\right]}$. With our payoffs, this is $\delta \ge \frac{5+\beta_E}{50.875}$. With $\delta = 0.75$, the inequality holds for $\beta_E \in [0, 33.15]$ and the incentive compatibility constraint is harder to satisfy as β_E increases in that range.
- **E2'.d)** The procurement manager does not want to deviate from "compete" if the supplier shared as long as $\frac{ISC_m+ISN_m}{2} \ge ISN_m + \beta_P$, or equivalently with our payoffs, $\beta_P \le \frac{112+56}{2} 56 = 28$.
- **E2'.e)** The procurement manager does not want to deviate from "compete" in the ambiguous node as long as $\frac{NC_m+NN_m}{2} \ge NN_m + \beta_P$, or equivalently with our payoffs, $\beta_P \le \frac{12+17}{2} 12 = 2.5$.

The conditions for the punishment state are as in the previous case and thus omitted.