

Designing Pricing Contracts for Boundedly Rational Customers: Does the Framing of the Fixed Fee Matter?

Teck-Hua Ho

Haas School of Business, University of California at Berkeley, Berkeley, California 94720,
hoteck@haas.berkeley.edu

Juanjuan Zhang

Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142,
jjzhang@mit.edu

The format of pricing contracts varies substantially across business contexts, a major variable being whether a contract imposes a fixed fee payment. This paper examines how the use of the fixed fee in pricing contracts affects market outcomes of a manufacturer-retailer channel. Standard economic theories predict that channel efficiency increases with the introduction of the fixed fee and is invariant to its framing. We conduct a laboratory experiment to test these predictions. Surprisingly, the introduction of the fixed fee fails to increase channel efficiency. Moreover, the framing of the fixed fee does make a difference: an opaque frame as quantity discounts achieves higher channel efficiency than a salient frame as a two-part tariff, although these two contractual formats are theoretically equivalent.

To account for these anomalies, we generalize the standard economic model by allowing the retailer's utilities to be reference dependent so that the up-front fixed fee payment is perceived as a loss and the subsequent retail profits as a gain. We embed this reference-dependent utility function in a quantal response equilibrium framework where the retailer is allowed to make decision mistakes due to computational complexity. The key prediction of this behavioral model is that channel efficiency decreases with loss aversion for sufficiently Nash-rational retailers. Consistent with this prediction, the estimated loss-aversion coefficient is 1.37 in the two-part tariff condition, significantly higher than 1.27 in the quantity discount condition. At the same time, loss aversion dominates contract complexity in explaining the data. Lastly, we conduct a follow-up experiment to confirm the central role of loss aversion as a behavioral driver. In one condition, the retailer becomes less loss averse when we temporally compress the fixed fee payment and the realization of retail profits, which supports the loss aversion theory. In the other condition, the retailer's contract acceptance rate does not decline when we reward the manufacturer a higher cash payment for each experimental point earned, which rules out the competing hypothesis that the retailer rejects contract offers due to fairness concerns.

Key words: fixed fee; two-part tariffs; quantity discounts; distribution channels; loss aversion; behavioral economics; experimental economics

History: Accepted by Detlof von Wintefeldt, decision analysis; received November 8, 2004. This paper was with the authors 2 years and 3 months for 1 revision. Published online in *Articles in Advance* March 12, 2008.

1. Introduction

In any market transaction, a seller must determine a way to charge a buyer. This transfer of payment can take various formats. Some firms simply adopt a linear pricing rule, where customers pay a constant price for each unit bought. In other settings, typically business-to-business contexts, the pricing format can be more complex. We can classify pricing contracts by whether they stipulate a fixed fee payment, and by the number of different marginal prices (i.e., the price for an additional unit) they charge. Figure 1 presents this two-by-two taxonomy: A *linear-price contract* does not impose a fixed fee, and charges a single marginal

price independent of the sales volume. A *two-part tariff* requires a fixed fee in addition to charging a single marginal price. In other words, the buyer must first incur a fixed fee when entering the purchase agreement, and when the transaction materializes, pay a constant price for each unit bought. A *block tariff* does not impose a fixed fee, but sets multiple marginal prices, so that the buyer pays a different unit price after the purchase volume exceeds some threshold. Lastly, a *three-part tariff* consists of a fixed fee and charges multiple marginal prices. In this paper, we study how the use of the fixed fee in pricing contracts affects market outcomes. To isolate the role of

Figure 1 A Taxonomy of Pricing Contracts

	Fixed fee not charged	Fixed fee charged
Number of marginal prices = 1	Linear-price contract	Two-part tariff (Quantity discount)
Number of marginal prices > 1	Block tariff	Three-part tariff

the fixed, we focus on the simplest class of pricing contracts—ones that charge a single marginal price.¹

There are different ways to frame the same fixed fee payment. One frame is the aforementioned two-part tariff scheme, where the fixed fee is labeled as the necessary entry payment for agreeing to purchase. Alternatively, this same pricing contract can be framed as a *quantity discount* schedule, where the average price per unit decreases with the quantity sold because of the spread of the fixed fee over more units. Our research question is thus twofold: firstly, we ask how the introduction of the fixed fee affects the market outcomes; secondly and more importantly, we ask whether the effects of the fixed fee are sensitive to framing.

We choose to answer the research question in a manufacturer-retailer dyad where standard economic theories yield sharp predictions on market outcomes. In this setting, a manufacturer sells a product to a retailer, which goes on and sells the product to a group of end consumers. Both firms must independently determine a pricing contract to maximize their own profit. It is assumed that the retailer adopts a linear pricing rule. If the manufacturer also employs a linear pricing rule and charges a uniform wholesale price w per unit, it can be shown that the total channel profit is less than that of an integrated channel where the two firms set their prices cooperatively to maximize their joint profit. This inefficiency occurs in an independent channel because neither firm takes into account the externality its pricing decision imposes on the other's profit. Consequently, both charge a higher price than if they were part of the integrated channel. This is known as the "double-marginalization" problem (Spengler 1950), named for the fact that the retail price facing the end consumers is suboptimally high due to the stacking of manufacturer and retailer profit margins.

A number of solutions have been proposed to solve the double-marginalization problem (see Tirole 1988 for a summary). They typically involve the manufacturer employing a more complex pricing schedule that induces the retailer to charge the retail price of

the integrated channel (e.g., Jeuland and Shugan 1983, Moorthy 1987). Introducing a fixed fee, for example, would in theory eliminate the inefficiency. The manufacturer can adopt a two-part tariff by imposing a fixed fee on the retailer (F) in addition to charging a constant wholesale price (w) for each unit. The total revenue to manufacturer is then $F + w \cdot q$, where q denotes the total quantity sold. If w equals the marginal cost of production, which removes the manufacturer's profit margin, the retailer would want to charge the retail price of the integrated channel. The manufacturer can then divide this maximized channel profit between the two firms by choosing a fixed fee F , taking into account the retailer's relative bargaining power. The quantity discount is advocated by economists as an equivalent way to restore channel efficiency. Consider a quantity discount schedule in which the retailer is charged a unit price given by $(F/q) + w$ for a sales quantity q . Note that for the same q the two-part tariff and the quantity discount contracts confer identical revenues to the manufacturer. In fact, the latter is simply a reframing of the former in average-cost terms. Economic theories predict that market outcomes should be invariant to this reframing of fixed fees given the mathematical equivalence between the two pricing contracts.

In summary, standard economic theories make two sharp predictions for the market outcomes of a channel that consists of a manufacturer-retailer dyad:

1. The introduction of the fixed fee restores channel efficiency to the level of the integrated channel.
 2. Channel outcomes are invariant to the framing of the fixed fee. Specifically, a two-part tariff and a quantity discount achieve the same channel outcomes.
- Whether the two hypotheses hold is an open empirical question. If they hold, the manufacturer should prefer the class of pricing contracts with a fixed fee payment and should ignore its framing. If they fail, it will be important to know which pricing contract and which way of framing yield the best market outcomes.

The study contains the first experimental test of the two hypotheses. In the laboratory, we randomly assign subjects into three treatment conditions corresponding to three types of manufacturer pricing contracts: the linear-price contract, the two-part tariff, and the quantity discount. In each condition, subjects act as either manufacturers or retailers and play the corresponding pricing game, motivated by substantial monetary incentives to make optimal decisions. The experimental results contradict the standard theories' predictions: the introduction of the fixed fee does not improve channel outcomes, and the quantity discount achieves higher channel efficiency than the two-part tariff.

¹ Lim and Ho (2007) study block tariffs and investigate how the number of blocks influences market outcomes.

To account for these anomalies, we develop a behavioral model in which utilities are reference dependent. In its two-part tariff framing, the fixed fee is paid up-front before sales materialize, and the retailer earns the retail margin. Consequently, the fixed fee payment and the retail margin are temporally separated so that the retailer may register them in different mental accounts (Thaler 1980). In particular, compared to the retailer's status quo before reaching the purchase agreement, the up-front fixed fee can register as a loss and the subsequent retail proceeds a gain. These two mental accounts may thus be given different utility weights, as measured by the loss-aversion coefficient. On the other hand, when the fixed fee is framed in terms of quantity discounts, the retailer does not experience a separate stage of fixed fee payment, and is more likely to integrate in the same mental account proceeds from sales and its total payments to the manufacturer. In this way, the quantity discount frame makes the gain-loss dichotomy more opaque. If this is the case, the psychological loss attached to the fixed fee will be more severe with the two-part tariff frame, forcing the manufacturer to lower the fixed fee and, to compensate, raise the wholesale price to the detriment of channel efficiency.

Apart from loss aversion, there is a second dimension along which two-part tariffs and quantity discounts may differ: contract complexity. To capture the effects of contract complexity, we further embed the reference-dependent utility function into a quantal response equilibrium framework (McKelvey and Palfrey 1995), which allows the retailer to be less "Nash rational" and make more mistakes when the decision task is more complex.² This behavioral model nests the standard economic model as a special case, and allows us to assess the relative impact of loss aversion and contract complexity on channel outcomes. We estimate the behavioral model using the experimental data. The estimated loss-aversion coefficient is indeed significantly higher in the two-part tariff frame than in the quantity discount frame. In addition, loss aversion dominates complexity aversion in explaining the data.

We conduct a follow-up experiment to confirm loss aversion as the primary driver of the experimental results, and to rule out fairness concern as an alternative explanation. We create two additional variants of the two-part tariff contract to achieve these purposes. In the first variant, we eliminate the participation decision stage when the loss perceived from the fixed fee registers, and ask the retailer to simply

choose a retail price. The retailer, however, still retains the option to turn down the contract offer by charging a maximal retail price that results in zero demand. By combining the purchase agreement and the sales realization stages, this contract variant encourages the retailer to include the fixed fee payment and the sales proceeds in one common mental account. If channel outcomes are indeed driven by loss aversion, we should see a lower loss-aversion coefficient in this variant than in the two-part tariff condition. The second variant is identical to the original two-part tariff except that we change the conversion rate between the experimental currency and the eventual cash payment, making each experimental point worth more to the manufacturer than to the retailer. Should subjects care about fairness, we would expect the retailer to be more prone to reject the manufacturer's offer in this new condition, although the degree of loss aversion should not be influenced. The experimental results support our behavioral model: the estimated loss-aversion coefficient is indeed lower in the first variant and remains the same in the second variant; the retailer does not reject offers more frequently in the second variant, contrary to the prediction of the fairness hypothesis.

The rest of this paper is organized as follows. Section 2 outlines the standard economic model and formulates the main hypotheses. Section 3 describes the design and implementation of the experiment. Section 4 reports the experimental results. Section 5 develops a behavioral model and estimates it using the experimental data. Section 6 presents the second study designed to explicitly test the loss-aversion hypothesis and to assess fairness concern as a competing explanation. Section 7 concludes the paper and discusses future research directions.

2. The Standard Economic Model and Hypotheses

In this section we present the standard economic model, which illustrates how the use of linear-price contracts can cause inefficiencies in an independent channel, and how the introduction of the fixed fee restores channel efficiency. We then formally state the channel outcome hypotheses derived from the standard economic model.

We consider a simple distribution channel that consists of a one-manufacturer-one-retailer dyad. The manufacturer, an upstream monopolist, produces a product at a constant marginal cost c . The manufacturer extends a take-it-or-leave-it offer to the retailer that specifies a uniform wholesale price w . The retailer, in turn, is a monopolist seller in the end consumer market. It incurs no additional selling costs, and charges a retail price p . Demand in the end consumer market is assumed to be common knowledge,

² This approach shares the same spirit with a number of recent advances in behavioral economics (for a survey, see Camerer et al. 2003 and Ho et al. 2006).

and takes the linear form of $q = d - p$, where d is the maximum possible demand and is assumed to be greater than c .³ The product has no salvage value; hence, the retailer will ensure that the quantity sold to end consumers is equal to the quantity purchased from the manufacturer.

If the channel is integrated (i.e., the manufacturer and the retailer cooperate to maximize their joint profit), a retail price p should be chosen to maximize the total channel profit given by $\pi(p) = (p - c)(d - p)$. It follows that this efficient retail price equals $(d + c)/2$, yielding an optimal sales quantity of $(d - c)/2$ and a maximized channel profit of $(d - c)^2/4$. The manufacturer and the retailer can then negotiate a transfer price w to divide the channel profit between them.

If the channel is independent, the manufacturer and the retailer choose their own prices, w and p , respectively, to maximize their individual profits. The manufacturer moves first and offers a wholesale price w . If the retailer accepts this offer, its profit will be $\pi_R(p) = (p - w)(d - p)$. The best-response retail price that maximizes the retailer's profit is $(d + w)/2$. Rationally anticipating the retailer's response, the manufacturer would choose a w to maximize its profit of $\pi_M(w) = (w - c)(d - (d + w)/2)$ in the first place. Therefore, the equilibrium wholesale price equals $(d + c)/2 > c$, inducing a retail price of $(3d + c)/4$, which is higher than the efficient price of an integrated channel. It follows that the manufacturer earns a profit of $(d - c)^2/8$, and the retailer earns $(d - c)^2/16$. The total profit realized in this independent channel is reduced to $3(d - c)^2/16$, representing only 75% of the integrated channel profit. This inefficiency stems from the externality each firm's pricing decision imposes on the other's profit. When the firms fail to internalize such externality in an independent channel, both choose a higher than optimal price margin. The final retail price, as a result of stacking the manufacturer and retailer margins, is higher than the efficient level. This inefficiency problem associated with linear-price contracts is thus known as the double-marginalization problem.

A well-known solution to this double-marginalization problem is to introduce a fixed fee. Specifically, the manufacturer can offer a two-part tariff that imposes a fixed "franchise fee" F in addition to charging a constant marginal wholesale price w . Once the retailer enters the purchase agreement, it is committed to paying this amount of fixed fee independent of the realized sales. If the retailer buys a quantity q , it incurs a total cost of $wq + F$, and thus earns a

profit of $\pi_R(p) = (p - w)(d - p) - F$. The best-response retail price equals $(d + w)/2$, and the retailer's profit is $(d - w)^2/4 - F$. If the retailer does not have an outside option (which is the case in the experiment setting), the manufacturer can charge a fixed fee of $(d - w)^2/4$, and appropriate the entire channel profit. Therefore, the manufacturer's optimization problem becomes $\max_w \pi_M = (w - c)(d - (d + w)/2) + (d - w)^2/4$. In equilibrium, $w = c$, inducing an efficient retail price of $(d + c)/2$ and an efficient channel profit of $(d - c)^2/4$. In essence, a two-part tariff contract restores channel efficiency by eliminating the manufacturer's margin and lowering the retail price. The manufacturer then relies on the fixed fee as its source of profit.

Economics literature often juxtaposes the above two-part tariff contract with its quantity discount equivalent. Consider the contract that specifies an average unit price of $(F/q) + w$, where F and w are nonnegative constants the manufacturer needs to determine, and q is the amount purchased by the retailer. This contract represents quantity discounts because the retailer lowers its average cost by buying more units. This quantity discount scheme is equivalent to the two-part tariff in that, for a given q , both contracts accord the manufacturer the same revenue $F + wq$. Therefore, a quantity discount would restore channel efficiency in the same way a two-part tariff does: the manufacturer charges $w = c$, induces the retailer to charge the efficient price, and then extracts the maximized channel profit by setting F appropriately.

In summary, the standard economic model yields two sharp predictions about the channel outcomes, formally stated as follows:

HYPOTHESIS 1. (THE EFFICIENCY HYPOTHESIS). *The introduction of the fixed fee improves channel efficiency.*

(a) **THE STRONG EFFICIENCY HYPOTHESIS.** *A two-part tariff (TPT) restores full channel efficiency.*

$$\pi_M^{\text{TPT}} + \pi_R^{\text{TPT}} = \frac{(d - c)^2}{4},$$

where M stands for manufacturer, and R stands for retailer.

(b) **THE WEAK EFFICIENCY HYPOTHESIS.** *A two-part tariff generates a higher channel profit than a linear-price (LP) contract.*

$$\pi_M^{\text{TPT}} + \pi_R^{\text{TPT}} > \pi_M^{\text{LP}} + \pi_R^{\text{LP}}$$

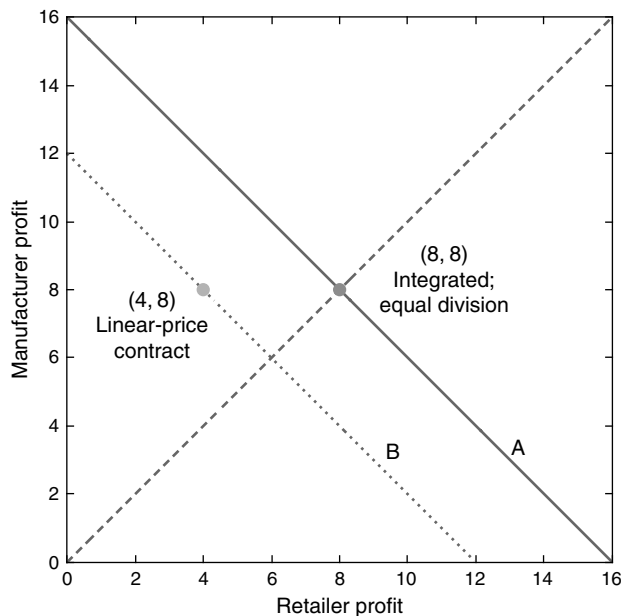
HYPOTHESIS 2. (THE FRAME INVARIANCE HYPOTHESIS). *Channel outcomes are invariant to the framing of the fixed fee.*

(a) *Total channel profit is invariant to the framing of the fixed fee. That is, a two-part tariff and a quantity discount (QD) are revenue equivalent.*

$$\pi_M^{\text{TPT}} + \pi_R^{\text{TPT}} = \pi_M^{\text{QD}} + \pi_R^{\text{QD}} = K, \quad 0 \leq K \leq \frac{(d - c)^2}{4}.$$

³Rey and Tirole (1986), Desai and Srinivasan (1995), Iyer and Villas-Boas (2003) explore the channel equilibrium in the presence of demand uncertainty.

Figure 2 The Efficiency and Frame Invariance Hypotheses



(b) The manufacturer’s share of the channel profit is invariant to the framing of the fixed fee. That is, the two-part tariff and the quantity discount are division equivalent.

$$\frac{\pi_M^{TPT}}{\pi_M^{TPT} + \pi_R^{TPT}} = \frac{\pi_M^{QD}}{\pi_M^{QD} + \pi_R^{QD}} = \alpha, \quad 0 \leq \alpha \leq 1.$$

Figure 2 illustrates the above hypotheses. We assume $d = 10$ and $c = 2$ so that the efficient channel profit equals 16. The manufacturer-retailer profit pairs obtained when using the two-part tariff contract should fall on line A if the strong efficiency hypothesis holds, and should fall northeast of line B that passes point (4, 8) (i.e., the profit pair when using the linear-price contract) if the weak efficiency hypothesis holds. If the frame invariance hypothesis holds, all profit pairs obtained using the two-part tariff and the quantity discount should fall on a line parallel to A, and should fall on a line that passes through the origin. The standard economic model predicts both hypotheses to hold so that all profit pairs under the two-part tariff and the quantity discount should fall on the same point—for example, point (8, 8) if we assume even division of the channel profit.

3. The Experiment

3.1. Experimental Design

We give the standard economic model its best shot by testing its two hypotheses in a well-controlled laboratory environment. If a theory fails in the lab, which emulates the cleanest possible market environment, we should question how well the theory applies in the field. We recruit human subjects who assume the

Table 1 Predictions of the Standard Economic Model

Variables	Treatment conditions		
	LP	TPT	QD
Wholesale price	6	2	2
Fixed fee	—	16	16
Retail price	8	6	6
Manufacturer profit	8	16	16
Retailer profit	4	0	0
Channel profit	12	16	16
Channel efficiency (%)	75	100	100
Manufacturer profit share (%)	66.7	100	100

Note. Assume $q = 10 - p$, $c = 2$, and retailer reservation utility equals 0.

role of either the manufacturer or the retailer, make the corresponding pricing decisions, and receive cash payments based on the profits they make. Subjects are randomly assigned to one of three treatment conditions: LP, TPT, and QD, where the channel contract to determine is a linear-price contract, a two-part tariff, and a quantity discount, respectively. The laboratory channel setting is identical to that in the theory section, with $q = 10 - p$ and $c = 2$. Table 1 summarizes the theoretical predictions for the three conditions. LP is predicted with a wholesale price of 6, a retail price of 8, and a total channel profit of 12. TPT and QD, on the other hand, predict the same wholesale price of 2, a retail price of 6, and a channel profit of 16. Because the retailer has no outside options in the lab, the manufacturer is predicted to charge a fixed fee of 16 in both TPT and QD conditions.

3.2. Experimental Procedures

We ran 10 experimental sessions, 2 for the LP condition, 4 for TPT, and 4 for QD. One hundred twenty undergraduate students at a West Coast university participated in the experiment.⁴ Most sessions had 12 subjects, and all sessions consisted of 11 decision rounds, so that each subject played the game 11 times. This design is meant to increase the number of observations and to capture any potential learning effects. The design also required a subject to be matched with a different partner in each round.⁵ Matching was anonymous. Subjects enrolled in the same session sat in cubicles separated by partition boards and did not know the identity of their partners. Actions were taken to prevent communication between subjects. The goal of unrepeated and

⁴ It is a common practice to use undergraduate students to test industrial organization theories (see Holt 1995). The results could, in principle, be replicated with managers. Several studies comparing professionals and students find little difference between the two groups (see Plott 1987, Ball and Cech 1996).

⁵ In sessions where only 11 subjects showed up, we had each subject play the game 10 times, with one subject sitting still in each round to ensure unrepeated matching. In sessions where more than 12 subjects participated, each subject still played only 11 times.

anonymous matching was to minimize collusion, reciprocity, and other dynamic strategic behaviors, so that each round was framed as a one-shot game, and so that the experiment setting resembled an independent channel.

Each session lasted for 90 minutes. In the beginning of a session, the experiment administrator explained the instructions to the subjects. (See the online appendix in the e-companion for the instructions used in the TPT condition. Instructions for other conditions are available upon request.)⁶ Subjects were subsequently assigned a set of exercises to ensure understanding of the task prior to the start of the experiment. At the beginning of each round, the administrator announced the subjects' random role assignment as either "player A" (the manufacturer) or "player B" (the retailer).⁷ Subjects then made decisions given their role, and received experimental point earnings based on the profits they made. The total point earnings across all rounds were converted to cash payments at the end of the session at the rate of \$0.20 per point. Each subject earned, on average, \$13. Subject decisions in each treatment condition are detailed as follows.

In the LP condition, player A moved first and chose PRICE A, an integer between 0 and 10, at which to offer the product to her partner, Player B. After receiving A's offer of PRICE A through the administrator, who acted as the mediator to maintain anonymity, player B must first decide whether to accept this offer. If she accepted, she must then choose PRICE B, an integer between 0 and 10, at which she sold the product to a group of end consumers. The quantity sold was determined by a simple chart (see Table 1 of the instructions). For each unit sold, player A earned (PRICE A – 2) points and player B earned (PRICE B – PRICE A) points. Each player's total point earnings for that round were calculated as their unit point earnings multiplied by the quantity sold. If, on the other hand, player B decided not to accept A's offer, this round ended and both players earned zero points.

The procedures in the TPT condition were similar to those in the LP condition except for the following: In addition to PRICE A, player A also had to choose a lump-sum FIXED FEE, a nonnegative integer, to charge her partner. Player B then made her decisions in two stages. First, she must decide whether to accept A's contract offer. If she did, she agreed to pay A the FIXED FEE no matter what, and proceeded to the second stage, where she chose a PRICE B for the

end consumers. A's point earnings were (PRICE A – 2) × QUANTITY + FIXED FEE, and B's point earnings were (PRICE B – PRICE A) × QUANTITY – FIXED FEE. If B turned down A's offer, the round ended and both players earned zero points.

In the QD condition, player A moved first and determined a pricing scheme. The pricing scheme had to be that the more B buys, the lower the average unit price (PRICE A). Specifically, the format of the pricing scheme was given as PRICE A = $x + y/\text{QUANTITY}$. Player A's decisions were thus to choose the values for the two coefficients x (an integer between 0 and 10) and y (a nonnegative integer). If B accepted this offer, she must choose a PRICE B for the end consumers. A's point earnings were (PRICE A – 2) × QUANTITY, and B's point earnings were (PRICE B – PRICE A) × QUANTITY. If B turned down A's offer, the round ended and both players earned zero points.

4. Results

Table 2 presents the summary statistics of the subjects' decision variables and the channel outcome variables. We discuss the experimental results by answering the following questions.

4.1. Does the Double-Marginalization Problem Exist?

The LP condition serves as a benchmark to test whether the double-marginalization problem does indeed exist. Table 3 reports the *t*-test results. The linear-price contract achieves 72.95% channel efficiency (defined as the actual channel profit divided by the integrated channel profit) over the entire

Table 2 Summary Statistics

Variables	Treatment conditions		
	LP <i>N</i> = 143	TPT <i>N</i> = 264	QD <i>N</i> = 242
Subject decision variables			
Wholesale price	5.47 (0.95)	3.96 (1.17)	3.41 (1.25)
Fixed fee	—	5.24 (2.32)	6.95 (4.17)
Acceptance (%)	93.71 (24.37)	74.24 (43.81)	82.23 (38.30)
Retail price (given acceptance)	7.75 (0.69)	6.86 (0.54)	6.71 (0.80)
Channel outcome variables			
Entire sample:			
Channel efficiency (%)	72.95 (23.22)	69.51 (41.27)	76.37 (36.18)
<i>M</i> Profit share (%)	57.40 (11.69)	63.86 (12.81)	67.91 (15.65)
Given acceptance:			
Channel efficiency (%)	77.85 (13.83)	93.62 (5.29)	92.87 (7.27)
<i>M</i> Profit share (%)	57.90 (11.91)	68.67 (11.45)	71.78 (14.60)

Notes. Values in parentheses are standard deviations. *M* profit share = manufacturer profit/channel profit. When both parties earn zero profits, *M* profit share = 50%.

⁶ An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

⁷ We avoided using the terms "manufacturer" and "retailer" to minimize anchoring from the real-world practice.

Table 3 Tests of the Standard Economic Model

Null hypothesis	<i>t</i> -statistic	<i>p</i> -value
(a) Linear-price contracts and the double-marginalization problem		
Entire sample		
Efficiency(LP) = 75%	−1.06	0.292
Efficiency(LP) = 100%	−13.93	0.000
Given acceptance		
Efficiency(LP) = 75%	2.38	0.019
Efficiency(LP) = 100%	−18.54	0.000
(b) The efficiency hypothesis		
Entire sample		
Efficiency(TPT) = 100%	−12.01	0.000
Efficiency(QD) = 100%	−10.16	0.000
Efficiency(TPT) = Efficiency(LP)	−1.08	0.283
Efficiency(QD) = Efficiency(LP)	1.13	0.259
Given acceptance		
Efficiency(TPT) = 100%	−16.87	0.000
Efficiency(QD) = 100%	−13.83	0.000
Efficiency(TPT) = Efficiency(LP)	12.59	0.000
Efficiency(QD) = Efficiency(LP)	11.55	0.000
(c) The frame invariance hypothesis		
Entire sample		
Efficiency(TPT) = Efficiency(QD)	−1.99	0.047
<i>M</i> 's profit share(TPT) = <i>M</i> 's profit share(QD)	−3.17	0.002
Given acceptance		
Efficiency(TPT) = Efficiency(QD)	1.18	0.240
<i>M</i> 's profit share(TPT) = <i>M</i> 's profit share(QD)	−2.35	0.019

sample, and 77.85% conditional on retailer acceptance, neither significantly different from the double-marginalization level of 75%. In fact, the average wholesale price 5.47 is substantially higher than the efficient level of 2 ($t = 43.75$, $p = 0.000$). As a result, the average retail price of 7.75 given retailer acceptance is significantly higher than the efficient level of 6 ($t = 29.63$, $p = 0.000$), leading to too-low sales compared to what is optimal for the entire channel. These results conform to the expectation that the double-marginalization problem occurs when only a linear-price contract is used.

4.2. Does the Efficiency Hypothesis Hold?

Panel (b) of Table 3 summarizes the test results of the efficiency hypothesis. Overall channel efficiency is 69.51% and 76.37% in the TPT and QD conditions, respectively, both significantly below 100%, failing the strong efficiency hypothesis. Moreover, overall channel efficiency in neither condition is better than that in the LP condition, failing the weak efficiency hypothesis.

Conditional on retailer acceptance, average channel efficiency is 93.62% in TPT and 92.87% in QD, both significantly higher than that in LP. Thus, the two-part tariff and the quantity discount do help alleviate

the double-marginalization problem when retailers are willing to accept the contract offers. However, full efficiency is not achieved, even conditional on retailer acceptance: the strong efficiency hypothesis is rejected at the $p = 0.000$ level for both the TPT and QD conditions.

The two-part tariff and the quantity discount fail to restore channel efficiency for two reasons. First, the wholesale price in both conditions deviates significantly from the efficient level. In TPT the average wholesale price is 3.96, higher than the efficient level of 2 ($t = 27.18$, $p = 0.000$). In QD the average wholesale price is 3.41, closer to the efficient level but still significantly higher ($t = 17.49$, $p = 0.000$). In fact, only 10.2% of manufacturers choose the efficient wholesale price of 2 in TPT, and 13.6% in QD. As a result, the average retail price given retailer acceptance is 6.86 in TPT and 6.71 in QD, both higher than the efficient level of 6 ($t = 22.39$, $p = 0.000$ for TPT, and $t = 12.48$, $p = 0.000$ for QD). Second, some retailers reject the contract offer: 74.24% of retailers accept in TPT, and 82.23% in QD, both lower than 100% at the $p = 0.000$ level ($t = -9.55$ and -7.22 , respectively). Consequently, the two-part tariff and the quantity discount turn out to be no more efficient than the linear-price contract in coordinating the channel.


4.3. Does the Frame Invariance Hypothesis Hold?

Panel (c) of Table 3 summarizes the test results of the frame invariance hypothesis. Although conditional on acceptance, channel efficiency does not differ significantly between TPT and QD, over the entire sample QD achieves a statistically higher level of channel efficiency than does TPT (76.37% versus 69.51%, $p = 0.047$). The differences in the channel decision variables between TPT and QD help explain this discrepancy: the average wholesale price is lower in QD than in TPT ($t = -5.13$, $p = 0.000$), the average fixed fee higher ($t = 5.65$, $p = 0.000$), and the acceptance rate higher ($t = 2.19$, $p = 0.029$).

To test the division equivalence property, we compute the ratio of manufacturer profit over channel profit. When the retailer turns down the offer and both parties receive zero profits, we set the manufacturer's profit share to 50%. Overall, the manufacturer's profit share is 67.91% in QD, higher than the share of 63.86% in TPT ($p = 0.002$). Conditional on retailer acceptance, the manufacturer's profit share is 71.78% in QD, higher than 68.67% in TPT ($p = 0.019$).

In sum, the data rejects the frame invariance hypothesis. The quantity discount contract generates higher channel efficiency than the two-part tariff overall, and gives the manufacturer a larger profit share.

Figure 3 Subject Decisions by Decision Round



4.4. Is There Evidence of Learning?

Figure 3 plots the average subject decisions by decision round. To statistically test for possible learning effects, we estimate a model of $x_t = \alpha + \beta \cdot (x_{t-1}/t) + \epsilon$, where x_t is the subject decision (i.e., the wholesale price, the fixed fee, whether to accept the offer, and the retail price given acceptance) in round t . α represents the steady-state value of the decision variable. β captures the learning effects: if β is significantly different from zero, subjects do learn over time and the decisions converge to the steady state α (see also Camerer and Ho 1999). The p values of the test $\beta = 0$ are 0.014 (LP), 0.205 (TPT), and 0.849 (QD) for the wholesale price, 0.015 (TPT) and 0.501 (QD) for the fixed fee, 0.219 (LP), 0.336 (TPT), and 0.982 (QD) for the acceptance decision, and 0.107 (LP), 0.478 (TPT), and 0.870 (QD) for the retail price given acceptance. As a result, there is no significant time trend in channel efficiency: the p value of the above test is 0.985 (LP), 0.298 (TPT), and 0.931 (QD). Overall, there is minimal evidence of learning across the decision rounds.

5. A Behavioral Model

The experimental results contradict the predictions of the standard economic model. The most puzzling empirical regularities are that the acceptance rate is lower, the wholesale price higher, and the fixed fee lower when the fixed fee is more salient. To account for these anomalies, we generalize the standard economic model by allowing the retailer's utility to be reference dependent. We embed this reference-dependent utility function in a quantal response equilibrium framework and estimate the resulting model using the experimental data.

5.1. The Equilibrium Two-Part Tariff

If the retailer views the status quo as a reference point when it evaluates a two-part tariff contract offer, the lump-sum fixed fee (F) represents an up-front loss, whereas the subsequent variable retail profit $(p - w)q$ is perceived as a gain. If the retailer is loss averse, it may weigh the up-front loss more heavily than the subsequent gain (Kahneman and Tversky 1979). Put differently, the retailer creates two separate mental

accounts for the fixed fee and the variable retail profit, and these accounts do not integrate dollar for dollar (Thaler 1985). This gain-loss dichotomy does not apply to the manufacturer because both the wholesale profit and the fixed fee are gains relative to its status quo.

Allowing the retailer's utility to be reference dependent, we model its decision in the same way as the standard economic model. We solve for the equilibrium by backwards induction. After observing the wholesale price and fixed fee offer, the retailer, if it accepts this offer, calculates the best-response retail price to maximize its reference-dependent utility:⁸

$$\max_p U_R = (p - w)(d - p) - \lambda F, \quad (1)$$

where $\lambda \geq 1$ stands for the loss-aversion coefficient, which is typically defined as the ratio between the marginal disutility of losses (in absolute value) and the marginal utility of gains. The retailer's best-response retail price conditional on acceptance is $p^*(w, F) = (d + w)/2$, which is the same as in the standard economic model because the fixed fee, whether weighted by λ or not, does not enter the first-order condition for the retail price. It follows that given any manufacturer offer (w, F) , the retailer's highest possible utility is $U_R^*(w, F) = (d - w)^2/4 - \lambda F$.

We model the retailer's acceptance decision in the "quantal response equilibrium" framework (McKelvey and Palfrey 1995). The central idea of the quantal response conceptualization of game equilibria is that decision makers are more likely to choose better options over worse ones, but are subject to random errors and may not always choose the best option prescribed by the Nash equilibrium. A widely adopted formulation of quantal response equilibria takes the logit form, where the idiosyncratic utility errors follow an i.i.d. extreme value distribution. The retailer's probability of accepting an offer (w, F) , given that it receives zero utility from turning down the offer, can be written as

$$\Pr(w, F) = \frac{\exp(\gamma \cdot U_R^*(w, F))}{1 + \exp(\gamma \cdot U_R^*(w, F))}, \quad (2)$$

where the logit precision parameter γ increases with the retailer's degree of "Nash rationality," and decreases in the complexity of the decision task.⁹ When

⁸ We assume the utility function to be kinked at the reference point and linear over the gain/loss domain. We adopt the linear specification because it is both tractable analytically and amenable to econometric estimation.

⁹ We simplified the decision task as much as possible in the experiment. For example, subjects were given a chart that depicted the linear relationship between demand and different integer price points (Table 1 of the attached instructions). However, we still allow the degree of contract complexity to be an empirical question and to be estimated from the experimental data.

$\gamma = 0$, the retailer makes a random choice between acceptance and rejection; when $\gamma = \infty$, the retailer accepts the offer whenever the highest possible utility it draws from the trade exceeds zero.

For the manufacturer, because both the wholesale margin and the fixed fee contribute positively to its well-being, they receive equal weight in the manufacturer's utility function. Hence, the manufacturer's optimization problem is equivalent to maximizing its expected profit, anticipating the retailer's acceptance and pricing decisions:

$$\max_{w, F} E \pi_M = \pi_M(w, F) \cdot \Pr(w, F), \quad (3)$$

where $\pi_M(w, F) = (w - c)(d - p^*(w, F)) + F$ represents the manufacturer's profit if the offer (w, F) is accepted by the retailer. An aggressive offer confers on the manufacturer a high $\pi_M(w, F)$, but has a lower chance of being accepted. The manufacturer thus needs to strike the right balance, taking into account how loss averse the retailer is. Lemma 1 presents the equilibrium manufacturer decisions as a function of the degree of retailer loss aversion and Nash rationality. Proposition 1 states the main prediction of the behavioral model.

LEMMA 1. *The equilibrium wholesale price increases with the retailer's degree of loss aversion, and does not depend on its level of Nash rationality. Formally,*

$$w^*(\lambda) = \frac{(d + c)\lambda - d}{2\lambda - 1}. \quad (4)$$

The equilibrium fixed fee decreases with the retailer's degree of loss aversion, and is jointly determined by the retailer's degree of loss aversion and its level of Nash rationality. Formally,

$$1 + \exp(\gamma \cdot U_R^*(w^*(\lambda), F^*)) = \gamma \cdot \lambda \cdot \pi_m(w^*(\lambda), F^*), \quad (5)$$


where

$$U_R^*(w^*(\lambda), F^*) = \frac{(d - w^*(\lambda))^2}{4} - \lambda F^* \quad \text{and} \\ \pi_m(w^*(\lambda), F^*) = (w^*(\lambda) - c) \frac{d - w^*(\lambda)}{2} + F^*.$$

PROPOSITION 1. *Equilibrium channel efficiency decreases with the retailer's degree of loss aversion when it is sufficiently Nash rational.¹⁰ Equilibrium channel efficiency always increases with the retailer's level of Nash rationality.*

¹⁰ The result that equilibrium channel efficiency decreases with loss aversion holds for most parameter values (see the e-companion for the exact conditions). For example, in the setting of this study, the result always holds when the loss-aversion coefficient is less than 1.34. For loss-aversion coefficients between 1.34 and 5, a Nash rationality coefficient greater than 0.41 is sufficient for the result to hold. The result holds for the parameter values estimated from the experimental data.

Figure 4 Equilibrium Channel Efficiency, Loss Aversion, and Nash Rationality



PROOF. See the online appendix in the e-companion.

The two plots in Figure 4 show how equilibrium channel efficiency $((d + w^*)/2 - c)(d - (d + w^*)/2) \cdot \Pr(w^*, F^*) / (d - c)^2 / 4$ varies with the retailer's loss-aversion coefficient λ and Nash rationality coefficient γ . For both plots we set $d = 10$ and $c = 2$. In the top plot when $\gamma = 1$, channel efficiency decreases as the retailer becomes more loss averse. When $\lambda \rightarrow \infty$, any offer charging a fixed fee will be rejected. The two-part tariff reduces to a linear-price contract, and channel efficiency approaches the double-marginalization level. The bottom plot displays how channel efficiency increases with γ , assuming $\lambda = 1$.

In sum, channel decisions and outcomes are jointly decided by the degree of retailer loss aversion and Nash rationality. In the following section, we estimate the behavioral model to assess the relative contribution of either factor in explaining the experimental results.

5.2. Estimation

We estimate the behavioral model on the experimental data using the maximum-likelihood method. We use subject decisions (the manufacturer's choices of w and F and the retailer's acceptance decision) to develop the likelihood function.¹¹

Individual observations of w_{it} and F_{it} are assumed to follow the joint normal distribution:

$$\begin{pmatrix} w_{it} \\ F_{it} \end{pmatrix} \sim N \left\{ \begin{pmatrix} w^* \\ F^* \end{pmatrix}, \begin{pmatrix} \sigma_w^2 & \rho_{wF} \sigma_w \sigma_F \\ \rho_{wF} \sigma_w \sigma_F & \sigma_F^2 \end{pmatrix} \right\}, \quad (6)$$

where i indexes the manufacturer-retailer pair, and t indexes the decision round. w^* and F^* are the equilibrium predictions of the behavioral model (see Lemma 1). The manufacturer's random decision errors, which capture the manufacturer's bounded rationality facing computational complexity, are distributed normally with mean 0 and variance σ_w^2 and σ_F^2 , respectively. These errors are assumed to be i.i.d. across pairs and across rounds.¹² In addition, the choices of w and F from the same manufacturer may be correlated. ρ_{wF} denotes this correlation coefficient.

The retailer's acceptance decision is modeled in the quantal response equilibrium framework as discussed before. The probability for retailer i to accept the contract offer in round t (that is, $accept_{it} = 1$) equals $\exp(\gamma \cdot U_R^*(w_{it}, F_{it})) / (1 + \exp(\gamma \cdot U_R^*(w_{it}, F_{it})))$, where $U_R^*(w_{it}, F_{it}) = (d - w_{it})^2 / 4 - \lambda F_{it}$ represents the highest possible utility retailer i receives, by charging the best-response retail price, from the offer (w_{it}, F_{it}) .¹³

The joint log-likelihood function is

$$LL(\lambda, \gamma, \sigma_w, \sigma_F, \rho_{wF})$$

¹¹ There are two reasons why we do not include retail price in the estimation: First, if the retailer makes an optimal pricing decision ($p^* = (d + w)/2$), the retail price is a derived variable from the wholesale price and contains no new information about the parameters of interest (i.e., λ and γ). Indeed, most retailers are able to charge the optimal price. For 90.4% of the entire sample the actual retail price p equals the best-response retail price p^* , for 6.5% of the sample $0 < |p - p^*| \leq 1$, for 2.7% of the sample $1 < |p - p^*| \leq 2$, and for 0.4% of the sample $|p - p^*| > 2$. The correlation between p and p^* over the entire sample is 0.809 ($p = 0.000$). Second, we ran a separate estimation including the actual retail prices, and obtained approximately the same parameter estimates.

¹² We also estimate the model using data from the first six rounds and the last five rounds separately to capture potential time trends in the parameters. The two sets of parameter estimates are not significantly different. In particular, the loss-aversion coefficient is 1.39 (TPT) and 1.27 (QD) for the first six rounds, and 1.38 (TPT) and 1.27 (QD) for the last five rounds.

¹³ Alternatively, when the retailer does accept the offer, actual retailer profit can be used instead of this best-response retailer profit. These two approaches generate almost identical estimates because most retailers are able to charge the optimal retail price.

$$\begin{aligned}
 &= \sum_{i=1}^I \sum_{t=1}^T \left\{ -\ln(2\pi) - \frac{1}{2} \ln |\Sigma| \right. \\
 &\quad - \frac{1}{2} \left(\begin{matrix} w_{it} - w^* \\ F_{it} - F^* \end{matrix} \right)' \Sigma^{-1} \left(\begin{matrix} w_{it} - w^* \\ F_{it} - F^* \end{matrix} \right) \\
 &\quad + \text{accept}_{it} \cdot \gamma \cdot U_R^*(w_{it}, F_{it}) \\
 &\quad \left. - \ln[1 + \exp(\gamma \cdot U_R^*(w_{it}, F_{it}))] \right\}, \quad (7)
 \end{aligned}$$

where

$$\Sigma = \begin{pmatrix} \sigma_w^2 & \rho_{wF} \sigma_w \sigma_F \\ \rho_{wF} \sigma_w \sigma_F & \sigma_F^2 \end{pmatrix}.$$

The log-likelihood function is the same for the quantity discount contract. However, we expect the estimated loss aversion coefficient λ to be smaller in the quantity discount condition because it contains an opaque frame of the fixed fee payment.

5.3. Estimation Results

Table 4 presents the estimation results.¹⁴ We first estimate the full model allowing both the loss-aversion coefficient λ and the Nash rationality coefficient γ to vary across treatment conditions. To assess the relative importance of loss aversion and contract complexity in explaining the data, we also estimate two nested models where we restrict γ or λ to be identical across conditions. The χ^2 values for the likelihood-ratio tests and the associated p -values are reported at the bottom. Both nested models are rejected at the $p = 0.05$ and $p = 0.001$ levels, respectively. We discuss the relative explanatory power of loss aversion and contract complexity in greater detail in the next section.

Notably, the loss-aversion coefficient is 1.37 in the TPT condition, and 1.27 in the QD condition, both significantly larger than 1 ($p = 0.000$).¹⁵ In other words, for every dollar paid as the fixed fee of a two-part

¹⁴ We pool the data from all treatment conditions and present the estimation results restricting the manufacturer’s decision error parameters (σ_w , σ_F , and ρ_{wF}) to be the same across conditions. We explain the “TPT-Compress” and “TPT-Unequal” conditions of “Experiment II” in the next section.

¹⁵ Typical estimates of the loss-aversion coefficient in the literature range from 1.3 to 2.6. For example, Bateman et al. (2004) report a loss-aversion estimate for goods of 1.3 after controlling for loss aversion of money. Hardie et al. (1993) find in a longitudinal study of brand choice that the effect of price increases is 1.66 times that of price cuts. Benartzi and Thaler (1995) show that a coefficient of about 2 can help resolve the equity premium puzzle. Tversky and Kahneman (1991) find 2 to be the approximate ratio of the slopes of the value function in loss and gain domains in riskless choice settings. Kahneman et al. (1990) find 2.29 to be the mean ratio of selling prices to buying prices in endowment-effect experiments. Camerer and Chua (2004) estimate a slope coefficient of 2.63 in savings contexts.

Table 4 Estimation Results

Parameters	Full model		γ identical across conditions		λ identical across conditions	
	Estimate	S. E.	Estimate	S. E.	Estimate	S. E.
Experiment I						
Loss aversion						
λ_{TPT}	1.37	(0.02)	1.40	(0.02)	1.30	(0.01)
λ_{QD}	1.27	(0.01)	1.25	(0.01)	—	—
Complexity						
γ_{TPT}	0.82	(0.06)	1.00	(0.05)	0.66	(0.05)
γ_{QD}	1.27	(0.13)	—	—	1.43	(0.09)
Experiment II						
Loss aversion						
λ_{TPT_Comp}	1.26	(0.01)	1.27	(0.02)	—	—
λ_{TPT_Uneq}	1.35	(0.02)	1.33	(0.03)	—	—
Complexity						
γ_{TPT_Comp}	0.98	(0.14)	—	—	1.02	(0.07)
γ_{TPT_Uneq}	1.25	(0.15)	—	—	1.07	(0.08)
Manufacturer decision errors						
σ_w	1.22	(0.03)	1.21	(0.03)	1.24	(0.03)
σ_F	3.37	(0.06)	3.37	(0.06)	3.39	(0.05)
ρ_{wF}	−0.70	(0.02)	−0.70	(0.02)	−0.71	(0.01)
LL	−3,085.32		−3,090.40		−3,096.48	
χ^2			10.16		22.32	
p -value			0.017		0.000	

Note. All parameter estimates are significant at the 0.01 level.

tariff, a retailer needs to receive \$1.37 in subsequent variable profit to be indifferent. However, the quantity discount reframing makes the gain-loss dichotomy less salient, resulting in significantly less loss aversion than the two-part tariff ($\chi^2(1) = 15.70$, $p < 0.001$).

The rationality/complexity coefficient γ is estimated to be 0.82 in the TPT condition, and 1.27 in the QD condition. Both estimates are significantly larger than zero ($p < 0.001$), which means that the retailer systematically prefers the option that gives a higher utility in both conditions. Also, the retailer seems significantly more rational in the QD condition than in TPT ($\chi^2(1) = 8.18$, $p < 0.005$), likely due to the fact that the TPT contract is computationally more complex.

The random errors in the wholesale price decisions are negatively correlated with the random errors in the fixed fee decisions ($\rho_{wF} = -0.70$). In other words, a lower-than-average wholesale price is accompanied by a higher-than-average fixed fee and vice versa, suggesting that the manufacturer uses these two contract instruments as substitutes for each other. This result is consistent with the prediction of the behavioral model that in equilibrium the wholesale price increases with loss aversion while the fixed fee decreases with it.

Overall, the estimation results suggest that the QD condition leads to higher channel efficiency than the TPT condition because the former induces a lower

loss-aversion coefficient and is computationally less complex. Next we shall determine the relative importance of the two factors and show that loss aversion is indeed more important than computational complexity in explaining the data.

6. Alternative Hypotheses

We conduct additional analytical, empirical, and experimental analyses to confirm that loss aversion is the primary explanation of the observed anomalies. We summarize these analyses below.

6.1. Loss Aversion vs. Contract Complexity

Our behavioral model relies on the premises that subjects are both loss averse and prone to mistakes when faced with complex contracts. Hence, it is worthwhile to determine the relative importance of the two factors in the context of this study. We establish the primary role of loss aversion in three ways.

First, we examine whether the retailer is more likely to make mistakes when the decision task is computationally more complex. Because most of the subjects who accepted the offer were able to charge the best-response retail price (see Footnote 11), we focus on the impact of contract complexity on the retailer's acceptance decision. If computational complexity is proportional to the nonlinearity of a pricing contract,¹⁶ the test becomes whether the acceptance rate is negatively correlated with the degree of nonlinearity, as measured by F/w . The correlation between acceptance and F/w is 0.000 ($p = 0.998$) in the TPT condition and -0.326 ($p = 0.000$) in the QD condition. Because the results are inconclusive, we conduct the following analyses.

Second, the quantal response equilibrium framework enables us to quantify the relative impact of contract complexity in statistical terms. If contract complexity is the primary driver of the empirical differences across the nonlinear pricing contract conditions, the estimated rationality/complexity coefficient γ should vary widely across conditions, and the model fit will suffer substantially if we restrict γ to be identical across conditions. The middle column of Table 4 shows that the nested model assuming γ to be identical across conditions can be rejected ($p = 0.017$). However, as the right column of Table 4 shows, the nested model that assumes an identical loss-aversion coefficient across conditions is even more strongly rejected ($p = 0.000$). Overall, the statistical fit of the two nested models suggests that contract complexity does help to explain the data, but loss aversion has greater explanatory power.

Third, we conduct an additional experiment to explicitly test the existence of loss aversion. We create

a new contract variant of the two-part tariff, called "TPT_Compress," where we compress the retailer's two decision stages into one. In particular, we eliminate the separate acceptance decision stage when the perceived loss from the fixed fee payment registers, and ask the retailer to only determine a retail price. However, the retailer still retains the option to reject the contract offer by charging a retail price of 10, which results in zero demand. This treatment encourages the retailer to create a common mental account for the fixed fee payment and the variable sales proceeds, thus making the gain-loss division less salient to the retailer. If loss aversion does drive subjects' behaviors in the lab, we should see a lower loss-aversion coefficient in the TPT-Compress condition.

We ran two sessions of the TPT-Compress condition. The experimental design and implementation was identical to that in the TPT condition except that the retailer was told to choose only a retail price. We estimate the behavioral model using data from the TPT-Compress condition. The middle panel of Table 4, labeled as Experiment II, presents the results. The loss-aversion coefficient in the TPT-Compress condition is 1.26, significantly lower than that in the TPT condition ($p = 0.000$), but very close to that in the QD condition ($p = 0.9124$).¹⁷ The results are supportive of loss aversion being at work.

6.2. Fairness Concerns

Recent research in behavioral economics has shown that people may not be purely self-interested and may care about the well-being of others (e.g., Rabin 1993, Fehr and Schmidt 1999).¹⁸ Fairness concerns have proven to be robust in ultimatum games (see Camerer 2003 for a comprehensive review). In a typical ultimatum game, a proposer makes a take-it-or-leave-it offer to a receiver about the division of a fixed pie. The receiver can either accept or reject this offer. If the receiver rejects the offer, both parties walk away empty-handed. Standard economic theories predict that the proposer should offer very little to the receiver, who should prefer accepting the small offer to receiving nothing. However, a number of experiments have found that the receiver typically rejects offers of less than 30% of the pie.

The decision task in our experiment has the flavor of an ultimatum game in that the manufacturer, like the proposer, makes a take-it-or-leave-it offer to the

¹⁷ For brevity, we omit the summary statistics of subject decisions, which are available upon request.

¹⁸ Cui et al. (2007) prove that a linear price contract can restore channel efficiency when the channel members are concerned about fairness. In our experiment, channel efficiency is not significantly above the double-marginalization level in the LP condition, suggesting that fairness is less likely to be at play in our experimental setting.

¹⁶ We thank a reviewer for suggesting this analysis.

Table 5 Comparison Between Accepted and Rejected Offers

Variables	Treatment conditions		Test for equality	
	TPT	QD	<i>t</i> -statistic	<i>p</i> -value
Entire sample				
Wholesale price	3.96	3.41	5.13	0.000
Fixed fee	5.24	6.95	−5.65	0.000
Best-response <i>R</i> profit	4.23	4.31	−0.39	0.696
Best-response <i>R</i> share (%)	26.57	28.20	−0.92	0.358
Accepted offers				
Wholesale price	3.74	3.56	2.00	0.046
Fixed fee	5.08	5.96	−3.79	0.000
Best-response <i>R</i> profit	4.89	4.65	1.23	0.220
Best-response <i>R</i> share (%)	32.16	30.40	1.51	0.131
Rejected offers				
Wholesale price	4.59	2.70	5.25	0.000
Fixed fee	5.71	11.53	−5.59	0.000
Best-response <i>R</i> profit	2.33	2.71	−0.81	0.420
Best-response <i>R</i> share (%)	10.46	18.02	−1.39	0.168

Notes. “Best-response *R* profit” refers to the retailer’s profit when it charges a best-response retail price given the manufacturer’s contract offer. “Best-response *R* share” refers to the retailer’s share in channel profit when it charges a best-response retail price given the manufacturer’s contract offer.

retailer, whereas the retailer, like the receiver, has the option to reject the offer if it perceives the offer as unfair.¹⁹ Therefore, fairness could potentially explain the difference in the acceptance rate between the two-part tariff and the quantity discount conditions if somehow the former leads the manufacturer to give more uneven offers. Table 5 shows the manufacturer’s offers in experimental currency to the retailer if the latter accepts and chooses the best-response retail price.²⁰ The average offers to the retailer are 4.23 in the TPT condition and 4.31 in QD. Correspondingly, the retailer is offered 26.57% of the channel profit in the TPT condition, and 28.20% in QD. Neither the retailer’s offered profits nor its offered shares are statistically different between TPT and QD. The same observation holds true if we break down the sample of offers into accepted ones and rejected ones: In the TPT condition, on average the accepted contracts offer a retailer profit of 4.89, and the rejected contracts offer 2.33, representing 32.16% and 10.46% of the channel profits, respectively. In the QD condition, the average offered retailer profits are 4.65 for the accepted offers and 2.71 for the rejected ones, representing 30.40% and 18.02% of the channel profit accordingly. Neither accepted nor rejected offers differ significantly between the two conditions. In sum, offers seem to be

¹⁹ One major difference is that the size of the pie in the channel game is not fixed exogenously, but is a function of the retail price.

²⁰ The subjects who played the retailer’s role were able to choose the best-response retail price once they accepted the offer (see Footnote 11). Hence their actual profits conditional on acceptance were very close to their offered profits.

“equally (un)fair” between TPT and QD, making fairness concerns insufficient in explaining the discrepancy of acceptance rates across conditions.

Fairness concerns alone also cannot explain why the wholesale price in the experiment is significantly higher than the marginal cost. If the retailer cares about fairness and rejects uneven offers, the manufacturer will offer the retailer a higher share of the channel profit than if the retailer is purely self-interested. However, the manufacturer can keep the wholesale price constant at the efficient level, thus maximizing the size of the pie, and then lower the fixed fee to fulfill the fairness goal. We prove this result analytically. Consider a model in which the manufacturer’s and the retailer’s utilities depend on both their own profit and the profit difference between the two firms.²¹ Specifically, we have

$$U_M = \pi_M - \theta \cdot (\pi_R - \pi_M)^+$$

$$U_R = \pi_R - \rho \cdot (\pi_M - \pi_R)^+,$$

where $\theta \geq 0$ and $\rho \geq 0$ measure the degree to which the manufacturer and the retailer, respectively, dislike being behind in payoff. For example, if the manufacturer appropriates a larger portion of the channel profit (i.e., $\pi_M > \pi_R$), the retailer will perceive a utility lower than what it gets from the pure monetary payoff (i.e., $U_R < \pi_R$), and is thus more likely to reject the offer. We show that it is optimal for the manufacturer to charge a wholesale price equal to the marginal cost (see the online appendix in the e-companion for the proof). The intuition is that fairness concerns do not contradict efficiency concerns; whereas fairness affects the equilibrium size of the fixed fee, efficiency requires that the equilibrium wholesale price remain at the marginal cost level. In other words, the observed departure of the wholesale price from the marginal cost suggests that another behavioral variable would explain the data better than fairness concerns.

The best way to test the existence of fairness concerns in our lab setting is to experimentally manipulate the salience of fairness and see whether the resulting behaviors change systematically. We create another contract variant of the two-part tariff, called “TPT-Unequal,” where we apply a higher experimental-currency-to-cash conversion rate for the manufacturer and keep that rate unchanged for the retailer. If the retailer is concerned about fairness, we should see

²¹ A general specification of fairness may incorporate three plausible behavioral assumptions: first, people care about their own payoffs; second, people dislike being behind (i.e., envy); and third, people dislike being ahead (i.e., guilt). Our central result that the equilibrium wholesale price equals the marginal cost is found to be robust across a variety of specifications.

a lower acceptance rate associated with this variant. If the manufacturer perceives the retailer to be concerned about fairness, it should tone down its offer to encourage acceptance.

We ran two sessions of the TPT-Unequal condition. The experimental design and implementation were identical to those of the TPT condition except that we changed the experimental-currency-to-cash conversion rate for the manufacturer. Each experimental point was worth 20 cents to the retailer, the same as in the TPT condition, but worth 30 cents to the manufacturer. Neither the average wholesale price nor the average fixed fee differs significantly from that in the TPT condition ($p = 0.888$ and 0.762 , respectively). However, contrary to the prediction of the fairness hypothesis, the acceptance rate in this new condition is 87.13%, significantly higher than the rate of 76.52% in the TPT condition ($p = 0.003$). The estimated loss-aversion coefficient is 1.35, which was remarkably close to the TPT level of 1.37 ($p = 0.484$) (see the middle panel of Table 4, labeled as Experiment II). The results suggest that the retailer has not rejected offers due to fairness concerns, nor has the manufacturer reacted to the retailer's possible need for fairness in making the offers.

7. Conclusions

This paper is the first to experimentally test the performance of the fixed fee in nonlinear pricing contracts. Standard economic theories make two sharp predictions: first, the introduction of the fixed fee improves channel efficiency; second, channel outcomes are invariant to the framing of the fixed fee. Data from a lab experiment contradict both predictions: nonlinear pricing contracts using fixed fees fail to improve channel efficiency. Moreover, an opaque frame of the fixed fee in terms of a quantity discount significantly improves channel efficiency compared to a salient frame of the fixed fee as in a two-part tariff contract.

To account for the experimental results, we generalize the standard economic model by allowing the retailer to be both loss averse and computational complexity averse. We solve this model analytically. The key result is that channel efficiency decreases with the degree of loss aversion for sufficiently Nash rational retailers. We estimate the generalized model using the experimental data. The loss-aversion coefficient is 1.37 in the two-part tariff condition, indeed significantly higher than 1.27 in the quantity discount condition. In addition, model fit comparison results suggest that loss aversion has more explanatory power than complexity aversion.

We conduct a series of additional analytical, empirical, and experimental studies to examine possible alternative hypotheses. In particular, we run a follow-up experiment with two new treatment conditions

featuring two contract variants of the two-part tariff. In one condition we manipulate the salience of loss aversion by temporally compressing the fixed fee payment and the sales proceeds realization. In the other condition we increase the cash value of each experimental point the manufacturer earns. The estimated loss-aversion coefficient does significantly drop in the first contract variant, further confirming the existence of loss aversion. The retailer did not reject the offers any more often in the second contract variant, contrary to the prediction of the fairness concern hypothesis.

Several directions of future research are possible. First, it would be interesting to study how the fairness perception of the retailer changes when the production cost structure is the manufacturer's private information. Second, our behavioral model can be extended to incorporate dynamics such as the evolution of loss aversion over time given subjects' past payoffs. It would also be interesting to test whether expertise helps to reduce the influence of framing. Last but not least, we need a formal theory that maps framing contexts to the degree of loss aversion. At the moment, we only give directional predictions when we experimentally make loss aversion salient or opaque. A theory that yields point predictions of the loss-aversion coefficient can be more powerful.

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Acknowledgments

The authors thank Colin Camerer, Taizan Chan, Kay-Yut Chen, Preyas Desai, Shane Frederick, Drew Fudenberg, Ganesh Iyer, Noah Lim, John Morgan, Drazen Prelec, José Silva, Duncan Simester, Miguel Villas-Boas, and Birger Wernerfelt for helpful suggestions. This paper has also benefited from the comments of seminar participants at the University of California at Berkeley, Pennsylvania State University, National University of Singapore, and Massachusetts Institute of Technology. The authors also thank the department editor, associate editor, and three anonymous reviewers for insightful comments that have helped improve this paper. This research is partly funded by Hewlett-Packard Laboratories.

References

- Ball, S. B., P.-A. Cech. 1996. Subject pool choice and treatment effects in economic laboratory research. *Research in Experimental Economics*, Vol. 6. JAI Press, Greenwich, CT, 239–292.
- Bateman, I., D. Kahneman, A. Munro, C. Starmer, R. Sugden. 2004. Is there loss aversion in buying? An adversarial collaboration. *J. Public Econom.* 89(8) 1561–1580.
- Benartzi, S., R. H. Thaler. 1995. Myopic loss aversion and the equity premium puzzle. *Quart. J. Econom.* 110(1) 73–92.
- Camerer, C. F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ.

- Camerer, C. F., Z. Chua. 2004. Experiments on intertemporal consumption with habit formation and social learning. Working paper, California Institute of Technology, Pasadena, CA.
- Camerer, C. F., T.-H. Ho. 1999. Experienced-weighted attraction learning in normal form games. *Econometrica* **67**(4) 827–874.
- Camerer, C. F., G. Loewenstein, M. Rabin. 2003. *Advances in Behavioral Economics*. Princeton University Press, Princeton, NJ.
- Cui, T. H., J. S. Raju, Z. J. Zhang. 2007. Fairness and channel coordination. *Management Sci.* **53**(8) 1303–1314.
- Desai, P. S., K. Srinivasan. 1995. Demand signalling under unobservable effort in franchising: Linear and nonlinear price contracts. *Management Sci.* **41**(10) 1608–1623.
- Fehr, E., K. M. Schmidt. 1999. A theory of fairness, competition, and cooperation. *Quart. J. Econom.* **114**(3) 817–868.
- Hardie, B. G. S., E. J. Johnson, P. S. Fader. 1993. Modeling loss aversion and reference dependence effects on brand choice. *Marketing Sci.* **12**(4) 378–394.
- Ho, T.-H., N. Lim, C. F. Camerer. 2006. Modeling the psychology of consumer and firm behavior with behavioral economics. *J. Marketing Res.* **43**(3) 307–331.
- Holt, C. A. 1995. Industrial organization: A survey of laboratory results. J. H. Kagel, A. E. Roth, eds. *Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 349–443.
- Iyer, G., J. M. Villas-Boas. 2003. A bargaining theory of distribution channels. *J. Marketing Res.* **40**(1) 80–100.
- Jeuland, A. P., S. M. Shugan. 1983. Managing channel profits. *Marketing Sci.* **2**(3) 239–272.
- Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* **47**(2) 263–292.
- Kahneman, D., J. L. Knetsch, R. H. Thaler. 1990. Experimental tests of the endowment effect and the Coase theorem. *J. Political Econom.* **98**(6) 1325–1348.
- Lim, N., T.-H. Ho. 2007. Designing price contracts for boundedly-rational customers: Does the number of blocks matter? *Marketing Sci.* **26**(3) 312–326.
- McKelvey, R. D., T. R. Palfrey. 1995. Quantal response equilibria for normal form games. *Games Econom. Behav.* **10**(1) 6–38.
- Moorthy, K. S. 1987. Managing channel profits: Comment. *Marketing Sci.* **6**(4) 375–379.
- Plott, C. R. 1987. Dimensions of parallelism: Some policy applications of experimental methods. A. E. Roth, ed. *Laboratory Experimentation in Economics: Six Points of View*. Cambridge University Press, New York.
- Rabin, M. 1993. Incorporating fairness into game theory and economics. *Amer. Econom. Rev.* **83**(5) 1281–1302.
- Rey, P., J. Tirole. 1986. The logic of vertical restraints. *Amer. Econom. Rev.* **76**(5) 921–939.
- Spengler, J. J. 1950. Vertical integration and anti-trust policy. *J. Political Econom.* **58**(4) 347–352.
- Thaler, R. H. 1980. Toward a positive theory of consumer choice. *J. Econom. Behav. Organ.* **1**(1) 39–60.
- Thaler, R. H. 1985. Mental accounting and consumer choice. *Marketing Sci.* **4**(3) 199–214.
- Tirole, J. 1988. *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.
- Tversky, A., D. Kahneman. 1991. Loss aversion in riskless choice: A reference-dependent model. *Quart. J. Econom.* **106**(4) 1039–1061.