An Experimental Investigation of Managing Quality Through Monetary and Relational Incentives

Andrew M. Davis
Johnson School of Management, Cornell University, adavis@cornell.edu, https://courses.cit.cornell.edu/amd365/

Kyle Hyndman
Jindal School of Management, University of Texas at Dallas, KyleB.Hyndman@utdallas.edu, http://www.hyndman-honhon.com/Main/Kyle_Hyndman.html

We investigate the efficacy of monetary and relational incentives for managing the quality of a product in a two-tier supply chain. In our setting, a retailer offers a supplier contract terms for a product, where the product can be low or high quality. The supplier can choose to exert high effort, which is costly but guarantees high quality, or low effort, which does not assure high quality with certainty. We compare how monetary incentives, such as a bonus that is paid to the supplier when high quality is received by the retailer, and relational incentives, such as the two parties engaging in a long-term relationship where there is the threat of punishment, affect overall quality and supply chain efficiency. Two of our primary results suggest that (1) relational incentives improve both quality and supply chain efficiency, regardless of whether a monetary incentive is present, and (2) when relational incentives are present, the impact of adding monetary incentives is non-monotonic: less efficient monetary incentives appear to crowd out the benefits of relational incentives leading even to a *reduction* in supply chain efficiency, whereas more efficient monetary incentives actually complement relational incentives and lead to significant *increases* in both quality and supply chain efficiency. We then proceed by demonstrating that a behavioral model of fairness can organize the data quite well.

Key words: behavioral operations; quality; incentives; supply chain contracting; experimental economics

1. Introduction

Retailers and manufacturers are constantly trying to identify ways to improve the quality of products they receive from suppliers. By doing so they not only increase their overall revenues, but also avoid significant costs and negative impacts to their reputation, such as when Hy Vee, a retail grocery chain with annual revenue of $9.3 billion, announced in 2016 that some of its private label products may be contaminated with Listeria monocytogenes (a bacteria that can cause serious or fatal infections), and was due to its supplier SunOpta (FDA 2016). When product quality concerns arise, the retailer or manufacturer is often held responsible, even when the defect may stem from the supplier. As a further example, in 2002 Ford publicly stated that its suppliers were responsible for Ford's recent decline in quality, claiming that 76% of quality issues came from suppliers (Sherefkin 2002). Nevertheless, in the year prior, Ford incurred annual losses of $5.45 billion and also ousted its CEO. Examples such as these make clear that it is in a retailer or manufacturer's best interest to utilize mechanisms which ensure suppliers provide high quality products.

The fundamental problem when managing quality in supply chains is due to the fact that the parties' incentives are misaligned. The retailer or manufacturer (hereafter "retailer" for simplicity) wants high effort from the supplier, so that quality is high, but high effort is costly for the supplier. This is further exacerbated by the fact that effort is rarely observable and, hence, not contractible. As a consequence, the retailer must consider alternative ways to address this misalignment of incentives (Girotra and Netessine (2014) p. 148). The most obvious alternative is a monetary incentive, such as when a payment to the supplier is made conditional on some variable – such as quality or profit – which is observable and positively correlated with effort. For instance, one recent article on how to manage supplier relationships in China suggests including monetary payments for performance (Chaudhuri 2013), and another specifically recommends offering a bonus when suppliers comply with quality requirements (ChinaImportal 2014). However, as we will demonstrate, monetary incentives may not always be desirable, as they may lead to a loss of supply chain efficiency, relative to the first-best outcome. On the other hand, if the parties are engaged in a long term relationship, then an alternative to monetary incentives is readily available – namely, relational incentives. In this case, the credible threat of future punishment, such as never buying from the supplier again, may be sufficient to induce high quality products from suppliers. Indeed, anecdotal evidence suggests that long-term contracts with favorable suppliers are preferred when attempting to ensure high quality products (Stallkamp 2005, CEB 2014).

In this study, we focus on two primary research questions for managing quality in supply chains. First, theoretically, under what conditions can relational incentives induce high quality from suppliers and improve supply chain efficiency as compared to monetary incentives? By addressing this first research question we contribute to the rich stream of theoretical work on quality management in supply chains (e.g. Reyniers and Tapiero (1995), Balachandran and Radhakrishnan (2005), Zhu et al. (2007), Chao et al. (2009)). For our second research question, we deviate from the theoretical literature and take a behavioral perspective: when allowing human decision makers to act as retailers and suppliers, how do monetary incentives, relational incentives, and a combination of both, affect overall quality and supply chain efficiency? Through answering this second research question we take an important first step in understanding how to manage quality in supply chains when behavioral factors may be at play.

We begin by first considering a setting where a retailer and supplier are engaged in a one-shot interaction. In this scenario, the primary mechanism a retailer has to induce high effort from suppliers is a monetary incentive. However, in such a setting, there will generally be a supply chain efficiency loss, relative to first-best where effort can be directly contracted. In particular, retailers are often larger than their suppliers, which can lead to a supply chain efficiency loss for two reasons.[1] First, suppliers may be more risk averse than retailers, and require extra compensation for any risk exposure, which reduces supply chain efficiency. Second, suppliers may have less favorable access to financing. If a retailer pays the supplier a bonus if agreed-upon performance metrics have been satisfied over a review period, then each dollar of the bonus paid by the retailer is valued at less than a dollar by the supplier, also leading to a supply chain efficiency loss. We will refer to this latter scenario as a "deferred payment mechanism." The deferred payment mechanism framework has been studied extensively in the modeling literature on quality management in supply chains (Babich and Tang 2012, Rui and Lai 2015), and has the flexibility to capture a number of features of other contracts. For instance, the deferred payment framework can effectively be used for investigating service-level agreements, which are used by 91% of organizations to manage suppliers (Oblicore 2007), contingency payment contracts, which have been the focus of studies on supplier responsibility (Chen and Lee 2016), and trade credit contracts, where retailers can evaluate the quality of the product during a review period before deciding to pay (Lee and Stowe 1993, Klapper et al. 2010). Given its ability to capture so many different contracts studied in the academic literature and its use in practice, we adopt the deferred payment mechanism as our lens for generating theoretical insights and studying human behavior on quality management in supply chains.

We subsequently consider the case in which a retailer and supplier are engaged in a long-term relationship. Past studies have pointed out a number of potential benefits of long term relationships in supply chains, such as facilitating trust and learning, joint-investments, and better alignment of actions (Han et al. 1993, Kalwani and Narayandas 1995). One key benefit of long term relationships for managing quality is that relational incentives may be used to induce high effort from suppliers, and reduce (or even eliminate), a retailer's reliance on monetary incentives. For example, a retailer may make an upfront unconditional payment each period – which means that the supplier is neither exposed to risk, nor forced to borrow at unfavorable terms to finance operations – and use the credible threat of future punishment (such as never buying from the supplier again) to induce the supplier to exert high effort. Indeed, in operations and supply chain management, several

[1] A recent empirical study by the World Bank analyzed supply chain contracts between retailers and suppliers from North America and Europe (Klapper et al. 2010). They found that retailers are typically quite large compared to suppliers: 84% of retailers have more than \$10 billion in sales, whereas 50% of suppliers have *less* than \$100 million in sales.

theoretical papers have studied the potential benefits of relational contracting (e.g., Taylor and Plambeck (2007a,b), Tunca and Zenios (2006), Gibbons (2005)).

We theoretically investigate these one-shot and long-term settings where a retailer seeks to contract with a supplier for delivery of a product. The product's quality is not immediately ascertainable, so the retailer employs a deferred payment mechanism comprised of a fixed payment and a monetary incentive - a conditional bonus payment - which is paid to the supplier when high quality is received, and is effectively discounted when the supplier (potentially) receives it. The supplier can exert high effort (which is costly) and ensure that a high quality product is produced, or she can exert low effort and probabilistically produce a high quality product. We show that in the one-shot setting the optimal contract leads to a loss in supply chain efficiency, relative to first-best where effort can be directly contracted. However, in our analysis of the indefinitely repeated game between the retailer and supplier, we provide a condition under which the retailer prefers to use relational incentives only – specifically, the threat that the retailer will no longer buy from the supplier in the future if low quality is observed – to induce high effort. Moreover, when this is the case, relational incentives generate the first-best level of supply chain efficiency. Thus, we show that relational incentives can substitute for monetary incentives and, when they do, enhance supply chain efficiency.

Because many quality and supply chain contracting decisions are made by managers in practice, we proceed by conducting a 2×3 between-subjects experiment, where human subjects act as retailers and suppliers. The first dimension we consider is one-shot versus repeated interactions, thus manipulating whether or not there is a relational incentive. The second dimension we vary is the monetary incentive. In our baseline control case, we omit bonuses entirely such that there is no monetary incentive, and retailers must rely solely on fixed payments. In the other two manipulations, we include the monetary incentive (i.e., the bonus) but vary its efficiency: in one set of treatments the bonus is heavily discounted when the supplier receives it, such that the relative efficiency of the monetary incentive is low. In the other set, the bonus is only modestly discounted before the supplier receives it, such that the relative efficiency of the monetary incentive is high (when speaking of the relative efficiency of the bonus, "high" and "low" do not have a concrete meaning in the one-shot case; however, as we will explain, our parameters were chosen because they generate different predictions in the case of repeated interactions).

Our experiment provides several interesting managerial insights. First, we observe that relational incentives (long-term relationships) increase both quality and supply chain efficiency, regardless of whether there is any type of monetary incentive. Furthermore, this improvement often comes without any cost to either party's profits: relational incentives are weakly Pareto improving. Second, in a one-shot setting without a monetary incentive, standard theory predicts that introducing any type of monetary incentive should lead to higher quality and supply chain efficiency. However, our data suggest that when the efficiency of the monetary incentive is low, it actually decreases supply chain efficiency by approximately 5%. Lastly, we observe that when relational incentives are present, the impact from adding monetary incentives is non-monotonic. In particular, when the efficiency of the monetary incentive is low, it crowds out the benefits of the relational incentive, actually leading to a 13.8% *reduction* in supply chain efficiency. On the other hand, when the efficiency of the monetary incentive is high, the monetary incentive appears to complement the relational incentive and increase quality significantly, contributing to a 20.5% *increase* in supply chain efficiency.

Another observation from our experiment is that retailers have a tendency to set contract terms which differ from the standard theoretical predictions. In particular, when the efficiency of the monetary incentive is low, retailers offer relatively larger fixed payments and relatively smaller bonuses than theory predicts, whereas the reverse is true when the efficiency of the monetary incentive is high. To explain this result, we posit that suppliers have a preference for fairness (Fehr and Schmidt 1999, Bolton and Ockenfels 2000). We incorporate this behavioral tendency into the original model, and show that it generates contract proposals which are similar to the ones we observe in our experimental data. As a further robustness check of this hypothesis, we conduct two additional one-shot experimental treatments where we automate the supplier's role so that no fairness concerns should exist. Indeed, in these automated-supplier treatments we observe that retailer offers differ significantly from the original treatments with human suppliers, in that their offers converge remarkably close to the standard theoretical predictions.

2. Related Literature

There are two primary streams of literature that relate most closely to our work: experimental papers which investigate principal-agent settings and supply chain management papers which focus on quality management issues.

Our study falls within the realm of principal-agent contracting, which has been studied in the experimental economics literature. One well-known application is that of gift-exchange, where a principal offers an agent a fixed wage, which will always be paid from the principal to the agent, and the agent exerts effort, which is costly. The normative theory suggests that principals should offer the lowest possible wage and agents exert the lowest possible effort, but many studies have found that principals offer non-zero wages, and agents, in turn, exert higher effort (see, e.g., Charness and Kuhn 2010, for a survey of the literature). Moreover, some of this literature has found that explicit monetary incentives may crowd out any intrinsic motivation or feelings of reciprocity, which may actually lead to lower effort. For instance, Fehr and Gächter (2000) showed that the virtuous trust-reciprocity cycle broke down when, with some probability, a fine could be imposed if the agent had shirked on the contractually specified effort. Irlenbusch and Sliwka (2005) also find evidence of crowding out in a standard principal-agent setting. In particular, Irlenbusch and Sliwka observe that effort levels decrease once the principal offers a piece rate in addition to the standard wage.[2]

The results from our baseline one-shot treatment, which does not include any monetary or relational incentive and is similar to the classic gift-exchange game, show that suppliers have little intrinsic motivation to exert high effort, and retailers respond with low fixed payments. Thus we find evidence that supports the standard theoretical prediction, in contrast to many gift-exchange experiments (Charness and Kuhn 2010). This result is analogous to recent work by Rubin and Sheremeta (2016), who also showed that the amount of reciprocity in a gift-exchange setting diminishes when the output of the agent is partially random, and not perfectly dependent on their effort decision. In our study, when the agent exerts high effort, high quality is guaranteed, while low effort probabilistically generates high quality. Hence, a retailer who receives high quality cannot be sure that it was due to high effort. Consistent with Rubin and Sheremeta (2016), this may diminish trust and reciprocity. Another relevant paper gift-exchange paper is Gächter et al. (2011), who also include treatments where subjects play a repeated game. Unlike us, they find that implicit incentives increase effort only if the contract is non-incentive compatible.

One of the common results across experimental studies that investigate principal-agent settings, is that subjects often exhibit social preferences (e.g., positive reciprocity, fairness preferences or other behavioral factors). For instance, Fehr and Schmidt (2000) consider a game where a principal can specify a wage and desired effort level for the agent, along with a bonus that is not contractually binding. Despite the bonus being non-binding, many agents exert high effort, and principals, in turn, award the optional bonus. Fehr and Schmidt account for this result by applying their own model of social preferences, finding that a large fraction of participants display fairness concerns,

[2] Crowding out has been documented in other areas as well, such as social psychology and public policy. For instance, Bowles (2008) discusses how policy makers should design laws and public policies, emphasizing that extrinsic incentives may fail when they undermine moral values and intrinsic motivations. Additionally, some studies have documented a "crowding in" effect. In particular, Bowles and Polanía-Reyes (2012) conduct an excellent review of fifty experiments on decision making, and find that in some studies, intrinsic and extrinsic incentives may actually complement each other such that there are further improvements, referred to as crowding in. This further motivates our study, as it is not necessarily clear whether certain incentives, such as relational and monetary ones, will be substitutes or complements, in a quality management context.

which has also been documented in other principal agent settings and games (e.g. Fehr and Schmidt (1999, 2004)). In our study, we too find evidence that fairness preferences, which leads retailers to offer contracts that differ from the standard theoretical predictions. Furthermore, in a follow-up set of treatments with automated suppliers, who do not exhibit any fairness concerns, we find that retailers' contract proposals are much closer to the standard theoretical predictions with self-interested suppliers, than in our main experiments with human suppliers.

From a supply chain management standpoint, there has been a considerable amount of theoretical research on managing quality. Mu et al. (2016) investigate how buyers can ensure that milk suppliers do not produce low quality milk, which can lead to hazardous health effects. Baiman et al. (2004) consider a model where a buyer manufactures a product which consists of multiple parts, each produced by a different supplier who can incur a cost to improve their component's quality, and evaluate different contractual structures and how they affect overall efficiency. Kaya and Özer (2009) investigate a setting where the supplier's cost of quality may not be known to the manufacturer, and show, among other things, how the manufacturer's pricing strategy affects quality risk. There are also a number of modelling papers on quality management that focus on cost sharing mechanisms between buyers and suppliers (e.g. Reyniers and Tapiero (1995), Balachandran and Radhakrishnan (2005), Zhu et al. (2007), Chao et al. (2009)). Of the theoretical work done on quality management in supply chains, the most relevant to our study is Babich and Tang (2012), who provide a theoretical analysis comparing inspections and deferred payment mechanisms, ultimately concluding that deferred payment mechanisms are preferred when attempting to prevent adulterated products. Despite the plethora of modeling studies on managing quality in supply chains, and the importance of the topic from a practical standpoint, there has been little research done from a behavioral operations perspective. Of the few behavioral works, Beer et al. (2015) administer a human-subject experiment related to quality, where they focus on the effect of symbolic supplier awards. Additionally, Nosenzo et al. (2016), conduct an experiment on inspection mechanisms in a repeated setting, and identify whether rewards or sanctions are best at preventing a worker from shirking. Compared to these works, we believe our study is unique in investigating monetary and relational incentives, and their effectiveness for managing quality and improving supply chain efficiency, from a behavioral standpoint.

3. Theoretical Analysis

We consider a situation in which a retailer seeks to procure a product from a supplier, the quality of which can either be high or low. The retailer values high quality at v_h and low quality at $v_l < v_h$.

We assume that the supplier can exert either high or low effort. The cost of effort is c_h or c_l with $c_h > c_l$. Suppose that exerting high effort leads to high quality with probability 1 and low effort leads to high quality with probability $\mu \in (0,1)$.[3] Note that the parameters $(v_h, v_l, c_h, c_l, \mu)$ are all exogenous. We proceed under the assumption that all players are self-interested, expected profit maximizers.

In order to induce the supplier to exert high effort, the supplier's payoff must be conditioned on the quality of the product received by the retailer. We adopt the deferred payment mechanism framework to generate insights. By deferring a portion of the supplier's payment until the quality of the product can be ascertained – and paying only if quality is high – the supplier can be incentivized to exert high effort.[4] To keep things simple, we operationalize this as follows. Suppose that the contract – chosen by the retailer and agreed to by the supplier – specifies a fixed payment, F, and a bonus, B, that is paid conditional on high quality being received by the retailer. While the cost of the bonus is B to the retailer, the benefit to the supplier is only αB, where $\alpha \in (0,1]$ is an exogenous parameter, which we refer to as the efficiency of the monetary incentive. Notice that this leads to a loss in supply chain efficiency, relative to first-best where effort can be directly contracted, because a fraction $(1-\alpha)$ of the deferred bonus is lost. In particular, the total supply chain surplus from the supplier choosing high effort is:

$$\begin{aligned}\text{Supply Chain Surplus} &= (v_h - F - B) + (F + \alpha B - c_h)\\ &= v_h - c_h - (1-\alpha)B.\end{aligned}$$

Therefore, supply chain surplus is maximized at $B = 0$. However, when $B = 0$, the supplier exerts low effort because her payoff does not depend on the quality provided and high effort is costly.

We first analyze the case in which the retailer and supplier interact once and refer to this as the *monetary incentives* case because only monetary incentives (i.e., B) can be used to incentivize high effort. This analysis is closely related to Babich and Tang (2012), whose framework is useful in that it allows us to manipulate the efficiency of the monetary incentive through the parameter α. We then consider the case in which the retailer and supplier are engaged in an interaction of indefinite duration. We refer to this as the *relational incentives* case because, in addition to

[3] As noted in the Introduction, our setting closely mimics a service-level agreement (SLA). In this case, one would, equivalently, assume that the retailer wishes to achieve some target service level over a certain review period. The supplier can then exert high effort and guarantee that the target is reached, or low effort, in which case the target is only reached probabilistically. Of course, in a typical SLA, the supplier would have many more actions at her disposal (e.g., a continuum of production levels). We restrict attention to binary decisions only in the interest of simplicity.

[4] In the SLA context, this can be equivalently stated as the supplier is paid a bonus after the review period *if* the target service level is achieved.

direct monetary incentives, high effort can be incentivized by implicit threats of punishment should the retailer receive low quality. We show that monetary and relational incentives are partially substitutable tools for inducing high effort. Because of this, we provide a condition under which the retailer prefers to rely on relational incentives only, by setting $B = 0$, and achieve full supply chain efficiency. When the condition is not satisfied, it may still be possible to incentivize high effort with $B = 0$, but to do so is more costly to the retailer than simply relying on monetary incentives.

3.1. Monetary Incentives: One-Shot Game

The one-shot game proceeds as follows. In the first stage, the retailer proposes a contract $(F, B) \geq 0$, with F and B as defined above. Upon observing the contract, the supplier either accepts or rejects the contract. If the supplier rejects the contract, then both players receive their outside option, $u_0 \geq 0$. If the supplier accepts the contract, then, in the final stage, the parties play the game depicted in Figure 1. That is, the retailer decides whether to buy or not buy the product and, simultaneously, the supplier decides whether to exert high effort or low effort.[5]

Figure 1 The Strategic Game and Payoffs Between a Retailer and Supplier Conditional on a Contract (F, B)

	High Effort (H)	Low Effort (L)
Buy	$v_h - F - B$, $F + \alpha B - c_h$	$\mu(v_h - B) + (1-\mu)v_l - F$, $F + \mu\alpha B - c_l$
Don't	u_0, u_0	u_0, u_0

In order to ensure that the retailer choosing to buy and the supplier choosing to exert high effort (Buy, H) is a subgame perfect equilibrium, the initial contract proposal by the retailer must ensure that (Buy, H) is a Nash equilibrium in the last stage and that both the supplier and the retailer earn at least their outside option. Therefore, the retailer solves the following problem:

$$\min_{F,B \geq 0} F + B \qquad s.t.$$

[5] The presence of the final stage, in which the retailer makes a buy/don't buy decision is non-standard in the principal-agent literature. This feature is necessary to maintain consistency with our repeated game setting (relational incentives) in which the retailer and supplier engage in a long term interaction, and in each period the retailer decides whether he wants to buy or not from the supplier. However, it is also consistent with many real world contracts where the retailer has the option to purchase or cancel an upfront order. For instance, Tsay and Lovejoy (1999) analyze quantity flexible contracts that allow a cancellation of (a portion of) an order and highlight their use by Toyota Motor Corporation, while Eppen and Iyer (1997) study backup agreement contracts where an initial order can be cancelled subsequently for a penalty fee, and their use in the fashion apparel industry. Our setting is similar in that we too assume powerful retailers who can choose to not purchase the product, except without penalty, and instead resulting in both parties receiving an outside option payoff. In any case, we also conducted an experimental session of the modified game in which the supplier's acceptance of the contract means that the retailer is committed to buying. The two treatments were qualitatively the same.

$$(IC)\;\; F+\alpha B-c_h \geq F+\mu\alpha B-c_l$$
$$(IR_s)\;\; F+\alpha B-c_h \geq u_0$$
$$(IR_r)\;\; v_h-F-B \geq u_0$$
$$(HE)\;\; v_h-F-B \geq \mu v_h+(1-\mu)v_l-(c_l+u_0)$$

(IC) is the incentive compatibility constraint, which ensures that the supplier prefers to exert high effort rather than low effort. This can be rewritten as $B \geq {}^{(c_h-c_l)}/_{\alpha(1-\mu)}$, and note that while the bonus depends on α, the effective payment to the supplier, αB, does not. (IR_s) is the supplier's individual rationality constraint, which ensures that the supplier earns at least as much as her outside option by exerting high effort. Together with the non-negativity constraint on F, this can be rewritten as $F \geq \max\{0, c_h+u_0-\alpha B\}$. (IR_r) is the retailer's individual rationality constraint, which ensures that the retailer prefers to induce high effort rather than take his outside option. Finally, (HE) ensures that the retailer prefers to induce high effort by offering the contract (F,B), rather than just providing a fixed payment equal to c_l+u_0 to cover the supplier's cost of low effort. In our experiment, we assume parameter values such that the retailer would rather not participate than accept a guarantee of low effort by the supplier. That is, $\mu v_h+(1-\mu)v_l-(c_l+u_0)<u_0$. In this case, (HE) is automatically satisfied whenever (IR_r) is satisfied. Notice that, so long as $c_h+u_0-\alpha B>0$, the fixed payment, F, is actually *independent* of the value of α. Thus we have,

Proposition 1. *To induce high effort, the retailer's contract must satisfy* $B \geq B^* = {}^{(c_h-c_l)}/_{\alpha(1-\mu)}$ *and* $F \geq F^* = \max\{0, c_h+u_0-\alpha B\}$*. If (*$IR_r$*) and (HE) are satisfied at* (F^*,B^*)*, then there exists a subgame perfect equilibrium in which (i) the retailer offers* (F^*,B^*)*, (ii) the supplier accepts, and (iii) the retailer chooses to buy and the supplier exerts high effort.*

Figures 2(a)–(b) plot the incentive compatibility and individual rationality constraints for the parameters used in our experiment – when we allow retailers to offer bonuses – and the regions where it is optimal for the supplier to reject, or exert low or high effort. The lower (lightly shaded; Region C) triangle represents the area where it is optimal for the supplier to reject, the upper-left (darkly shaded; Region B) area represents the area where the supplier should accept the contract and exert low effort, while the lower-right (moderately shaded; Region A) quadrilateral represents the area where the supplier should accept the contract and exert high effort. We also plot (dashed line) the retailer's iso-cost line at the optimal solution, which passes through the intersection of the (IC) and (IR_s) curves.

Figure 2 **Incentive Compatibility and Individual Rationality For Our Experimental Parameters**

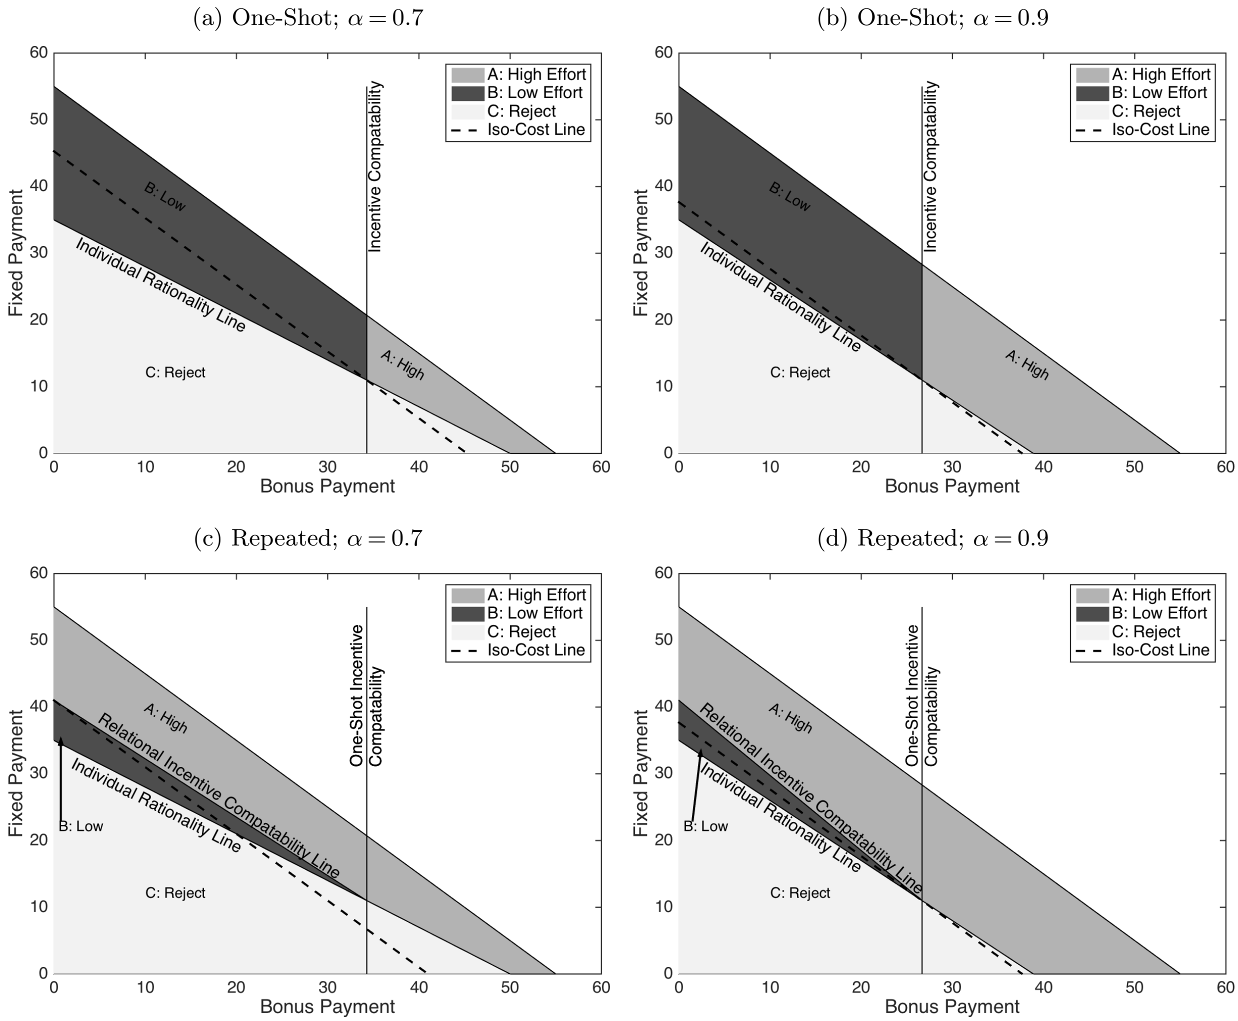

Note: The equation of the outermost line is given by $F = v_h - u_0 - B$. This is because a retailer would never offer so much total compensation that he does not receive at least his outside option.

3.2. Relational Incentives: Indefinitely Repeated Game

Our goal in this section is to determine conditions under which relational incentives substitute for monetary incentives and lead to greater supply chain efficiency. As we will see, while relational incentives expand the region for which the supplier prefers to exert high effort, these incentives are costly for the retailer to provide. Consequently, depending on certain conditions, the retailer may prefer to rely on relational or monetary incentives.

We modify the one-shot game as follows. A retailer and a supplier are paired together for an indefinite duration. As before, there is an initial contracting phase in which the retailer and supplier come to terms on a contract, which specifies F and B. If the contract is rejected, then one period passes and the retailer proposes another contract. Following the conclusion of a contract, in each of an indefinite number of periods, the two parties play the stage game as in Figure 1. For ease of analysis, we will assume that players discount the future with common factor $\delta \in (0,1)$.[6] We also assume that the retailer employs a trigger strategy such that in every subsequent period after observing low quality, the retailer chooses to not buy.

We first provide intuition for why the retailer may prefer either monetary or relational incentives, and then conduct a more formal analysis. To this end, suppose that the contract between the retailer and supplier is given by $(F^r > 0, B^r = 0)$, and the retailer chooses to buy in every period so long as high quality has always been received in the past; otherwise he doesn't buy. For $(F^r, 0)$ to induce high effort, it must be that $F^r - c_h > u_0$, since otherwise, the retailer cannot punish the supplier for providing low quality. In contrast, the equilibrium contract of the one-shot game, $(F^m, B^m) > 0$, that we solved for in the previous section ensures that $F^m + \alpha B^m - c_h = u_0$; that is, the supplier is pushed down to her outside option. The net cost of switching from monetary to relational incentives is $(F^r - c_h - u_0) - (1-\alpha)B^m$. The first term comes from the fact that the supplier must be paid strictly above her outside option, while the second term comes from the fact that, with $B^r = 0$, there is no supply chain efficiency loss due to a deferred bonus payment. In our formal analysis, we will determine a condition for which this expression is positive or negative.

The general problem of the retailer who wants to induce high effort in an indefinitely repeated environment is to solve the following problem:

$$
\begin{aligned}
&\min_{F,B\geq 0} F + B \qquad s.t.\\
&(IC^r)\ \frac{F+\alpha B - c_h}{1-\delta} \geq F - c_l + \mu\left(\alpha B + \delta\frac{F+\alpha B - c_h}{1-\delta}\right) + (1-\mu)\delta\frac{u_0}{1-\delta}\\
&(IR_s)\ F + \alpha B - c_h \geq u_0\\
&(IR_r)\quad v_h - F - B \geq u_0\\
&(HE)\quad v_h - F - B \geq \mu v_h + (1-\mu)v_l - (c_l + u_0)
\end{aligned}
$$

Notice first that if (IC^r) is satisfied with $B = 0$, then (IR_s) will also be satisfied. With some effort, one can rewrite (IC^r) as:

$$
\begin{aligned}
\delta F + \alpha B &\geq \frac{1}{1-\mu}\left(c_h - c_l + \delta(u_0 + c_l) - \delta\mu(c_h + u_0)\right)\\
&= \frac{c_h - c_l}{1-\mu} + \frac{\delta(c_l - \mu c_h)}{1-\mu} + \delta u_0.
\end{aligned}
$$

[6] Note that δ and α may differ. The latter captures several factors, such as the *difference* in financing rates between the retailer and supplier, or the length of the review period used to determine the product's quality, while the former includes factors such as the *absolute* financing rate – for a single time period – of the supplier and an exogenous probability of dissolution of the relationship, among others. Therefore, it is possible that $\delta \lesseqgtr \alpha$.

Thus the retailer's optimization problem is one of cost-minimization with perfect substitutes.

Figures 2(c)–(d) plot the incentive compatibility and individual rationality constraints for the supplier and divides the parameter space into regions of high effort, low effort and rejection. The retailer's iso-cost (dashed line) curve at the optimal solution is also plotted. Consider panel (c), which corresponds to the case of relatively inefficient monetary incentives ($\delta > \alpha$). While the retailer could still incentivize high effort with conventional monetary incentives, he prefers to set $F > 0$, $B = 0$ and use the threat of punishment to incentivize high effort. This is because the supplier cares sufficiently about the future so that the increased total compensation necessary to encourage high effort is not too large to negate the savings from eliminating the inefficient bonus. On the other hand, in panel (d), which corresponds to the case of relatively efficient monetary incentives ($\delta < \alpha$), the retailer prefers to rely on monetary incentives because relying on relational incentives only would require too large a fixed payment to incentivize high effort. Therefore, we have demonstrated the following:

PROPOSITION 2. *The retailer's preferred subgame perfect equilibrium contract satisfies the following:*

(a) If $\delta > \alpha$, the retailer relies exclusively on relational incentives. That is, $F = \frac{c_h - c_l}{\delta(1-\mu)} + \frac{(c_l - \mu c_h)}{1-\mu} + u_0$ and $B = 0$. In this equilibrium, 100% supply chain efficiency is achieved.[7]

(b) If $\delta < \alpha$, the retailer relies on monetary incentives. That is, the equilibrium contract is identical to the one-shot game. In particular, $F = \max\{0, c_h + u_0 - \alpha B\}$ and $B = \frac{c_h - c_l}{\alpha(1-\mu)} > 0$. In this equilibrium, there is a supply chain efficiency loss of $(1-\alpha)B$.

That is, if $\delta > \alpha$, then the supplier cares sufficiently about the future so that the retailer does not need to raise the supplier's per-period payoff too far above her outside option so as to negate the savings from setting the bonus equal to zero. On the other hand, when $\delta < \alpha$, the supplier is relatively impatient, which makes the retailer's cost of relational incentives ($F^r - c_h > u_0$) too great relative to the efficiency loss due to monetary incentives ($(1-\alpha)B$). Because of this, the retailer prefers to revert to the equilibrium contract of the one-shot game.

REMARK 1. To be sure, there exist other equilibria of the repeated game that can be supported as subgame perfect equilibria. Our analysis has simply provided a characterization of the retailer's *most-preferred equilibrium*, which is a plausible starting point given his first-mover advantage in

[7] Note that, when $\delta > \alpha$, the contract $(F > 0, B = 0)$ backed up by a grim-trigger strategy if low quality is observed is not renegotiation proof. One can easily show that incorporating the possibility for renegotiation makes it more difficult to support an efficient contract with $B = 0$, and the difficulty is increasing in the expected profits that the supplier can extract in the renegotiation. We believe it is unlikely that the ability to renegotiate would substantially influence our experimental results, though this is an interesting avenue for future research.

choosing the contract terms. For example, since (Don't Buy, Low Effort) is always an equilibrium of the post-contracting subgame so long as $\mu(v_h - B) + (1-\mu)v_l - F \leq u_0$, the supplier can credibly demand a higher payoff than in the retailer's preferred equilibrium with the threat that she will reject any less generous contract and choose low effort in any post contracting subgame which does not provide the desired level of compensation. Thus, when the retailer and the supplier are in a long-term relationship, the supplier, potentially, has more bargaining power than when they are in a one-shot relationship. Because of this, even when $\delta < \alpha$, supply chain surplus may be higher in the repeated game than in the one-shot game.

4. Experimental Design

Our experiment consisted of a 2×3 between-subjects design in which we varied, along the first dimension, the scope of retailer and supplier interaction. In one setting, the players were engaged in a one-shot interaction ("one-shot") so that the only incentives the retailer had to induce high effort were direct monetary incentives. In the other setting, the retailer and supplier interacted in an indefinitely repeated game ("indefinite"), such that the retailer could use relational incentives to induce high effort.

In the second dimension of our experiment, we manipulated the monetary incentive. In all three variations the retailer could always offer an unconditional fixed payment, F, to the supplier. However, in two of these variations, we allowed the retailer to also propose a conditional bonus B, where the efficiency of the bonus was set so that a 1 experimental currency unit (ECU) bonus payment by the retailer was worth $\alpha < 1$ ECU to the supplier. In these two conditions, we set $\alpha = 0.7$ or $\alpha = 0.9$, representing a relatively less efficient or more efficient monetary incentive. In the third condition, we did not allow retailers to offer a bonus to the supplier. Therefore, this third condition is a kind of gift exchange game because, in the absence of a conditional bonus, the supplier has no incentive to exert high effort in a one-shot game. We refer to the three conditions as NM (no monetary incentive), LM (low monetary incentive; $\alpha = 0.7$) and HM (high monetary incentive; $\alpha = 0.9$), respectively.

In all six treatments, the value to the retailer of high quality was $v_h = 65$, and the value of low quality was $v_l = 0$. For the supplier, the cost of high effort was $c_h = 25$ and the cost of low effort was $c_l = 13$. High effort would guarantee high quality, whereas low effort would lead to high quality with 50% chance. In all treatments, both players had an outside option of 10.

In each treatment we placed subjects into cohorts of eight to 14. Subjects did not know the size of their cohort nor who was in it. Half of the subjects were given the role of the retailer and the

Table 1 Experimental Design and Number of Participating Subjects

		Monetary Incentive		
		NM (None)	LM (Low) ($\alpha = 0.7$)	HM (High) ($\alpha = 0.9$)
Relationship	One-Shot ($\delta = 0$)	30	40	44
	Indefinite ($\delta = 0.8$)	44	36	46

other half were given the role of the supplier. Roles were fixed for the duration of the experiment. In the one-shot treatments, subjects played the game described above for 30 periods, with random rematching each period. In our indefinite treatments, we followed the implementation introduced by Roth and Murnighan (1978). In particular, subjects were told that they would participate in 10 cycles, with each cycle consisting of an unknown number of periods. At the beginning of each cycle subjects were randomly matched into pairs consisting of one retailer and one supplier. These pairs were maintained for the entire cycle. Subjects were told that at the end of each period, the cycle would continue for another period with probability $\delta = 0.8$ and would terminate with complementary probability. Upon termination of a cycle, subjects would be randomly rematched and a new cycle would begin.[8] A summary of our experimental design and number of participating subjects is displayed in Table 1. There were a total of 240 subjects across all six treatments.

In each period of the one-shot treatments, and the first period of each cycle in the indefinite treatments, the retailer would propose contract terms: either F in the NM conditions, or F and B in the LM/HM conditions. Upon observing the proposed terms, the supplier could accept or reject the terms. If the supplier chose to accept the contract, then the retailer and the supplier would simultaneously choose Buy/Don't Buy or High/Low Effort, respectively. This concluded one period of the one-shot game. In the indefinite game, in every subsequent period of a cycle, if the supplier agreed to the retailer's proposal, then the contract would remain in force and the players would continue to only make a purchase or effort decision in each period until the cycle ended. Conversely, if the supplier chose to reject the retailer's proposal in the first period of the cycle, then both players would receive their outside option for that period and, in the next period, the retailer would propose a new contract. Once a contract was accepted it would remain in force for the remainder of the cycle, though it is possible that a cycle would end without a contract ever being accepted. Thus, this reflects a scenario in which contracting is costly so that once an agreement is reached, it remains in force. An alternative setting is one in which the retailer can make a new

[8] We pre-drew the length of each cycle according to this 80% continuation probability rule ($\delta = 0.8$). All sessions had the same number of cycles and cycle lengths; however, we permuted the order of the cycles across the cohorts. Therefore, all indefinite sessions consisted of 10 cycles (of varying lengths) and the total number of periods was 45.

contract proposal every period. This is interesting but also substantially more complicated because of the expanded strategy space, and is beyond the scope of this paper.

In order to ensure that complexity was not a driver of any results and that the subjects understood the task, we provided them with a decision support tool. Specifically, when making an offer during the contracting stage, retailers had the ability to enter hypothetical contract terms and observe the potential profits for themselves and their supplier. These profits were depicted in two separate 2×2 tables on the screen (each similar to Figure 1), where one table illustrated the retailer's profits, and the other depicted the supplier's profits. In the case of buy and low effort, we displayed the profits for both potential outcomes (whether high quality or low quality was the outcome) so as to not prime subjects with respect to risk. After the retailer made an offer to the supplier, the supplier could observe these same tables for the offered contract, when deciding to accept or reject the contract. Additionally, if the contract was accepted by the supplier, then both players could see the same tables for the accepted contract during their simultaneous decisions (Buy/Don't Buy or High/Low Effort). Table 2 depicts the theoretical predictions for our experiment, based on the assumption of self-interested expected profit maximizers. As noted earlier, for the indefinite treatments, the predictions are stated in terms of the retailer's preferred equilibrium.

Table 2 Experimental Predictions

	One-Shot			Indefinite$^\sharp$		
	NM	LM	HM	NM	LM	HM
% Supply Chain Efficiency†	50.00	74.29	93.33	100	100	93.33
% High Effort	0	100	100	100	100	100
Supplier Profit	10.00	10.00	10.00	16.00	16.00	10.00
Retailer Profit	10.00	19.71	27.33	24.00	24.00	27.33
Fixed Payment (F)	0.00	11.00	11.00	41.00	41.00	11.00
Effective Bonus (αB)*	0.00	24.00	24.00	0.00	0.00	24.00

$^\sharp$ We report the retailer's preferred equilibrium. As noted in Remark 1, there are other equilibria in which the supplier extracts a larger share of the surplus.
† We measure supply chain efficiency as the ratio $\frac{\pi^r+\pi^s}{v_h-c_h}$, where π^i are the expected profits of firm i. When high effort is an equilibrium, this expression becomes $\frac{v_h-c_h-(1-\alpha)B}{v_h-c_h}$.
* We measure the *effective* bonus payment as αB.

After completing all periods of the game, subjects also participated in a risk elicitation task, similar to that in Holt and Laury (2002), where they decided between a series of 50-50 binary lotteries. One lottery was fixed from question-to-question and was relatively safe, whereas the other lottery was more risky, and the payoff if the good state occurred was increasing from question-to-question. Subjects were paid for one randomly selected lottery question.

All sessions were conducted at a university located in the northeast United States, where participants were students, mostly undergraduates. Cash was the only incentive offered, with average earnings of \$30, which was based on the profits from all rounds of the main experiment, profits from the risk aversion exercise, and a \$5 show-up fee. Sessions took approximately 80 minutes, and were implemented using zTree (Fischbacher 2007).

5. Results

Here we present a top-down summary of our experimental findings, beginning with outcome statistics, then proceeding to retailer contracts and supplier decisions. In the treatments, there was some evidence of experience effects, however, they do not affect any of the primary results. Therefore, our analysis includes all data, and we provide additional tables, which only include the final third of decisions, in Appendix A. For presentation purposes, we consider our main experimental results here, and relegate additional results to Appendix B. Some of these primary results include (1) relational incentives increase supply chain efficiency and quality, regardless of whether there is a monetary incentive, (2) in a one-shot setting, introducing a low efficiency monetary incentive actually reduces supply chain efficiency, contrary to theory, (3) also counter to theory, when relational incentives are present, adding a low efficiency monetary incentive decreases supply chain efficiency, and essentially crowds out any benefit from relational incentives, whereas adding a high efficiency monetary incentive actually complements relational incentives and leads to even more favorable performance, and (4) retailers appear to make contract offers that deviate from the normative benchmarks, and suppliers tend to make decisions that lead to more equitable divisions of profit.

5.1. Outcomes

In Figure 3 we plot the average expected profit earned by retailers and suppliers, and supply chain efficiency, based on decisions from all six treatments. The dark shaded rectangle (and the number inside) gives the supplier's average expected profit while the light-shaded rectangle (and number inside) gives the retailer's average expected profit. The number on top of the bars represent the average supply chain efficiency for the treatment.

As can be seen in Figure 3, long-term indefinitely repeated interactions always lead to higher supply chain efficiency than one-shot interactions (relative increases of 17.6% in NM, 6.6% in LM, and 8.5% in HM). Taking the session average as the unit of observation, a simple ANOVA model with factors for the scope of interaction (one-shot vs indefinite) and the type of monetary incentive (NM, LM and HM), demonstrates that supply chain efficiency is indeed significantly higher for the indefinite treatments ($p = 0.019$). In terms of each party's profits, retailers benefit from long-term indefinite interactions, with large increases in profits in the NM and LM treatments (7.81

Figure 3 **Expected Supplier Profits, Retailer Profits, and Supply Chain Efficiency in All Treatments.**

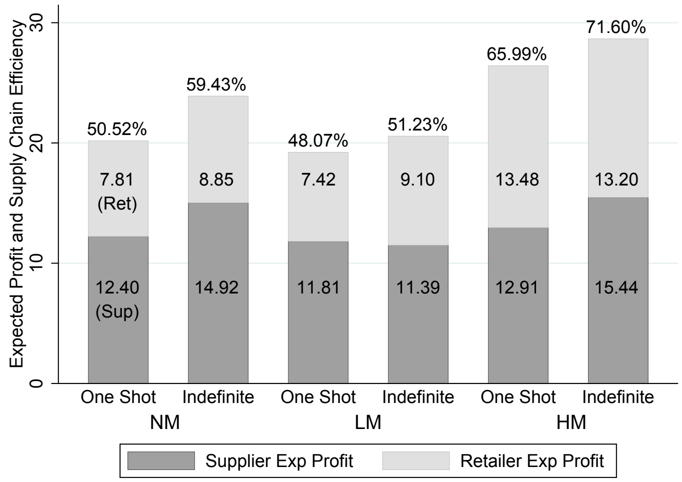

versus 8.85 in NM, and 7.42 versus 9.10 in LM) and a negligible decrease in the HM treatment (13.48 versus 13.20), although this decrease is not significant ($p = 0.817$). Turning to suppliers, they also benefit from indefinite interactions, with large increases in profits in the NM and HM treatments (12.40 versus 14.92 in NM, and 12.91 versus 15.44 in HM), and a negligible decrease in the LM treatment (11.81 versus 11.39), but again this decrease is not significant ($p = 0.776$). Thus, it appears that relational incentives lead to a weak Pareto improvement in average profits.

The same ANOVA also shows that the HM treatments achieve significantly higher supply chain efficiency than either the NM or the LM treatments (in both cases, $p \ll 0.01$). Somewhat surprisingly, and contrary to the theoretical predictions, the one-shot LM treatment actually achieves the lowest efficiency of all treatments – roughly 5% lower than even the NM treatment (48.07% versus 50.52%, $p = 0.065$). This suggests that, contrary to theory, having no monetary incentive may actually be better than having a less efficient monetary incentive.

Another interesting result from Figure 3 is that adding a monetary incentive to a setting which already has relational incentives, yields a non-monotonic effect on supply chain efficiency. Specifically, when comparing the indefinite NM treatment to the indefinite LM treatment, the latter actually yields a lower supply chain efficiency, by 13.8% (59.43% versus 51.23%). On the other hand the indefinite HM treatment yields a significantly better supply chain efficiency than any other treatment, and is 20.5% higher than the indefinite NM treatment (59.43% versus 71.60%). This result parallels some of the results on intrinsic and extrinsic motivation, which often find that the presence of a small extrinsic incentive may crowd out any intrinsic motivation, but a large extrinsic incentive may complement it (Gneezy and Rustichini 2000, Bowles 2008). Our data may be analogous in that the benefits of relational incentives may be crowded out by a less efficient monetary incentive (indefinite NM versus indefinite LM), but the presence of an efficient monetary incentive appears to complement relational incentives (indefinite HM treatment).

To see where the differences in supply chain efficiency come from, Table 3 reports the average frequency of each of the possible outcomes – outside option (either because the supplier rejected or the retailer chose not to buy), the retailer chose buy and the supplier exerted low effort, and the retailer chose to buy and the supplier exerted high effort. Note that quality is directly tied to the effort exerted by the supplier, so we omit it for brevity.

Table 3 Frequency of Outcomes in All Treatments (%)

			Monetary Incentive		
			NM	LM	HM
Relationship	One-Shot	Outside Option	59.33	38.83	23.64
		Buy, Low Effort	38.67	47.50	32.42
		Buy, High Effort	2.00	13.67	43.94
	Indefinite	Outside Option	46.57	44.07	18.84
		Buy, Low Effort	33.74	31.98	26.38
		Buy, High Effort	19.70	23.95	54.78

It is immediately apparent that the indefinite treatments induce substantially more high effort, and thus yield higher quality, across all types of monetary incentives (19.70% versus 2.00% in NM, 23.95% versus 13.67% in LM, and 54.78% versus 43.94% in HM). Furthermore, in the NM and HM treatments the frequency of both the outside option and low effort, are lower in the indefinite treatments than in the one-shot treatments. Comparing the NM and LM treatments, in both one-shot and indefinite settings, we observe a higher frequency of high effort in the LM treatment, yet Figure 3 showed that supply chain efficiency is actually lower in the LM treatments. This is due to two reasons. First, under one-shot interactions, low effort is more frequent in the LM treatment (47.50%), which is the worst outcome from a supply chain perspective. Second, as we will see, in the indefinite LM treatments, retailers chose to rely on positive (less efficient) bonuses, rather than relational incentives, thus reducing supply chain efficiency. Lastly, note that for both the one-shot and indefinite treatments, the HM condition achieves the highest frequency of high effort, 43.94% and 54.78%, contributing to higher quality and supply chain efficiency.

5.2. Retailer Contracts

Table 4 reports the average contract terms proposed by retailers. First, the average contract in the indefinite treatments provides significantly higher total compensation than in the one-shot treatments (35.87 versus 23.02 in NM, 34.27 versus 33.60 in LM, and 40.73 versus 37.88 in HM, ANOVA using session averages, $p = 0.004$). This suggests that suppliers may, in fact, use their enhanced bargaining power in the indefinite setting (by being able to reject contracts and receive a new offer, next period, by the same retailer) to extract a larger share of the surplus.[9] In addition, although the theory predicts no difference in total compensation under one-shot interactions for the LM and HM treatments (35.00 in both cases), total compensation is significantly higher in the HM treatment (37.88 versus 33.60, ANOVA, $p = 0.010$). Moreover, both of these numbers differ from the theoretical prediction (Wilcoxon signed rank test, $p = 0.068$).[10]

Table 4 Summary of Contract Terms in All Treatments

			Monetary Incentive		
			NM	LM	HM
Relationship	One-Shot	Total Compensation	23.02	33.60	37.88
		Fixed Payment (F)	23.02	15.66	7.84
		Effective Bonus (αB)	0.00	17.94	30.04
	Indefinite	Total Compensation	35.87	34.27	40.73
		Fixed Payment (F)	35.87	12.74	14.24
		Effective Bonus (αB)	0.00	21.53	26.49

Next, in Table 4, consider the average values for the fixed payment and effective bonus in the one-shot LM and HM treatments. Again, while theory predicts no difference between the parameters across these two treatments, in the one-shot LM treatment, we see that the fixed payment is higher ($F = 15.66$), and the effective bonus is lower ($\alpha B = 17.94$), than the theoretical predictions of $F = 11.00$ and $\alpha B = 24.00$. On the other hand, the situation is reversed in the one-shot HM treatment – retailers set the fixed payment too low ($F = 7.84$) and rely too much on the effective bonus ($\alpha B = 30.04$). One possibility is that retailers recognize the inefficiency of the bonus in the LM treatment and so offer higher fixed payments, hoping that the suppliers will be intrinsically motivated to exert high effort. However, as the supply chain efficiency results in the previous subsection illustrated, and as we will show when analyzing suppliers, they do not appear to be intrinsically motivated. Instead, in Section 6 we will show how a model of fairness preferences may be able to explain this discrepancy in retailer offers.

Lastly, in Table 4, focusing on the HM treatment, and comparing the average contract parameters for one-shot and indefinite interactions, we see that while total compensation increases by almost 3 ECU (37.88 versus 40.73; rank sum test, $p = 0.021$), the conditional bonus actually decreases and

[9] Indeed, if we compare the total compensation from an initial rejected offer with the retailer's next proposal, in all incentive conditions, compensation increases.

[10] With four independent sessions, this is the strongest possible rejection using session averages as the unit of observation.

the average fixed payment increases.[11] By relying more on the fixed payment (relational incentive) and less on the bonus (monetary incentive), in the indefinite HM treatment, it appears that retailers recognize that they can use relational incentives to reduce the cost of inducing high effort. This contributes to the favorable supply chain efficiency results observed earlier in the indefinite HM treatment.

While useful from an aggregate standpoint, Table 4 masks a great deal of heterogeneity in contract terms proposed by retailers. To get a sense of this, for the LM and HM treatments, Figure 4(a) plots the subject average fixed payment and effective bonus.[12] The first thing to notice is that there is a similar, negative relationship between the fixed payment and effective bonus in all treatments. Although, there are vertical shifts in the indefinite treatments due to the higher total compensation offered. It is important to note that those subjects who offer large effective bonuses and small fixed payments not only provide stronger incentives for high effort, but are also more generous to suppliers in that total compensation is higher (i.e., the slopes are greater than -1).

In Figure 4(b), we plot the relationship between the average effective bonus offered and the average quality received. Recognizing that higher effective bonuses mean stronger incentives and, given the relationship between F and αB, higher total compensation, retailers should receive high quality more frequently.[13] As can be seen, in all treatments there is a similar, positive relationship between the average effective bonus and quality. Therefore, the reason why retailers in the one-shot LM treatment do not receive high quality frequently can be explained, in part, by the weak incentives and low total compensation. In the HM treatment, we also see that, for the same effective bonus, retailers receive higher quality under indefinite interactions than one-shot interactions.

Finally, in Figure 4(c), we plot the relationship between the average effective bonuses and the retailers' average profit. Here we see a stark difference between the LM and HM treatments, for both one-shot and indefinite interactions. In the LM treatments, although the higher effective bonus leads to higher quality, retailer profits do not increase. This is largely due to the bonus being relatively inefficient in the LM treatment, where increasing the effective bonus by 1 ECU costs the retailer $1/0.7 \approx 1.429$ ECU. On the other hand, in the HM treatments, the relationship between

[11] A rank-sum test comparing the effective bonus in the HM treatment across one-shot and indefinite interactions does not reject that they are equal ($p = 0.149$); however, the same test on the fixed payment does reject that they are equal ($p = 0.043$).

[12] While there is heterogeneity in the NM treatments, it is less interesting because retailers have only one contract parameter to set. In the interest of space, we omit this analysis.

[13] Unfortunately, the fact that retailers who offer higher bonuses also offer higher total compensation makes it difficult to distinguish between whether it is the incentives alone or something else that drives the increased effort provision. We will return to this in the next section.

Figure 4 **Heterogeneity in Contract Terms, Quality Received and Retailer Profits in the LM and HM Treatments**

(a) Fixed and Effective Bonus Payments

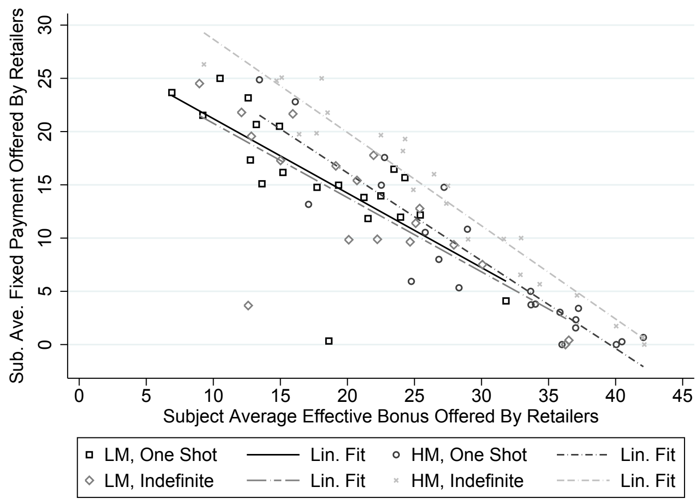

(b) Quality Received and Effective Bonus Payments

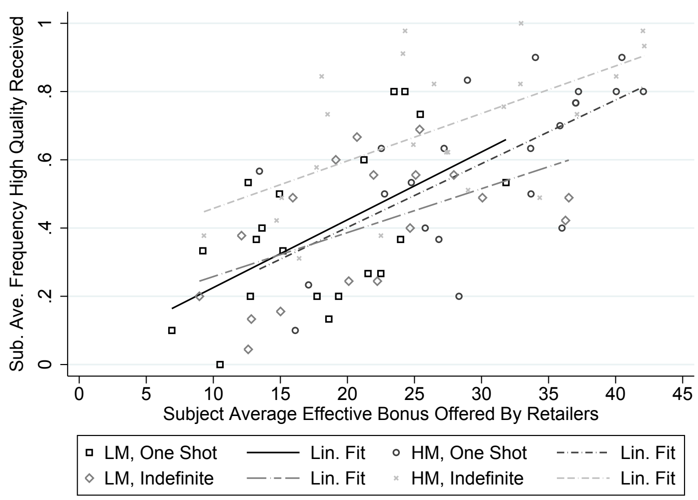

(c) Retailer Profit and Effective Bonus Payments

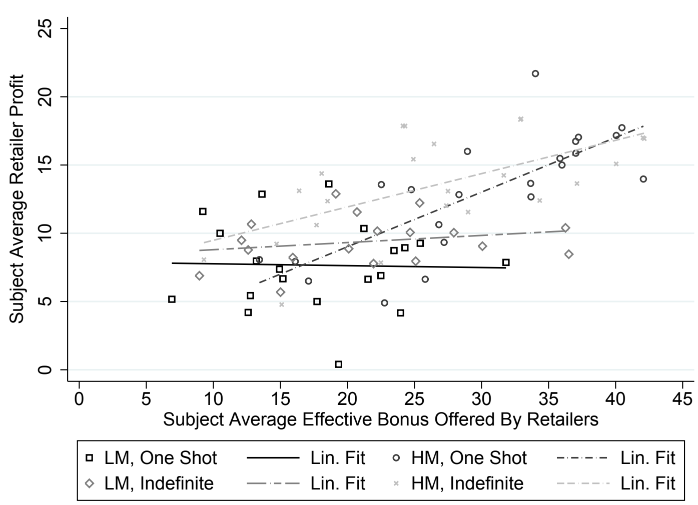

effective bonuses and retailers' profit is positive. As Figure 4(b) showed, the quality increase is similar, but now a 1 unit increase in the effective bonus only costs retailers $1/0.9 \approx 1.11$ ECU.

5.3. Supplier Decisions

We now turn to supplier decisions, who face a two-part decision – whether to accept or reject the contract and, conditional on accepting the contract, whether to exert high or low effort. In the interest of space, we do not report regressions for the decision to accept but we note that both the fixed payment and the effective bonus are (approximately) equally important. Specifically, an extra unit of the effective bonus increases the chance of a supplier accepting the contract by between 3.9% and 4.8%, while an extra unit increase in the fixed payment increases the chance of accepting by between 3.2% and 5.9%.

The more interesting supplier decision is to exert high effort upon accepting the contract. We report these regression results in Table 5. In the one-shot LM and HM treatments, suppliers only increase their effort provision when the effective bonus increases (marginal effect of 0.032, $p < 0.01$) but are unresponsive to increases in the fixed payment (marginal effect of 0.011, $p > 0.10$).[14] That is, there does not appear to be any underlying gift exchange reciprocity motive. In the NM treatment, at first glance, there appears to be some evidence for reciprocity: higher fixed payments lead to higher effort provision, but the size of the effect is virtually zero (marginal effect of 0.004, $p < 0.01$).

Table 5 Supplier's Decision to Exert High Effort (Random Effects Logit Regression, Marginal Effects)

	One-Shot				Indefinite			
Parameter	LM/HM		NM		LM/HM		NM	
Fixed Payment (F)	0.011	(0.007)	0.004^{***}	(0.001)	0.039^{***}	(0.012)	0.058^{***}	(0.006)
Effective Bonus (αB)	0.032^{***}	(0.008)			0.059^{***}	(0.009)		
Risk Attitude	-0.018	(0.014)	-0.004^{***}	(0.001)	0.000	(0.011)	0.043^{***}	(0.015)
Period	-0.001	(0.001)	-0.003^{*}	(0.002)				
Cycle					0.004	(0.007)	-0.005	(0.003)
Cycle Period					-0.005	(0.004)	-0.036	(0.040)
Observations	973		353		1574		183	

Note 1: *, ** and *** denote significance at the 10, 5 and 1% levels. The table reports the marginal effects from random effects logit models with standard errors in parentheses which are clustered at the session level.
Note 2: Risk attitudes are measured such that higher numbers indicate a greater willingness to take risk.

Proceeding with Table 5, in all of the indefinite treatments, both the fixed payment and the effective bonus have a positive effect on suppliers' willingness to exert high effort. This is further

[14] In the interest of parsimony, we pooled the LM and HM treatments. In regressions not shown, we included treatment interactions for both F and αB. Only for the indefinite treatment did we find that subjects respond to F significantly less in the LM treatment, which is consistent with Figure 4(b). In all other cases, there was no significant difference in F or αB for the LM and HM treatments.

evidence that relational incentives are effective at inducing high effort, and thus high quality, from suppliers: the suppliers anticipate that if low quality is observed (because they chose low effort), then they will lose out on the high fixed payment in subsequent periods of the cycle. Note, however, that in the LM/HM treatments, we can reject that the coefficients on F and αB are the same ($p \ll 0.01$). Thus bonuses are more effective at inducing high effort, which could explain why retailers still rely on them even when relational incentives are available.

While Table 5 and our discussion of acceptance decisions speaks to the general response of contract parameters to accept/reject or exert high/low effort, it does not speak to how suppliers behave when the contract is in one of the three specific ranges (i.e., reject, accept/high effort or accept/low effort). Table 6 focuses on the most interesting region – namely, when the contract was such that high effort was optimal – and reports the frequency with which suppliers take each possible action. First, across all treatments, suppliers reject a non-negligible fraction of contracts where high effort is optimal, and the fraction is substantially higher in the indefinite treatments than the one-shot treatments. This could be due to fairness concerns because, while not depicted, the total compensation offered to the supplier is significantly lower for offers that are rejected than accepted (Wilcoxon test, $p \ll 0.01$). Moreover, the fact that the frequency of rejection (conditional on an offer being made) is so high in the indefinite treatments (e.g. 45.66% in indefinite LM) may reflect suppliers using their extra bargaining power (i.e., the ability to reject) to extract more from retailers.

Table 6 Frequency of Suppliers' Actions When High Effort was Optimal in All Treatments (%)

(a) One-Shot

	Monetary Incentive		
	NM	LM	HM
Reject		11.29	15.08
Low Effort	n/a	45.97	24.83
High Effort		42.74	60.09
Freq. High Effort Opt.	0.00	20.67	68.33

(b) Indefinite[1]

	Monetary Incentive		
	NM	LM	HM
Reject[2]	4.11	45.66	23.14
Low Effort[3]	33.12	50.47	31.66
High Effort[3]	66.88	49.53	68.34
Freq. High Effort Opt.	31.72	61.85	90.53

Note 1: In the Indefinite treatments, we use (IC^r) as our definition for incentive compatibility, for relational incentives.
Note 2: After a contract was accepted, it would remain in force until the end of the cycle. Therefore, we must look at rejections in periods in which an offer was made, which is what the table reports.
Note 3: The numbers for "Low Effort" and "High Effort" are conditional on the offer being accepted. As such, a direct comparison between one-shot and indefinite treatments is difficult (which is not our primary purpose).

We also see in Table 6 that suppliers frequently exert low effort, despite high effort being optimal. This is especially the case in the LM treatment where it is the most frequent supplier decision (45.97% in one-shot LM and 50.47% in indefinite LM). Such choices could also be driven by fairness concerns. For instance, let $\Delta_h = (v_h - F - B) - (F + \alpha B - c_h)$ denote the payoff difference between the retailer and supplier conditional on high effort. Similarly, let $\Delta_l = (.5v_h - F - .5B) - (F + .5\alpha B - c_l)$ denote the payoff difference conditional on low effort. Then, if $\Delta_l < \Delta_h$, it means that choosing low effort tilts the payoffs to be relatively more equitable than if the supplier chose high effort. Indeed, given the parameters of the experiment, $\Delta_l < \Delta_h$ when $B < 89/(1+\alpha)$, which was true for nearly all contract proposals in our experiment.[15]

6. Behavioral Explanation: Fairness Concerns

Thus far we have observed considerable differences in our data and those predicted by self-interested, expected payoff maximization. For example, in the one-shot treatments, retailers set F too high and αB too low in the LM treatment, and conversely, set F too low and αB to high in the HM treatment. Here, we show that social preferences – in particular, inequality aversion (c.f., Fehr and Schmidt (1999), Bolton and Ockenfels (2000)) in which a player suffers a utility loss when her payoff deviates from her opponent – can capture a number of features seen in our data. Indeed, we have already seen fairness may be at play for suppliers: when high effort is optimal, suppliers frequently reject offers or choose to exert low effort, where low effort leads to smaller differences in payoffs between the players (and rejections equalize payoffs).

To formalize this, we assume that the supplier is inequality averse; in particular, with an aversion to disadvantageous inequality, which we parameterize by $\lambda^d \geq 0$ and a coefficient of advantageous inequality aversion of $\lambda^a \geq 0$.[16] We assume that the supplier takes an *ex ante* approach to inequality aversion (see Saito (2013) and the references therein for more details on fairness preferences in risky environments, including evidence for ex ante notions of fairness). We assume that the retailer does not suffer from inequality aversion. In this case, the retailer's problem is:

$$\begin{aligned}
&\min_{F,B\geq 0} F + B \qquad s.t.\\
(IC)\quad & F + \alpha B - c_h - \lambda^d[\Delta_h]^+ + \lambda^a[\Delta_h]^- \geq F + \mu\alpha B - c_l - \lambda^d[\Delta_l]^+ + \lambda^a[\Delta_l]^-\\
(IR_s)\quad & F + \alpha B - c_h - \lambda^d[\Delta_h]^+ + \lambda^a[\Delta_h]^- \geq u_0\\
(IR_r)\quad & v_h - F - B \geq u_0\\
(HE)\quad & v_h - F - B \geq \mu v_h + (1-\mu)v_l - (c_l + u_0)
\end{aligned}$$

where $[X]^+ := \max\{0, X\}$, $[X]^- := \min\{0, X\}$, $\Delta_h = v_h + c_h - 2F - (1+\alpha)B$ and $\Delta_l = \mu v_h + c_l - 2F - \mu(1+\alpha)B$.

[15] An alternative explanation might be risk seeking behavior by suppliers. If a supplier is risk seeking, he may prefer to choose low effort (since it is risky) even though a risk neutral or risk averse supplier would exert high effort. However, as seen in Table 5, there is no consistent evidence that subjects who are more willing to take risk are less likely to exert high effort.

[16] The literature often assumes that $\lambda^a \leq \lambda^d$; disadvantageous inequality is worse than advantageous inequality.

Working with (IC), we can rewrite it as:

$$B - \frac{\lambda^d}{\alpha(1-\mu)}\left([\Delta_h]^+ - [\Delta_l]^+\right) + \frac{\lambda^a}{\alpha(1-\mu)}\left([\Delta_h]^- - [\Delta_l]^-\right) \geq \frac{c_h - c_l}{\alpha(1-\mu)}.$$

If $\lambda^a = \lambda^d = 0$, so that the supplier is purely self-interested, then we have the familiar incentive compatibility constraint in Section 3.1: $B^* \geq (c_h - c_l)/(\alpha(1-\mu))$, and we can find F^* via the individual rationality constraint. However, when $(\lambda^d, \lambda^a) > 0$, both the IC and IR constraints change, necessitating changes in F and B.

In Figure 5(a)-(b) we plot regions of the parameter space where F and B take different values relative to the case of no fairness preferences, (F^*, B^*) for the LM and HM treatments. The region **B** covers the area relatively close to the origin, yet has $F \leq F^*$ and $B > B^*$. This is because, when fairness preferences are not too strong, it is actually in the retailer's interest to reduce inequality by raising incentives – B (and also lowering F). Eventually, however, because of the inefficiency of the bonus, this becomes too costly for the retailer. Therefore, she will remove payoff inequality by raising $F > F^*$ to equalize expected payoffs and reduce B to make sure that incentives are aligned. Interestingly, the effect of advantageous inequality aversion is to weaken the role of the bonus in inducing high effort. This is because, by exerting low effort, it becomes more likely that the supplier is in an advantageous position relative to the retailer, which she also dislikes. Hence, it may be that $B < B^*$. This is the region depicted by **A** in Figure 5(a)-(b), where $F > F^*$ and $B \leq B^*$. There is also a small region, **C**, where $(F, B) \gg (F^*, B^*)$. This happens when advantageous inequality aversion is relatively large and disadvantageous inequality aversion is relatively small.

Panel (c) considers the NM treatment. In Appendix C we show that the supplier can be induced to exert high effort if and only if she cares sufficiently about *advantageous* inequality aversion. For the parameters of our experiment, this corresponds to $\lambda^a > 24/89 \approx 0.2697$ and this is depicted as region **D**. On the other hand, in region **E** the retailer cannot induce high effort and so prefers to take his outside option.

To summarize, in the LM one-shot treatment, our experimental results illustrated that retailers set F too high and αB too low, relative to the case of no fairness preferences, which corresponds to the large region **A** in Figure 5(a). For the one-shot HM treatment, we observed that F was too low and αB was too high, relative to the case without fairness preferences, which corresponds to region **B** in Figure 5(b). Finally, for the one-shot NM treatment, retailers set F very low and almost never induced high effort, which corresponds to region **E** in Figure 5(c). As can be seen, there is a substantial region of the (λ^a, λ^d) parameter space in the three plots for which these types of contract offers can be satisfied. For instance, if both fairness parameters are, say 0.24, in both

Figure 5 The Influence of Fairness on Fixed and Effective Bonus Payments For Our Experimental Parameters

(a) One-Shot LM Treatment

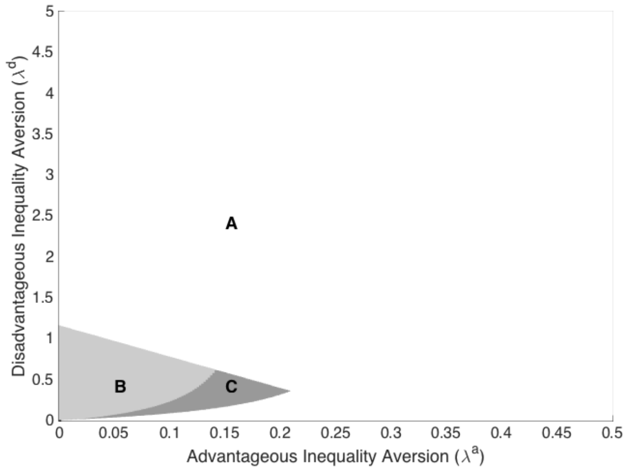

(b) One-Shot HM Treatment

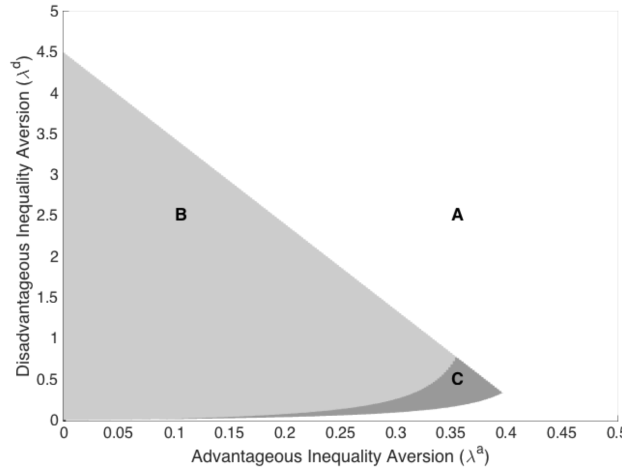

Legend

A: $F > F^*$ & $B \leq B^*$
B: $F \leq F^*$ & $B > B^*$
C: $F > F^*$ & $B > B^*$
D: Can induce high effort with F
E: Can't induce high effort with F

Note: F^* and B^* represent the optimal fixed payment and bonus for the one-shot game for a supplier without inequality aversion.

(c) One-Shot NM Treatment

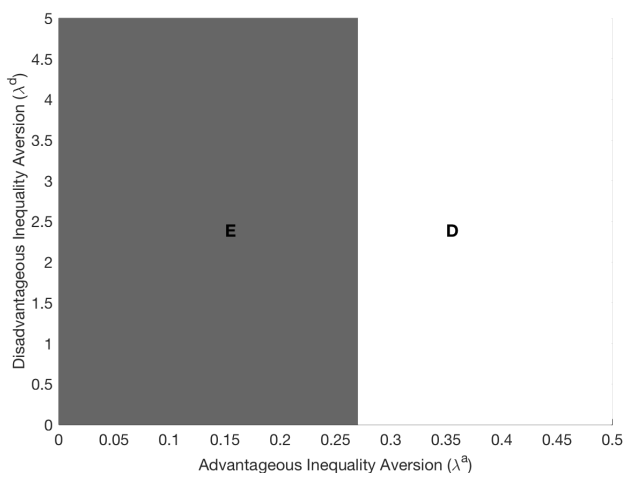

LM and HM scenarios, this leads to predicted contract parameters of $F = 14.25$ and $\alpha B = 22.43$ in LM, and $F = 10.15$ and $\alpha B = 27.60$, while making high effort impossible to implement in NM, which are qualitatively consistent with our data. Indeed, any pair (λ^a, λ^d) in the intersection of region **A** from panel (a), region **B** from panel (b) and region **E** from panel (c) will generate this pattern for F and αB that we observe.

6.1. Estimation

To estimate supplier decisions, we assume that suppliers are prone to errors, and apply the familiar logit structure. Specifically, conditional on an offer, (F, B), and fairness preference parameters, (λ^a, λ^d), let $u_j(F, B, \lambda^a, \lambda^d)$ for $j \in \{h, l, r\}$ denote the utility from exerting high effort (h), low effort (l) or rejecting the offer (r). Also let $u_a(F,B,\lambda^a,\lambda^d) := \max\{u_h(F,B,\lambda^a,\lambda^d), u_l(F,B,\lambda^a,\lambda^d)\}$ denote the utility from accepting the offer and then following through with the optimal action. The probability that the supplier accepts the offer is then:

$$\Pr(\text{Accept}|F,B,\lambda^a,\lambda^d) = \frac{1}{1+e^{\gamma_a(u_r(F,B,\lambda^a,\lambda^d)-u_a(F,B,\lambda^a,\lambda^d))}},$$

where γ_a is the rationality parameter associated with the accept/reject decision.

Conditional on acceptance, we assume that the probability of choosing high effort is given by:

$$\Pr(\text{High}|\text{Accept},F,B,\lambda^a,\lambda^d) = \frac{1}{1+e^{\gamma_e(u_l(F,B,\lambda^a,\lambda^d)-u_h(F,B,\lambda^a,\lambda^d))}},$$

where γ_e is the rationality parameter associated with the effort decision.[17] Given this formulation it is straightforward to derive the likelihood function, which can then be maximized to obtain maximum likelihood estimates of the rationality parameters, γ_a and γ_e, as well as our parameters of interest – fairness preferences: (λ^a,λ^d). We estimate the model separately for each subject and we impose the common restriction that $\lambda^a \leq \lambda^d$.

In Table 7 we summarize the results of our estimation in terms of the fraction of suppliers for whom we can reject self-interested expected payoff maximization in favor of one which incorporates inequality aversion. Overall, for 47.5% of subjects we can reject the joint hypothesis that $\lambda^a = \lambda^d = 0$, and across treatments, the fraction varies from 31.8% to 68.2% of suppliers. This fraction is consistent with other principal-agent studies, such as Fehr and Schmidt (2000), who showed that approximately 40% of subjects in their experiments showed signs of inequality aversion. Beyond this, about 34.2% of suppliers appear to be governed by advantageous inequality aversion, as well as disadvantageous inequality aversion. These findings corroborate what some of our earlier descriptive analysis showed – that at least some suppliers have concerns for fairness. In the next section, we further investigate the role that inequality aversion has on behavior in our environment by automating the role of suppliers and, therefore, eliminating the need for retailers to respond to the fairness concerns of suppliers.

[17] This formulation assumes that the supplier is myopic: When making the decision to accept or reject the offer, she compares the utility from rejecting with the utility from subsequently choosing the best action if she accepts. However, once the offer is accepted, she is prone to making decision errors in choosing between high and low effort.

Table 7 **Summary of Estimation Results Incorporating Fairness Concerns (Subject-Level Estimation)**

	Fraction Rejecting $\lambda^a = \lambda^d = 0$		Fraction Rejecting $\lambda^a = 0$	
Treatment	5% Level	10% Level	5% Level	10% Level
One-Shot NM	33.3%	40.0%	26.7%	40.0%
One-Shot LM	40.0%	50.0%	20.0%	30.0%
One-Shot HM	31.8%	31.8%	4.6%	4.6%
Indefinite NM	68.2%	72.7%	63.6%	63.6%
Indefinite LM	38.9%	44.4%	27.8%	27.8%
Indefinite HM	65.2%	69.6%	56.5%	56.5%
Overall	47.5%	52.5%	34.2%	37.5%

Note 1: In all cases, we imposed the restriction that $0 \leq \lambda^a \leq \lambda^d$.
Note 2: The tests are based on likelihood ratio tests where we impose the given restriction. Significance is according to the $p-$value from a χ^2-test with degrees of freedom equal to the number of restrictions.

7. Further Exploration of Fairness Concerns

7.1. A Treatment With Automated Suppliers

As a further robustness check for whether fairness preferences may account for our results, we conducted two additional experimental treatments in which the supplier side of the interaction was automated. These automated suppliers were programmed in a way to maximize their own expected profits, and thus, did not have any inequality aversion. The experimental protocols for the two treatments were identical to our original experiment, other than automating the suppliers, which retailers were aware of. We ran two treatments: one-shot LM and one-shot HM, with 19 and 15 subjects, respectively. To the extent that retailer decisions are closer to the normative benchmark theoretical predictions, it is suggestive that inequality aversion by human suppliers may be the cause for the outcomes we observed in our main treatments.

Figure 6 provides a snapshot of the observed retailer contract terms across all periods of the automated supplier treatments. As one can see, initially, there are some experience effects for both parameters in each treatment, but by the second half of decisions, in both LM and HM, the values of the fixed payment and effective bonus converge quite close to the standard theoretical benchmarks of $F = 11$ and $\alpha B = 24$. A two-sided $t-$test comparing these values with the equilibrium predictions yields only one weakly significant difference ($p = 0.303$ for F and $p = 0.081$ for αB in LM, and $p = 0.305$ for F and $p = 0.508$ in HM), with each subject constituting an independent observation. These levels of significance drop further if one omits the first half of decisions, when experience effects were quite salient ($p = 0.714$ for F and $p = 0.335$ for αB in LM, and $p = 0.553$ for F and $p = 0.641$ in HM). Lastly, it should come to no surprise that these retailer offers differ substantially from our original treatments with human suppliers, and that the corresponding supply chain efficiency results in the automated supplier treatments are more than 17% higher than with human suppliers.

Figure 6 Retailer Contracts in the Automated Supplier Treatments

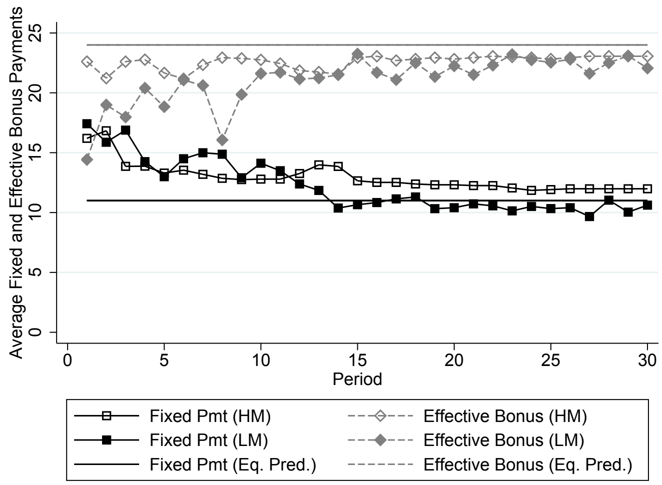

One other interesting result from these automated supplier treatments concerns the variance of decisions. With human suppliers we reported a great deal of heterogeneity in contract parameters (recall Figure 4). In contrast, in the automated supplier treatments, there appears to be less heterogeneity – both within and between subjects – in the proposed contract parameters. Relatedly, if we account for experience effects and omit the first half of decisions, only two retailers (out of 42) could be viewed as making the equilibrium contract offer in the human supplier treatments, but in the automated supplier treatments 15 (out of 34) are essentially making the equilibrium contract offer.[18] Thus, subjects are more quick to settle upon a particular set of contract parameters and the chosen parameters are substantially closer to the equilibrium when suppliers are automated to maximize their expected profit, obviating the need for retailers to account for any concern for fairness preferences.

7.2. Discussion

We would be remiss if we did not admit that other explanations may exist for the data in our original experiment. For instance, one may note that while decisions in the automated supplier treatments are close to the standard theoretical predictions, they are not perfectly in line with them. Therefore, it is possible that some other potential behavioral issues, such as bounded rationality or satisficing for retailers, may be at play. We should also note that in addition to removing fairness preferences, the automated treatments also eliminated the possibility of supplier errors, which should facilitate learning by retailers (although, the supplier's decision is not overly complicated, considering the

[18] We base this on the Euclidean distance of the subject average contract parameters from the theoretical prediction and say that a subject is close to the equilibrium if the distance is less than 3.

decision support). Nevertheless, we believe that our analysis of fairness preferences in the previous section, and the automated supplier treatments of this section, suggest that inequality aversion is useful in organizing the data and the corresponding results.

One may also wonder whether retailers – through suboptimal contract choices – are partially responsible for the patterns in our original experimental data. While we cannot completely eliminate this explanation, it is unlikely to be a primary explanation, as retailers generally behaved in a rational way, given the preferences of the suppliers, in each treatment. For example, take the one-shot NM treatment. Since it is possible for retailers to induce high effort with generous fixed payments if suppliers have sufficient advantageous inequality aversion, one may wonder why retailers did not simply offer more favorable fixed payments. However, our earlier results show that only approximately 27% of suppliers display such fairness preferences. Therefore, when faced with this distribution of supplier types, it is actually optimal for a retailer to take the outside option rather than offer a generous fixed payment, which is consistent with our data. Conversely, in those treatments where a large portion of suppliers exhibited fairness concerns, retailers appear to deviate from the normative theory and instead make offers which correctly take these preferences into account.

8. Concluding Remarks

Retailers and manufacturers are always searching for effective ways to ensure high quality products from suppliers. This is due to the fact that they are essentially the "face" of the product, and are often held responsible for any quality issues, even when those quality issues stem from suppliers. Two tools that are available for retailers and manufacturers to improve quality from suppliers are monetary incentives, such as a bonus that is awarded when high quality is received, and relational incentives, such as when the retailer (or manufacturer) can threaten to no longer buy from the supplier in the future. In this study we investigate the efficacy of monetary and relational incentives for managing high quality products from a supplier, and their resulting impact on supply chain efficiency.

While theoretically favorable at inducing high effort and high quality products by suppliers, one drawback of monetary incentives in a one-shot context is that they are inefficient from a supply chain perspective. We theoretically showed that under certain conditions, it is possible to avoid this supply chain inefficiency if the two parties engage in an long-term interaction, where relational incentives exist.

We then conduct a controlled human-subject experiment, between a retailer and supplier, that tests these equilibrium predictions, and find a number of interesting insights. Of those insights, a primary one is that, whether or not a monetary incentive is present, relational incentives increase quality and supply chain efficiency, and are weakly Pareto improving. This is consistent with a number of well-known examples in practice, which have highlighted the mutual benefits of close-knit, long-term relationships, such as Chrysler (particularly in the 1990's), Toyota, Visteon, Dell, John Deere, and Motorola (Stallkamp (2005), Fung et al. (2008), Girotra and Netessine (2014)). Another useful managerial implication is that the effect of monetary incentives is not necessarily monotonic: a less efficient monetary incentive does not lead to any improvement in supply chain efficiency (and can even crowd out benefits from relational incentives when they are used simultaneously), but a more efficient monetary incentive leads to gains in both quality and supply chain efficiency (and can be even further enhanced with relational incentives). We then proceed to demonstrate that a model of inequality aversion can account for the decisions we see in our data.

How can managers use these results in practice? First, if the two parties find themselves in a one-shot short-term relationship, they would benefit from moving to a long-term relationship, regardless of whether they are using some sort of monetary incentive. Second, if a long-term relationship is not possible, the two parties should ensure that any monetary incentive is not significantly discounted by the time the supplier receives it. Some ways to achieve a more efficient monetary incentive is to reduce the length of the review horizon, or for retailers (or manufacturers) to offer their own financing to suppliers to help offset any large upfront costs. Third, if a retailer or manufacturer is not in an existing contract, and is searching for a supplier, they may consider partnering with a slightly larger supplier that has reasonable access to financing, and in a long-term relationship. This last implication provides some initial evidence that runs counter to conventional thought, which suggests that retailers and manufacturers should partner with small suppliers, thus having more power in negotiations.

We believe that our work is one of the first to take a behavioral operations perspective on managing quality in supply chains. As such, there are a number of opportunities for future research. For one, we use the framework of deferred payment mechanisms, as it provides the flexibility to capture many contracting features, such as different efficiency levels of the monetary incentive, and can easily be applied in both one-shot and repeated settings. Indeed, if we had neglected to include both a relatively less efficient and more efficient monetary incentive, we would not have been able to recognize the non-monotonic effect of the monetary incentive, thus leading to erroneous conclusions. That being said, it may be useful to take a deeper dive into more specific contracts, such as trade credit. In addition, other settings and opportunities may exist for managing quality, such as when joint investments are permitted in initiatives like training, which can impact the probability of achieving high quality outcomes. Also, we assume that the retailer can observe the supplier's quality with certainty. Another extension would be to consider a situation where the retailer cannot easily identify if the supplier is responsible for any quality issues in a final product (Balachandran and Radhakrishnan 2005). Lastly, it may also be interesting to see how renegotiation or how allowing the two parties to bargain in a more unstructured environment may affect outcomes. All of these, and undoubtedly more, are exciting opportunities for future work on quality in supply chains.

In conclusion, our study provides evidence that retailers and suppliers can achieve win-win outcomes with relational incentives, and even further gains with highly efficient monetary incentives, as they lead to higher quality, higher supply chain efficiency, and a weak Pareto improvement for both party's payoffs. Additionally, higher quality products may lead to further benefits beyond those captured in our study, such as improved brand recognition and reputation for the retailer, which in turn, can lead to an even healthier relationship between the retailer and supplier in the supply chain.

Acknowledgments

We gratefully acknowledge the comments of three anonymous referees, an associate editor and the department editor (Serguei Netessine) for constructive comments, which have greatly improved the paper. We also thank conference participants at the 2015 INFORMS Annual Meeting for their helpful comments. We would like to thank the Clifford Whitcomb Faculty Fellowship and the Samuel Curtis Johnson Graduate School of Management, both of Cornell University, and the Naveen Jindal School of Management of the University of Texas at Dallas, for their generous financial support.

References

Babich, Volodymyr, Christopher S. Tang. 2012. Managing opportunistic supplier product adulteration: Deferred payments, inspection, and combined mechanisms. *Manufacturing & Service Operations Management* **14**(2) 301–314.

Baiman, Stanley, Serguei Netessine, Howard Kunreuther. 2004. Procurement in supply chains when the end-product exhibits the 'weakest link' property. SSRN Working Paper (2077640).

Balachandran, Kashi R., Suresh Radhakrishnan. 2005. Quality implications of warranties in a supply chain. *Management Science* **51**(9) 1266–1277.

Beer, Ruth, Hyun-Soo Ahn, Stephen Leider. 2015. The signaling and incentive effects of supplier awards. Working Paper, University of Michigan.

Bolton, Gary E., Axel Ockenfels. 2000. Erc: A theory of equity, reciprocity and competition. *American Economic Review* **90** 166–193.

Bowles, Samuel. 2008. Policies designed for self-interested citizens may undermine "the moral sentiments": Evidence from economic experiments. *Science* **320** 1605–1609.

Bowles, Samuel, Sandra Polanía-Reyes. 2012. Economic incentives and social preferences: Substitutes or complements? *Journal of Economic Literature* **50**(2) 368–425.

CEB. 2014. 8 steps for managing quality in contract manufacturing. *CEB Blogs, Procurement and Operations,* *https://www.cebglobal.com/blogs/8-steps-to-manage-quality-in-contract-manufacturing/* .

Chao, Gary H., Seyed M.R. Iravani, R. Canan Savaskan. 2009. Quality improvement incentives and product recall cost sharing contracts. *Management Science* **55**(7) 1122–1138.

Charness, Gary, Peter Kuhn. 2010. Lab labor: What can labor economists learn from the lab? Orley Ashenfelter, David Card, eds., *Handbook of Labor Economics*, vol. 4A. Amsterdam: North Holland, 229–330.

Chaudhuri, Saabira. 2013. Protect yourself when outsourcing to China. *Wall Street Journal,* *http://www.wsj.com/articles/SB10001424127887323681904578639461757495312* .

Chen, Li, Hau L. Lee. 2016. Sourcing under supplier responsibility risk: The effects of certification, audit and contingency payment. *Forthcoming, Management Science* .

ChinaImportal. 2014. Managing supplier relationships in china – 7 keys to success. *ChinaImportal, Strategy and Processes,* *http://www.chinaimportal.com/blog/managing-supplier-relationships-in-china-7-keys-to-success/* .

Eppen, Gary D., Ananth V. Iyer. 1997. Backup agreements in fashion buying - the value of upstream flexibility. *Management Science* **43**(11) 1469–1484.

FDA. 2016. Hy-vee voluntarily recalls six trail mix products due to possible health risk. *FDA,* *http://www.fda.gov/Safety/Recalls/ucm502543.htm* .

Fehr, Ernst, Simon Gächter. 2000. Cooperation and punishment in public goods experiments. *American Economic Review* **90**(4) 980–994.

Fehr, Ernst, Klaus M. Schmidt. 1999. A theory of fairness, competition and cooperation. *Quarterly Journal of Economics* **114** 817–868.

Fehr, Ernst, Klaus M. Schmidt. 2000. Fairness, incentives, and contractual choices. *European Economic Review* **44** 1057–1068.

Fehr, Ernst, Klaus M. Schmidt. 2004. Fairness and incentives in a multi-task principal-agent model. *Scandinavian Journal of Economics* **106**(3) 453–474.

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

Fung, Victor K., William K. Fung, Yoram Wind. 2008. *Competing in a Flat World: Building Enterprises for a Borderless World*. Pearson Education, Inc.

Gächter, Simon, Esther Kessler, Manfred Königstein. 2011. The role of incentives and voluntary cooperation for contractual compliance. IZA Discussion Paper No. 5774.

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

Girotra, Karan, Serguei Netessine. 2014. *The Risk-Driven Business Model: Four Questions that Will Define Your Company*. Harvard Business School Publishing Corporation.

Gneezy, Uri, Aldo Rustichini. 2000. Pay enough or don't pay at all. *The Quarterly Journal of Economics* **115**(3) 791–810.

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

Holt, Charles A., Susan K. Laury. 2002. Risk aversion and incentive effects. *American Economic Review* **92**(5) 1644–1655.

Irlenbusch, Bernd, Dirk Sliwka. 2005. Incentives, decision frames, and motivation crowding out - an experimental investigation. IZA Discussion Paper No. 1758.

Kalwani, M.U., N. Narayandas. 1995. Long-term manufacturer supplier relationships: do they pay off for supplier firms? *Journal of Marketing* **59**(1) 1–16.

Kaya, Murat, Özalp Özer. 2009. Quality risk in outsourcing: Noncontractibile product quality and private quality cost information. *Naval Research Logistics* **56** 669–685.

Klapper, Leora, Luc Laeven, Raghuram Rajan. 2010. Trade credit contracts. *Policy Research Working Paper 5328, The World Bank, Washington, DC* .

Lee, Yul W., John D. Stowe. 1993. Product risk, asymmetric information, and trade credit. *Journal of Finance and Quantitative Analysis* **28**(2) 285–300.

Mu, Liying, Milind Dawande, Xianjun Geng, Vijay Mookerjee. 2016. Milking the quality test: Improving the milk supply chain under competing collection intermediaries. *Management Science* **62**(5) 1259–1277.

Nosenzo, Daniele, Theo Offerman, Martin Sefton, Ailko van der Veen. 2016. Discretionary sanctions and rewards in the repeated inspection game. *Management Science* **62**(2) 502–517.

Oblicore. 2007. 2007 service-level management survey: Results, trends and analysis. Tech. rep., Oblicore, Inc.

Reyniers, Diane J., Charles S. Tapiero. 1995. The delivery and control of quality in supplier-producer contracts. *Management Science* **41**(10) 1581–1589.

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.

Rubin, Jared, Roman Sheremeta. 2016. Principal-agent settings with random shocks. *Management Science* **62**(4) 985–999.

Rui, Huaxia, Guoming Lai. 2015. Sourcing with deferred payment and inspection under supplier product adulteration risk. *Production and Operations Management* **24**(6) 934–946.

Saito, Kota. 2013. Social preferences under risk: Equality of opportunity versus equality of outcome. *American Economic Review* **103**(7) 3084–3101.

Sherefkin, Robert. 2002. Ford blames its suppliers for quality problems, demands improvement. *Automotive News,* *http://autoweek.com/article/car-news/ford-blames-its-suppliers-quality-problems-demands-improvement* .

Stallkamp, Thomas T. 2005. *SCORE! A Better Way to Business*. Pearson Education, Inc.

Taylor, T.A., E.L. Plambeck. 2007a. Simple relational contracts to motivate capacity investment: Price only vs. price and quantity. *Manufacturing & Service Operations Management* **9**(1) 94–113.

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

Tsay, Andy A., William S. Lovejoy. 1999. Quantity flexible contracts and supply chain performance. *Manufacturing & Service Operations Management* **1**(2) 89–111.

Tunca, T.I., S.A. Zenios. 2006. Supply auctions and relational contracts for procurement. *Manufacturing & Service Operations Management* **8**(1) 43–67.

Zhu, Kaijie, Rachel Q. Zhang, Fugee Tsung. 2007. Pushing quality improvement along supply chains. *Management Science* **53**(3) 421–436.

Appendix. Additional Analysis and Details

A. Statistics for Only Final Third of Decisions

Here we briefly present summary statics when omitting the first two-thirds of decisions: the first 20 periods in the one-shot treatments, and the first 7 cycles of the indefinite treatments. As mentioned previously, while there is evidence of experience effects in early rounds, the results presented in the main body of the manuscript continue to hold (and in many cases, are stronger).

Table A.1 **Supply Chain Efficiency in All Treatments (%), Last 10 Periods / 3 Cycles**

		Monetary Incentive		
		NM	LM	HM
Relationship	One-Shot	49.88	47.97	69.57
	Indefinite	58.03	53.62	74.47

Table A.2 **Frequency of Outcomes in All Treatments (%), Last 10 Periods / 3 Cycles**

			Monetary Incentive		
			NM	LM	HM
Relationship	One-Shot	Outside Option	63.33	44.50	20.91
		Buy, Low Effort	36.00	42.00	26.36
		Buy, High Effort	0.67	13.50	52.73
	Indefinite	Outside Option	43.70	37.44	14.43
		Buy, Low Effort	39.26	28.31	24.16
		Buy, High Effort	17.04	34.25	61.41

Table A.3 **Summary of Contract Terms in All Treatments, Last 10 Periods / 3 Cycles**

			Monetary Incentive		
			NM	LM	HM
Relationship	One-Shot	Total Compensation	19.36	34.01	38.26
		Fixed Payment (F)	19.36	14.01	5.12
		Effective Bonus (αB)	0.00	20.00	33.14
	Indefinite	Total Compensation	34.88	34.60	41.47
		Fixed Payment (F)	34.88	12.64	13.45
		Effective Bonus (αB)	0.00	21.96	28.02

B. Detailed Statistics for All Treatments

In the main text we focus on our main experimental results and metrics. Here we provide additional statistics, such as incentive compatibility (IC) rates and individual rationality (IR) rates.

Table B.1 Detailed Statistics in the One-Shot Treatments

(a) All Periods

	Monetary Incentive		
Variable	NM	LM	HM
Fixed Payment (F)	23.02	15.66	7.84
Effective Bonus (αB)	—	17.94	30.04
% IC	0.00	24.67	71.52
% IR	16.00	58.67	89.85
% IC and IR	0.00	20.67	68.33
% Reject	43.56	26.67	19.24
% High Effort (Cond. Accept)	5.51	22.50	55.35
% Retailer Buys (Cond. Accept)	72.05	83.41	94.56
Supplier Exp Profit	12.40	11.81	12.91
Retailer Exp Profit	7.81	7.42	13.48
Supply Chain Efficiency	50.52	48.07	65.99

(b) Last 10 Periods

	Monetary Incentive		
Variable	NM	LM	HM
Fixed Payment (F)	19.36	14.01	5.12
Effective Bonus (αB)	—	20.00	33.14
% IC	0.00	31.00	86.36
% IR	5.33	57.50	96.36
% IC and IR	0.00	25.50	82.73
% Reject	54.00	27.50	17.73
% High Effort (Cond. Accept)	1.45	25.52	64.09
% Retailer Buys (Cond. Accept)	79.71	76.55	96.13
Supplier Exp Profit	10.90	11.15	12.63
Retailer Exp Profit	9.05	8.04	15.20
Supply Chain Efficiency	49.88	47.97	69.57

Table B.2 Detailed Statistics in the Indefinite Treatments

(a) All Periods

Variable	Monetary Incentive NM	LM	HM
Fixed Payment (F)	35.87	12.74	14.24
Effective Bonus (αB)	—	21.53	26.49
% IC (one-shot)	0.00	43.58	60.78
% IC^r	31.72	62.10	90.53
% IR	62.12	66.79	95.36
% IC and IR	0.00	39.63	60.58
% IC^r and IR	31.72	61.85	90.53
% Reject (Cond. Offer)	46.96	56.79	29.10
% High Effort (Cond. Accept)	34.90	36.42	63.82
% Retailer Buys (Cond. Accept)	63.89	72.36	88.61
Supplier Exp Profit	14.92	11.39	15.44
Retailer Exp Profit	8.85	9.10	13.20
Supply Chain Efficiency	59.43	51.23	71.60

(b) Last Three Cycles

Variable	Monetary Incentive NM	LM	HM
Fixed Payment (F)	34.88	12.64	13.45
Effective Bonus (αB)	—	21.96	28.02
% IC (one-shot)	0.00	40.18	72.15
% IC^r	25.56	67.12	94.97
% IR	55.93	72.15	98.32
% IC and IR	0.00	38.81	72.15
% IC^r and IR	25.56	67.12	94.97
% Reject (Cond. Offer)	39.78	50.00	25.84
% High Effort (Cond. Accept)	32.19	45.76	69.09
% Retailer Buys (Cond. Accept)	65.24	77.40	92.73
Supplier Exp Profit	14.79	11.52	15.68
Retailer Exp Profit	8.41	9.93	14.11
Supply Chain Efficiency	58.03	53.62	74.47

C. The NM Treatment with Fairness Concerns

If we consider the NM treatment, then we can show that there is a threshold, $\bar{\lambda}^a$, such that the supplier can be incentivized to exert high effort if and only if $\lambda^a \geq \bar{\lambda}^a$.

We can rewrite the constraints faced by the retailer as:

$$(IR_s)\ F - c_h - \lambda^d[v_h + c_h - 2F]^+ + \lambda^a[v_h + c_h - 2F]^- \geq u_0$$

$$(IC)\ \ - c_h - \lambda^d[v_h + c_h - 2F]^+ + \lambda^a[v_h + c_h - 2F]^- \geq -c_l - \lambda^d[\mu v_h + c_l - 2F]^+ + \lambda^a[\mu v_h + c_l - 2F]^-$$

$$(IR_r)\ v_h - F \geq u_0$$

We first show that the IC constraint can never be satisfied if $v_h + c_h - 2F > 0$ and $\mu v_h + c_l - 2F > 0$. In this case IC becomes:

$$-c_h - \lambda^d(v_h + c_h - 2F) \geq -c_l - \lambda^d(\mu v_h + c_l - 2F),$$

which we can rewrite as:

$$\lambda_d \leq \frac{c_l - c_h}{(1-\mu)v_h + c_h - c_l} < 0.$$

However, this is impossible since, by assumption, $\lambda^d \geq 0$ to capture disadvantageous inequality aversion.

Now, assume that $v_h + c_h - 2F \geq 0$ and $\mu v_h + c_l - 2F < 0$. In this case, the IC constraint can be written as:

$$F \geq \frac{c_h - c_l + (v_h + c_h)\lambda^d + (\mu v_h + c_l)\lambda^a}{2(\lambda^a + \lambda^d)}.$$

Observe that this is valid if and only if $F \leq {}^{(v_h+c_h)}/_2$. With some work, it can be shown that

$$\frac{c_h - c_l + (v_h + c_h)\lambda^d + (\mu v_h + c_l)\lambda^a}{2(\lambda^a + \lambda^d)} \leq \frac{(v_h + c_h)}{2}$$

if and only if

$$\lambda^a \geq \frac{c_h - c_l}{v_h(1-\mu) + c_h - c_l}.$$

We must then check that the individual rationality constraint is satisfied. This amounts to:

$$F \geq \frac{u_0 + c_h + \lambda^d(v_h + c_h)}{1 + 2\lambda^d} \in [u_0 + c_h, {}^{(v_h+c_h)}/_2]$$

for all $\lambda^d \geq 0$.

Finally, suppose that $v_h + c_h - 2F \leq 0$. Then the IC constraint becomes:

$$-c_h + \lambda^a(v_h + c_h) \geq -c_l + \lambda^a(\mu v_h + c_l),$$

which can be rewritten as

$$\lambda^a \geq \frac{c_h - c_l}{v_h(1-\mu) + c_h - c_l}.$$

Thus we have shown that the supplier can be induced to choose high effort if and only if $\lambda^a \geq \bar{\lambda}^a = \frac{c_h - c_l}{v_h(1-\mu)+c_h-c_l}$. In this case, the retailer sets:

$$F = \max\left\{\frac{u_0 + c_h + \lambda^d(v_h + c_h)}{1 + 2\lambda^d}, \frac{c_h - c_l + (v_h + c_h)\lambda^d + (\mu v_h + c_l)\lambda^a}{2(\lambda^a + \lambda^d)}\right\}.$$