

Making Sense of Non-Binding Retail-Price Recommendations*

Stefan Buehler[†] Dennis L. Gärtner[‡]

February 2009

Abstract

This paper provides a theoretical rationale for non-binding retail price recommendations (RPRs) in vertical supply relations. Analyzing a bilateral manufacturer-retailer relationship with repeated trade, we show that linear relational contracts can implement the surplus-maximizing outcome. If the manufacturer has private information about production costs or consumer demand, RPRs may serve as a communication device from manufacturer to retailer. We characterize the properties of efficient bilateral relational contracts with RPRs and discuss extensions to settings where consumer demand is affected by RPRs, and where there are multiple retailers or competing supply chains.

Keywords: vertical relationships, relational contracts, asymmetric information, price recommendations.

JEL Classification: D23; D43; L14; L15.

*The authors thank Daniel Halbheer, Martin Kolmar, Johan Lagerlöf, Gerd Muehlheusser, Georg Nöldeke, as well as seminar participants at the University of Basel, the University of Lugano, the University of St. Gallen, the University of Zürich, the CIE Workshop 2008 at the Copenhagen Business School, and at the DIW conference on Market Power, Vertical Restraints, and Competition Policy in Berlin for very helpful discussions. The usual disclaimer applies.

[†]University of St. Gallen, ENCORE, Centre for Industrial Economics, University of Copenhagen; stefan.buehler@unisg.ch.

[‡]University of Zürich; dennis.gaertner@soi.uzh.ch.

1 Introduction

Retail Price Recommendations (RPRs) from manufacturers to retailers are ubiquitous, and they come in many forms: They are printed on the packing of consumer goods (groceries, body care, etc.), listed on internet platforms and commercial websites (e.g., ebay.com, bmwusa.com), and attached as “sticker prices” to durable goods displayed on retailer premises (e.g., automobiles, home appliances).¹ Yet, RPRs may also be privately communicated to retailers and thus be unobservable to consumers. A key feature of RPRs is that they are non-binding in nature, that is, in contrast to Resale Price Maintenance (RPM), the manufacturer does not retain the right to control the retail price (Mathewson and Winter, 1998, 58).

Despite the ubiquity of RPRs, it is probably fair to say that the economic rationale for making RPRs is not very well understood. Why do manufacturers recommend retail prices if retailers are free to ignore their recommendations? The literature suggests two answers to this question. First, there may be a *behavioral motive* if RPRs directly affect consumers’ willingness to pay. Assuming that consumers suffer from loss aversion (Tversky and Kahneman, 1991; Kahneman et al., 1991) if the effective retail price exceeds the recommended retail price, Puppe and Rosenkranz (2006) show that a monopolistic retailer may voluntarily adhere to the recommended price.² Second, there may be an *anticompetitive motive*. If RPRs facilitate collusion among retailers (Bernheim and Whinston, 1985; Mathewson and Winter, 1998; Faber and Janssen, 2008), retailers may voluntarily adhere to them.³ Both motives indicate that, to make sense of RPRs, it is crucial to understand the conditions under which retailers voluntarily adhere to such recommendations.⁴

In this paper, we propose an alternative approach that relies neither on the behavioral nor on the anticompetitive motive. Instead, we add two important ingredients to the analysis that have found little interest in previous work. First, we observe that vertical supply relationships usually offer repeated trade opportunities. That is, supply relations are typically long-termed. Long-term supply relations are common for the provision of branded consumer products (e.g., cosmetics, watches, and clothing) and even for entire retail chains.⁵ Second, we note that manufacturers are often better informed about their own production costs than retailers. This seems

¹Recommended retail prices are also known as manufacturer’s suggested retail prices (MSRPs).

²These authors demonstrate that, if consumers’ loss aversion becomes very strong, the manufacturer recommends a price close to the monopoly price.

³Recent work by Jullien and Rey (2007) further suggests that, if retailers adhere to RPRs, manufacturers may also employ RPRs to facilitate collusion among manufacturers.

⁴Note that if one is willing to assume that manufacturers have “other” (i.e., non-modeled) means of pressuring retailers into adherence to RPRs, there is no clear distinction between RPRs and RPM.

⁵For instance, Marks & Spencer, one of the UK’s leading retailers, is found to have little turnover in its suppliers, and “some of its relationships are more than 100 years old” (Kumar, 1996, 105).

particularly relevant in the context of the provision of assembled products such as automobiles, computers, or mobile phones.

We show that these two ingredients together suggest a novel and rather different motive for making RPRs: With repeated interaction and asymmetric information about production costs, RPRs may serve as a *communication device* from manufacturer to retailer. More specifically, RPRs may be part of an (implicit) relational contract (Levin, 2003), communicating private information from manufacturer to retailer which is crucial for joint-surplus maximization. In this relational contract, the comparative statics of wholesale and retail margins, respectively, are very different from the standard setting without repeated trade. That is, along the equilibrium path of the efficient relational contract, an increase in marginal cost leads to an *inverse variation of wholesale and retail margins*, whereas the standard setting would predict a proportional variation.⁶

The inverse variation of wholesale and retail margins is an implication of our finding that the efficient relational contract must be structured such that the retailer's profit is constant (i.e., independent of the manufacturer's marginal cost) along the equilibrium path. This finding follows from the fact that the manufacturer will recommend the retail price which maximizes joint surplus (thereby truthfully revealing marginal cost) only if he is made residual claimant to the effect of cost savings on joint surplus—which necessarily requires that the retailer's profit is independent of marginal cost. Since the retail price which maximizes joint surplus is increasing in marginal cost, retail profits can only be independent of marginal cost if the retail margin is increasing (rather than decreasing) in marginal cost, offsetting the adverse demand effect generated by an increase in marginal cost. Together with the standard result that the total margin is decreasing in marginal cost, this implies an inverse variation of wholesale and retail margins.

The relational contract which we characterize in this paper has a number of desirable features. First, the contract solves a complex vertical coordination problem employing solely (i) a simple linear wholesale price scheme, (ii) linear RPRs, and (iii) an implicit understanding of the costs and benefits of adhering to the contract. The simplicity of the contract may explain the prevalence of linear supply schemes despite the well-known double marginalization problem in vertical supply relations.⁷ It is also consistent with Kumar's (1996, 105) observation that "companies that base their relationships on trust either have minimal contracts or do away

⁶Bresnahan and Reiss (1985) examine the proportionality of wholesale and retail margins in the automobile industry. Steiner (1993) and Lal and Narasimhan (1996) highlight the possibility of an inverse relationship between wholesale and retail margins, but these authors ignore the role of RPRs in determining margins.

⁷The classic reference on double marginalization is Spengler (1950). See Tirole (1988) for a textbook treatment.

with contracts altogether”.⁸ Second, implementing the profit-maximizing outcome places a very low ‘computational burden’ on the retailer: All the retailer needs to do is follow the manufacturer’s recommendation. This is particularly convenient in industries where production costs are volatile.⁹ Third, the relational contract produces more realistic predictions on how the surplus is split between manufacturer and retailer than the standard one-shot setting without repeated trade.

We consider two major extensions to our basic setting. First, we study the case where RPRs not only convey information to the retailer, but directly affect consumer demand as well.¹⁰ We show that the relational contract with RPRs can still implement the surplus-maximizing outcome, but the effective retail price may deviate systematically from the RPR along the equilibrium path. This result is consistent with the notion that consumer demand may be stimulated by ‘moon pricing’ (the practice of setting fictitiously high RPRs so as to fool consumers into thinking they are buying at bargain prices). Second, we consider a setting where both the costs of production *and* consumer demand may change over time. This is natural if the periods in our analysis are interpreted as product cycles, where a new period represents the introduction of a new variety (a new pharmaceutical drug, a new car model, etc.). We show that our results generalize naturally to this setting, even if the manufacturer has superior information on consumer demand (e.g., from pre-launch R&D and marketing studies). The key difference to the basic setting is that the RPR now has the more complex function of communicating information on both production costs and consumer demand.

We also explain how to alleviate some limitations of our approach. In particular, we discuss the requirement that parties are sufficiently patient to make the relational contract self-enforcing, the confinement to linear tariffs, and the extent to which our analysis generalizes beyond bilateral relationships.

This paper is related to the literature on relational contracts and collusion. Baker et al. (2002) and Levin (2003) investigate how history-contingent strategies in repeated games can substitute for court-enforceable contracts. In contrast to our setting, transfers are allowed in every stage game in these papers. Athey and Bagwell (2001) also consider how relational contracts can substitute for transfers, but they focus on ‘optimal’ horizontal collusion and abstract from vertical relationships.

⁸This author also notes that the majority of wholesalers in Japan actually operate without (explicit) contracts.

⁹Using data from the Dutch gasoline market, Faber and Janssen (2008) examine whether RPRs simply summarize how retailers should adjust retail prices to be in line with frequently changing input prices. Their results suggest that RPRs also help to coordinate retail prices.

¹⁰The classical, behavioral reasoning is that the RPR’s effect on consumer demand derives from ‘mental accounting’ on consumers’ behalf (cf. Thaler, 1985). Alternatively, one might imagine that the RPR contains actual informational value for rational consumers in a setting with asymmetric information.

Nocke and White (2007), in turn, consider vertical relationships, but their focus is on showing that vertical integration facilitates horizontal upstream collusion, while we analyze how RPRs can help establish vertical collusion (i.e., implement the surplus-maximizing outcome) between manufacturer and retailer.

Our analysis is also related to the literature on repeated sequential games, as the vertical supply relationship under study gives rise to a sequential stage game where the manufacturer moves before the retailer. Wen (2002) provides a Folk theorem for repeated sequential stage games, and Mailath et al. (2008) study optimal punishment in repeated extensive-form games with impatient players.¹¹ In our setting, the retailer can punish observable deviations by the manufacturer immediately (i.e., within the same stage game), so that neither the standard Folk theorem (Fudenberg and Maskin, 1986) nor Abreu's (1988; 1986) result on simple penal codes apply. Nevertheless, we show that a simple linear relational contract with RPR can implement the surplus-maximizing outcome.

Finally, our paper contributes to the literature on vertical contracting under asymmetric information. One strand of this literature focuses on the case where the retailer is privately informed and studies RPM or quantity fixing arrangements, respectively (Gal-Or, 1991; Blair and Lewis, 1994; Martimort and Piccolo, 2007). Another strand of the literature assumes that the manufacturer has private information on the demand for a new product and examines the role of signaling and screening from the perspective of marketing research (Chu, 1992; Desai and Srinivasan, 1995; Lariviere and Padmanabhan, 1997; Desai, 2000). None of these papers studies the role of RPR for vertical contracting under asymmetric information.

The remainder of the paper is structured as follows. Section 2 introduces the analytical setup and analyzes the full information version of the repeated sequential game. Section 3 examines the repeated sequential game under the assumption that the manufacturer is privately informed about the cost of production and derives the key results of our analysis. Section 4 provides a number of relevant extensions and limitations. Section 5 concludes.

2 Repeated Trade and Relational Contracts with Complete Information

This section shows how relational contracts can resolve the double marginalization issue in a world of complete information and repeated trade: The prospect of continued future cooperation can entice parties to take actions which, rather than max-

¹¹Other important contributions to the literature on repeated sequential games include Rubinstein and Wolinsky (1995) and Sorin (1995).




Figure 1: The Vertical Structure of the Supply Chain.

imize own short-run profits, achieve joint-surplus maximization. We also argue, however, that, when information is complete, this is possible *without* the explicit communication of RPRs, which sets the stage for our later analysis of repeated vertical trade under *asymmetric* information in Section 3.¹²

2.1 The Basic Stage Game

We embed our analysis in the classical model of double marginalization (Spengler, 1950), where a manufacturer M and a retailer R (the ‘vertical supply chain’) trade to serve a retail market. M produces an intermediate good at constant marginal costs $c \geq 0$. R in turn can transform this intermediate good into a final good (for simplicity, using a costless 1:1-technology), for which he faces a consumer market with demand $D(p)$, where $D'(p) < 0$, and $D''(p) < 0$ for all p (see Fig. 1). Trade proceeds as follows: M offers a wholesale price w at which R can buy an arbitrary amount of units. R then sets a retail price p , after which demand $D(p)$ materializes. Payoffs are thus $\pi^R(w, p) = (p - w) \cdot D(p)$ and $\pi^M(w, p) = (w - c) \cdot D(p)$, respectively.

Let $\bar{\pi}^M$, $\bar{\pi}^R$ denote payoffs in the (subgame-perfect) equilibrium of this game, and \bar{p} , \bar{w} the corresponding equilibrium prices. By the familiar double-marginalization logic, $\bar{\pi}^M + \bar{\pi}^R < \pi^* \equiv \max_p (\pi^R + \pi^M)$: Vertical externalities lead to a retail price \bar{p} which *exceeds* the price $p^* \equiv \arg \max_p (p - c) \cdot D(p)$ which an integrated monopolist would charge.

¹²Formally, this section’s analysis is a special case of that in Section 3. Where therefore keep the formal exposition as compact as possible and omit proofs.

2.2 Repeated Play and Relational Contracts

Several remedies to the double-marginalization problem have been suggested in the literature, such as the use of non-linear tariffs (i.e., letting M charge a fixed fee), or RPM (i.e., assigning M the right to control retail prices).¹³ This paper investigates an alternative route to avoiding double marginalization by way of ‘relational contracts’ (cf. Baker et al., 2002) in the context of *repeated* play of the above game. This approach is motivated by two observations: First, real-world relationships between suppliers and retailers typically *are* long-termed. Second, price-recommendations may embody the communication of an implicit ‘threat’ from M to R , in the spirit of “charge this retail price, or else...”. Such a threat only makes sense in a world where M has a chance to *react* to R ’s choice of p in the future.

Technically, we consider a repeated game in which the above game (the ‘stage game’) is repeated *ad infinitum* in periods $t = 0, 1, 2, \dots$, and parties discount future payoffs at rate $\delta \in (0, 1)$. Following Levin (2003), a ‘relational contract’ represents a complete plan for the relationship, specifying each party’s action for any possible history of the game. It is ‘self-enforcing’ if it describes a perfect public equilibrium of the repeated game.

It is easy to see that one equilibrium of this game is an infinite repetition of the stage-game equilibrium, resulting in payoffs $\bar{\pi}^M, \bar{\pi}^R$ in every period. There are, however, other equilibria. Particularly, consider the following ‘trigger-strategy’ equilibrium:

- M sets $w = \hat{w}$ for any history which does *not* contain M having set $w \neq \hat{w}$ or R having set $p \neq p^*$ in any previous moves; M sets $w = \bar{w}$ (the one-shot equilibrium wholesale price) otherwise;¹⁴
- R sets $p = p^*$ for any history which does *not* contain M having set $w \neq \hat{w}$ or R having set $p \neq p^*$ in any previous moves; for all other histories, R plays (myopic) best response to the w set by M in the current period.

It is easily checked that, for adequate choice of \hat{w} (so that each party’s equilibrium stage-game payoff exceeds its equilibrium payoff in the one-shot stage game) and for sufficiently patient parties (δ close enough to 1), the above strategies form a subgame-perfect equilibrium.

The following features of this equilibrium are worth noting: First, the sum of equilibrium payoffs is maximal in every round (because $p = p^*$), so double marginalization is eliminated. Second, the level of \hat{w} reflects an (implicit) agreement

¹³Note that the double-marginalization problem will also disappear if R has full bargaining power (i.e., if he can make a take-it-or-leave-it offer concerning the terms of trade).

¹⁴Observe that, even if M cannot directly observe p , he is able to infer it from his realized profits at the end of the stage game because $D' < 0$.

on how the surplus from cooperation is to be divided between M and R , but does not affect the size of the surplus (provided that \hat{w} falls within the bounds described above). As such, the repeated setting is more flexible concerning the distribution of surplus than the non-repeated setting, where full bargaining power somewhat artificially rests with M (who can make a take-it-or-leave-it offer to R). Third and finally, the specification of R 's strategy differs slightly from the usual formulation of trigger strategies because the stage-game itself is an extensive- rather than normal-form game (essentially requiring optimality of R 's action not only contingent on actions taken in previous rounds but also contingent on M 's action in the *current* round).

Naturally, there are many more equilibria in this infinitely repeated game,¹⁵ which leads to the usual issue of equilibrium selection. As we will argue below, the focus on *efficient* equilibria can be justified by the fact that, for δ large enough, the joint surplus from the relational contract can be split among parties in an arbitrary way (see Theorem 1 in Levin, 2003, for a similar argument).

2.3 Enter Retail Price Recommendations

To investigate the role of RPRs in this setting, assume now that, in every round, in addition to naming the wholesale price w , M can name a *recommended retail price* $\tilde{p} \in \mathbb{R}$. What role might this non-binding communication play?

In the context of our above cooperative equilibrium, a first superficial interpretation might view the RPR as an explicit communication of M 's strategy to R : M might recommend $\tilde{p} = p^*$ in every round so as to make explicit that R setting any other price will terminate cooperation. Strictly speaking, however, there is no need for such communication in equilibrium: Nash equilibrium requires each party's strategy to be optimal given *correct* beliefs about the strategy of the other. Consequently, in equilibrium, there is no point in such communication. In other words, efficient collusion between M and R ('supply chain efficiency') can be achieved with or without RPR because R *knows* the efficient retail price p^* (and he is assumed to know that this is the focal point of coordination).¹⁶

Essentially, however, the remainder of our analysis will argue that this is an artifact of the assumption of complete information. In reality, trade between M

¹⁵Observe that we cannot immediately apply to the standard Folk theorem to describe possible equilibrium payoffs because the stage game is an extensive-form rather than a normal-form game (see Wen, 2002, on applying the Folk theorem to repeated extensive-form games).

¹⁶Given the multiplicity of equilibria, one might see a role for RPRs as a device for equilibrium selection. Absent any clear theoretical underpinning, we shall not further investigate this role here. Moreover, observe that communication of an RPR can only help in *one* dimension of the coordination problem: focussing on *efficient* equilibria (the more 'obvious' part of the coordination problem), but not on the problem of coordinating on the terms by which surplus is split (i.e., the equilibrium level of w).

and R takes place in an ever changing environment where production cost, demand conditions, and thus the efficient retail price vary from period to period. Moreover, it seems plausible that the manufacturer is better informed about some aspects of this environment—particularly his costs of production. As we will see, repetition of trade still allows parties to coordinate on an efficient outcome by means of a relational contract. However, with M holding private information, some form of communication from M to the R —such as by means of RPRs—becomes crucial.

3 RPRs with Privately Known Production Costs

In this section, we extend our analysis to a setting where (i) the manufacturer's costs of production vary from period to period, and (ii) the manufacturer holds private information on these costs. Concerning the former, note that we need not literally interpret our setting as one in which it is precisely the same good which is being sold in every period. It may indeed be more natural to interpret the periods of the model as 'product cycles' where, in each period, the retailer introduces a new or improved version of the product traded in the supply chain (such as a new drug, a new book, or a new car model).¹⁷

In this situation, surplus-maximizing coordination between M and R crucially requires some way for the former to communicate the current level of costs to the latter. This section formalizes the point that this can be accomplished by means of RPRs.

3.1 The Setup

In each stage game, we now assume the following sequence of events:

1. marginal costs c are (independently) drawn from some publicly known distribution $F(c)$ over $[c, \bar{c}]$ and observed only by M ;^{18,19}
2. M sets a wholesale price w and communicates an RPR \tilde{p} to R ;
3. R sets a retail price p ;
4. stage game profits $\pi^M = (w - c) \cdot D(p)$ and $\pi^R = (p - w) \cdot D(p)$ are realized.

¹⁷This interpretation of course suggests the possibility of *demand* varying from period to period as well. We will return to this in Section 4.1.2.

¹⁸Equivalently, the model captures a situation in which (i) R can observe M 's costs c , but (ii) c is the sum of inherent efficiency and cost-reducing efforts (both privately known to M). This isomorphism is completely analogous to that between the Baron and Myerson (1982) model and the Laffont and Tirole (1986) model, respectively, in the context of monopoly regulation.

¹⁹The simplifying assumption that costs are independent across periods can easily be relaxed.

Observe that, given w , R 's payoff in the stage game is independent of c , so that there are no signaling effects: the equilibrium of the stage game corresponds to that under common knowledge of c . In extension to Section 2's analysis, we let $\bar{w}(c)$ and $\bar{p}(w)$ denote M and R 's respective equilibrium strategies in the stage game, and we let $\bar{\pi}^M(c)$ and $\bar{\pi}^R(c)$ denote their resulting equilibrium payoffs.²⁰ Note that $\bar{\pi}^M(c) \geq 0$ and $\bar{\pi}^R(c) \geq 0$ for all c , as each party can always ensure itself a non-negative margin (M by setting w high enough, R by setting p high enough).

Moreover, we let $\pi(p, c) \equiv (p - c) \cdot D(p)$ denote joint surplus for any cost realization c and any price p set by R , and we let $p^*(c) \equiv \arg \max_p \pi(p, c)$ and $\pi^*(c) \equiv \pi[p^*(c), c]$ denote the joint-surplus maximizing price and the maximal joint surplus, respectively. In the following, for brevity, we will frequently refer to p^* as the *efficient* retail price; that is, efficiency will refer to the *supply chain's* surplus being maximized.

Furthermore, we will make the following assumption:

Assumption 1. $D[p^*(c)] > 0$ for all $c \in [\underline{c}, \bar{c}]$.

By assumption 1, M 's private information concerns only the scale at which the market should efficiently be supplied—not *whether* it should be supplied or not.

As above, collusion between M and R will be desirable because $\bar{\pi}^M(c) + \bar{\pi}^R(c) < \pi^*(c)$ for any c due to double marginalization. However, the problem with efficient collusion is that the efficient retail price p^* depends on c . Indeed, to the extent that joint-profit maximization corresponds to the standard monopoly-pricing problem, the following comparative statics are immediate (see the Appendix for a formal proof):

Lemma 1. $0 < \partial p^* / \partial c < 1$, and $\partial \pi^* / \partial c < 0$.

Thus, the joint-surplus maximizing price p^* is increasing in c , whereas both the overall margin $p^* - c$ and joint profits π^* are decreasing in c .

Consequently, R can set the efficient retail price p^* only if there is some form of communication about c from M to R . In the following, we will investigate how a relational vertical contract with RPRs can overcome the problems of informational dispersion and double marginalization, thereby ensuring efficiency of the supply chain.

²⁰Formally, equilibrium strategies in the stage game are $\bar{p}(w) \equiv \arg \max_p (p - w) \cdot D(p)$, $\bar{w}(c) \equiv \arg \max_w (w - c) \cdot D(\bar{p}(w))$, and equilibrium payoffs are $\bar{\pi}^M(c) \equiv [\bar{w}(c) - c] \cdot D[\bar{p}(\bar{w}(c))]$, $\bar{\pi}^R(c) \equiv [\bar{p}(\bar{w}(c)) - \bar{w}(c)] \cdot D[\bar{p}(\bar{w}(c))]$.

3.2 The Relational Vertical Contract

Generally, the above repeated game will again host a myriad of equilibria. In the following, we will focus on a special class of equilibria, which will also make clearer the role of the RPR in the game:

Definition 1. An *efficient relational contract with self-enforced price recommendations (ERCP)* is an equilibrium in which, for some function $w(c)$, parties take the following actions along the equilibrium path:

- (a) in any period with cost realization c , M sets $w = w(c)$ and recommends $\tilde{p} = p^*(c)$;
- (b) R sets $p = \tilde{p}$.

Conceptually, ERCPs are a subset of equilibria which establish a self-reinforcing, implicit mutual agreement between M and R , where this agreement substitutes for a court-enforceable contract asking (i) R to adhere to M 's price recommendation \tilde{p} , and (ii) M to recommend the *efficient* retail price p^* and to adhere to a certain schedule in setting w , where this schedule is a function of true costs c . Indeed, keeping this interpretation in mind will be useful for the following analysis.

Our focus on *efficient* relational contracts is motivated by the fact that, as we shall see shortly, if players are sufficiently patient, then there *exists* an ERCP. Moreover, we will show that, given sufficient patience, ERCPs can split this efficient surplus in an essentially arbitrary way (by choice of the $w(c)$ -schedule). Consequently, the issues of efficiency and distribution are separable in the sense that any *inefficient* equilibrium is strictly Pareto dominated by some ERCP.^{21, 22}

ERCPs leave two elements of players' strategies unspecified: (i) parties' off-equilibrium strategies, and (ii) $w(c)$, the wholesale price schedule—which essentially describes how equilibrium profits are to be shared. Thus completed, agents' strategies must form an equilibrium in the sense that parties find it unprofitable to deviate from their respective strategies both on and off the equilibrium path.

To understand *how* the relational contract can discipline deviations from the equilibrium path, it is important to observe that there are two classes of deviations: (i) deviations which are (eventually) observed by the other party, and (ii) deviations

²¹See Levin (2003) for a similar argument. The main difference is that Levin finds this separability between efficiency and distribution for arbitrary *fixed* discount factors, because his model allows for fixed transfers between parties in each period. We could reproduce Levin's exact argument in our model by giving parties a one-off chance for initial fixed transfers before the first round of trade ('installment transfers').

²²Another restriction implicit in our definition of ERCP is that the schedule by which w is set (i.e., the agreed-upon division of surplus) is assumed stationary. Given that parties discount at the same rate, this restriction involves no loss of generality: It does not affect the level of δ required for an ERCP with arbitrary division of expected discounted payoffs to be sustainable.

which are not. More specifically, R can deviate by setting a retail price p which does *not* correspond to the price \tilde{p} recommended to him by M . Such a deviation will be observed by M at the end of the current stage game.²³ For M , in turn, there exist both observable deviations and unobservable deviations. *Observable* deviations consist in proposing a combination of wholesale price w and RPR \tilde{p} which is ‘unreasonable’ in the sense that this combination never occurs along the equilibrium path for *any* realization of c . More formally, R will immediately detect a deviation if M sets a w and recommends a \tilde{p} such that there exists no $\tilde{c} \in [\underline{c}, \bar{c}]$ with $w = w(\tilde{c})$ and $\tilde{p} = p^*(\tilde{c})$. *Unobservable* deviations for M , on the other hand, consist in setting a w and recommending a \tilde{p} for which there exists a $\tilde{c} \in [\underline{c}, \bar{c}]$ such that $w = w(\tilde{c})$ and $\tilde{p} = p^*(\tilde{c})$.²⁴

Intuitively, *observable* deviations can be disciplined by appropriate trigger strategies which trigger future losses (‘punishments’) for these deviations if parties are sufficiently patient (and short-term deviation payoffs are bounded). For the moment, we shall therefore assume that parties avoid observable deviations and consider how the remaining unobservable deviations can be disciplined. We will make more precise how observable deviations can be disciplined in Section 3.4 below.

3.3 Unobservable Deviations: Truthful Revelation of c

Formally, unobservable deviations involve M setting \hat{p} and w such that there exists $\tilde{c} \in [\underline{c}, \bar{c}]$ with $w = w(\tilde{c})$ and $\hat{p} = p^*(\tilde{c})$. Such deviations are equivalent to M submitting a false cost report $\tilde{c} \in [\underline{c}, \bar{c}]$ to R and, given M ’s private knowledge of c , are not detectable by R . Moreover, given that $p^*(c)$ is strictly monotone in c (see Lemma 1), adherence to the relational contract (i.e., reporting the efficient price) is equivalent to requiring a *truthful* cost report from M in every period. As the following result shows, this in turn places heavy restrictions on parties’ equilibrium payoffs:

Proposition 1. *In any ERCP in which both parties avoid all observable deviations, M will recommend $\tilde{p} = p^*(c)$ if and only if R ’s equilibrium profit is independent of c .*

See the Appendix for the proof.

²³Even if M cannot directly observe the retail price set by R , given that $D' < 0$, he can always perfectly infer it from observation of his profits.

²⁴More specifically, such deviations are not detectable *with certainty*. In our repeated context, R may of course collect the implicit costs reports \tilde{c} made in every period, compare the series of reports to the prior distribution $f(c)$, and initiate punishment based on some statistical test for whether the reported \tilde{c} ’s represent independent draws from $f(c)$. However, such a strategy necessarily involves a positive probability of punishment *on the equilibrium path* and is thus weakly dominated.

The simple intuition underlying Proposition 1 is that M will find it optimal (in fact, *strictly* optimal) to truthfully reveal his costs c only if he is made *residual claimant* to the effect of cost-savings on joint surplus—which requires that R 's equilibrium profits π^R be independent of c .²⁵

By Proposition 1, ERCPs can be fully parameterized by R 's (constant) equilibrium rent $\pi^R \in \mathbb{R}$: the level of this rent is the only remaining degree of freedom in designing the ERCP.

3.4 Observable Deviations and Trigger Strategies

Given that unobservable deviations are disciplined (by making M residual claimant), it remains to make sure that parties avoid observable deviations as well. In the following, we consider trigger strategies in which players revert *indefinitely* to their respective equilibrium strategies of the stage game ($\bar{w}(c)$ and $\bar{p}(w)$, respectively) as soon as *somebody* observably deviates from the equilibrium path.²⁶

For M , observable deviations consist in setting a w and recommending a \tilde{p} which do not match for *any* c . Such deviations will immediately be detected by R (i.e., within the stage game) and, according to his trigger strategy, cause immediate reversion to his best response in the one-shot game, $\bar{p}(w)$. Consequently, M 's best observable deviation has him set $\bar{w}(c)$ and earn stage-game equilibrium profits $\bar{\pi}^M(c)$. His incentive constraint can therefore be written as

$$\bar{\pi}^M(c) - \pi^M(c) \leq \frac{\delta}{1-\delta} \mathbb{E}[\pi^M(c) - \bar{\pi}^M(c)]. \quad (1)$$

Since (1) must hold for all c , it must also hold in expectation (over c), implying $\mathbb{E}[\pi^M(c) - \bar{\pi}^M(c)] \geq 0$,²⁷ so that the right-hand side of (1) is non-negative. Thus, for any δ , deviation can only be profitable in the first place for c such that $\pi^M(c) < \bar{\pi}^M(c)$. Moreover, since net deviation gains (the left-hand side of (1)) are bounded,²⁸ it immediately follows that, for $\mathbb{E}[\pi^M(c)] > \mathbb{E}[\bar{\pi}^M(c)]$ (if M 's expected gains from the relational contract are strictly positive), there exists a δ large enough

²⁵This result is reminiscent, for instance, of the Loeb-Magat proposal in the context of monopoly regulation (Loeb and Magat, 1979): One way to induce a monopolist with private cost information to set socially efficient prices is to award the firm a transfer the size of consumer surplus for any price it chooses, thus making the firm residual claimant on social surplus.

²⁶We do not claim that these off-equilibrium strategies constitute optimal penal codes in the sense of establishing equilibrium at a minimal required discount factor. As shown by Mailath et al. (2008), Abreu's (1988) classic result on optimal penal codes is not generally applicable in the context of repeated *extensive-form* stage games.

²⁷Observe that this condition is equivalent to M 's *ex-ante* participation constraint, describing his willingness to enter the relational contract (rather than repeated one-shot play) in the first place. Thus, M 's participation is implied by his incentive constraints.

²⁸Recall that c is drawn from a compact set.




Figure 2: Achievable Average Payoffs for δ Large Enough.

to satisfy condition (1) for all c . Thus, provided sufficient patience, M will voluntarily adhere to the relational contract even if current profits $\pi^M(c)$ fall below his equilibrium profits from the one-shot stage game: The future long-run gains from cooperation will make it worth incurring this loss. For future reference, note in particular that M can be induced to adhere to the relational contract even if current profits $\pi^M(c)$ are *negative*.

For R , in turn, any deviation (i.e., setting $p \neq \bar{p}$) will earn him the expected payoff of the stage-game equilibrium, $E_c[\bar{\pi}^R(c)]$, in all future periods. Since this future stream of expected payoffs is independent of the deviation chosen (and the current level of c), R 's optimal deviation will simply have him play his best response in the one-shot game, $p = \bar{p}(w(c))$, which will earn him $[\bar{p}(w(c)) - w(c)] \cdot D[\bar{p}(w(c))]$ instead of π^R in the deviation period. Hence, R 's incentive constraint can be formulated as

$$[\bar{p}(w(c)) - w(c)] \cdot D[\bar{p}(w(c))] - \pi^R \leq \frac{\delta}{1-\delta} \left\{ \pi^R - E[\bar{\pi}^R(c)] \right\}, \quad (2)$$

where the left-hand side is non-negative (for any $w(c)$, R can always obtain π^R by setting $p = p^*(c)$), implying $\pi^R \geq E[\bar{\pi}^R(c)]$.

Moreover, the left-hand side of (2) is again bounded, so that, for $\pi^R > E[\bar{\pi}^R(c)]$, there necessarily exists a δ large enough so that condition (2) is satisfied for all c .²⁹

In sum, this establishes the following result:

Proposition 2. *If parties are sufficiently patient, any ERCP with $E[\pi^M(c)] > E[\bar{\pi}^M(c)]$ and $\pi^R > E[\bar{\pi}^R(c)]$ (i.e., such that each party's expected equilibrium payoffs strictly exceed expected payoffs from the stage game equilibrium) can be supported.*

Fig. 2 illustrates this result: Through different choices of the $w(c)$ -schedule (representing moves *along* the bold line in Fig. 2), different ERCPs with different

²⁹Again, R 's participation constraint is implied by his incentive constraints (see Fn. 27).

divisions of expected discounted payoffs can be implemented. The shaded area shows expected payoffs resulting from *inefficient* equilibria. Fig. 2 visualizes our earlier claim (used to justify our focus on ERCP) that, provided sufficient patience, any equilibrium is strictly Pareto-dominated by some ERCP.

In analogy to the standard Folk result, different ERCPs will of course require different critical discount factors. Generally, a high discount factor is required whenever a party's short-run gains from deviation are high relative to its future gains from cooperation. Specifically, this means that higher discount factors will be required to implement ERCPs with a more *asymmetric* division of the surplus from cooperation, leading to ERCPs toward the ends of the bold line in Fig. 2 (we will further discuss the role of sufficient patience in Section 4.2.1 below).

3.5 The Role of RPRs in the Relational Contract

The above relational contract suggests the following economic rationale for RPRs: They permit M to (implicitly) communicate the current level of costs c , and thereby enables R to set the conditionally optimal retail price p .

Strictly speaking, however, this additional communication is unnecessary if M 's equilibrium announcement of w permits R to perfectly infer the current c in every round, which in turn is the case if $w(c)$ is strictly monotone in c .³⁰ As the following result (proven in the Appendix) shows, there are situations in which the $w(c)$ -schedule is non-monotone:

Lemma 2. $\text{sign}(\partial w / \partial c) = \text{sign}[\pi^M(c)]$.

Thus, the retail price will be increasing (decreasing) in c whenever M 's stage-game payoff is positive (negative). To understand this, recall that the $w(c)$ -schedule must be chosen so as to induce M to recommend the optimal price $p^*(c)$. Now, for any c and starting from $\hat{p} = p^*(c)$, a marginal increase in the price recommendation will decrease demand. For a fixed $w = w(c)$, this will decrease his profits if his markup $w(c) - c$ is positive, and increase it if the markup is negative. To neutralize this effect, recommending a higher price (pretending to have higher costs) must therefore be accompanied by a rise in w if M 's markup is positive, whereas it must be accompanied by a *fall* in w if M 's markup is negative.

³⁰More specifically, for any ERCP (and the accompanying $w(c)$ -schedule), rather than ask R to adhere to the recommendation \tilde{p} , the relational contract could instead ask R to infer c from w in each period and set the conditionally optimal $p^*(c)$. Equivalently, this could be achieved by specifying a mapping from w to p as part of the relational agreement, to which R adheres. In either case, the price recommendation \tilde{p} (or more generally, communication of more than w from M to R in every period) would become superfluous.




Figure 3: Comparing Profits and Prices: ERCP vs. One-Shot Equilibrium.

To understand why $\pi^M(c)$ may well be negative (for certain values of c), recall from our previous discussion in Section 3.4 that this is unproblematic so long as $E[\pi^M(c)] > E[\bar{\pi}^M(c)]$, so that participation in the ERCP is profitable for M in the long run. Note however that, since $\pi^M(c) = \pi^*(c) - \pi^R$, where $\partial\pi^*/\partial c < 0$ by Lemma 1 and $\partial\pi^R/\partial c = 0$ by Proposition 1, $\pi^M(c)$ is strictly decreasing in c , which implies that $\partial w/\partial c$ can change signs at most once—from positive to negative (see the illustration in Fig. 3, panel (a)).

In sum, while the $w(c)$ -schedule is not *necessarily* non-monotone, it may well be in situations where M 's expected share in joint profits from the relationship are low—as might be expected if M has low bargaining power at the time the relational contract is negotiated. In such situations, the RPR \tilde{p} forms an indispensable part of the efficient relational agreement in terms of allowing M to signal current cost conditions to the retail-price setter R .

3.6 The Economics of the ERCP

Our above explanation of RPRs rests on two ingredients: (i) repeated vertical trade, and (ii) private information of the manufacturer on his production costs. To understand the joint economic role of these two ingredients, recall from Section 2 that repetition of trade *alone* does not justify the use of RPRs because parties may achieve joint-surplus maximization without it. Second, it is easy to see that, in a one-shot version of the game in which M privately knows c , there is no use for RCPs either: Given any announced wholesale price w , R 's optimal choice of p is unaffected by his belief about c , so that R will disregard any communication about c . Hence, both ingredients are indeed vital to our interpretation of RPRs as a signal

from manufacturer to retailer: M 's private knowledge on production costs gives him something to signal about in the first place, whereas it is the repetition of trade which makes R care about such a signal.

Eventually, given these two ingredients, RPRs may be seen as an instrument which, together with a simple linear transfer scheme between M and R (and together with an 'implicit understanding' between the two) provides an observationally simple solution to a complex coordination problem.

There are however two minor caveats to this interpretation of RPRs. *First*, RPRs (i.e., communication of the efficient retail price $p^*(c)$) are only one of many possible ways for M to signal information on c to R . Formally, communication of any other strictly monotone transformation of c will do as well. *Second*, as seen in Section 3.5, given communication of $w(c)$ and the single-peakedness of the $w(c)$ -schedule, any additional *binary* form of communication (i.e., 'c is low' vs. 'c is high') in fact suffices for R to correctly infer c .

Concerning these caveats, we shall for now contend ourselves by the fact that, given that *some* additional communication concerning c is necessary, an RPR seems like a very natural instrument to achieve this. Indeed, this form of communication puts the least computational burden on the retailer, who must simply follow the manufacturer's recommendation.³¹ We will return to these caveats in Section 4, where we will argue that (i) the form of communication used becomes unambiguous once consumers respond to the RPR (as is assumed elsewhere in the literature), and that (ii) the RPR can in fact be used to communicate more than just cost-information from M to R .

3.7 Implications for Markups at the Wholesale and Retail Level

A key feature of our above analysis is that equilibrium pricing both at the wholesale and at the retail level is not driven by myopic concerns, but rather by concerns for future cooperation. As such, the analysis' predictions on how margins at the retail and wholesale level vary in c are quite different from those derived in a static framework.

Specifically, in our model, markups at the wholesale level are driven by Proposition 1, which establishes that R 's equilibrium profits must be independent of c . For the assumed linear transfer scheme, R 's equilibrium profits under the relational contract are $\pi^R(c) = [p^*(c) - w(c)] \cdot D[p^*(c)]$. Thus, the requirement that $\pi^R(c)$ be constant in c immediately translates into a requirement on R 's equilibrium markup $p^*(c) - w(c)$: To keep $\pi^R(c)$ constant, R 's equilibrium markup must be inversely

³¹More precisely, RPRs will be the efficient communication device in a world with boundedly rational retailers. This role of RPRs has previously been suggested by Faber and Janssen (2008).

related to equilibrium demand:

Corollary 1. *The wholesale-price schedule $w(c)$ in the relational contract must be such that R 's equilibrium markup, $p^*(c) - w(c)$, varies in inverse proportion to demand. Equivalently, $w(c)$ must be of the form*

$$w(c) = p^*(c) - \pi^R / D[p^*(c)] \quad (3)$$

for some $\pi^R \in \mathbb{R}$.

Since equilibrium demand is strictly decreasing in c (recall that $\partial p^* / \partial c > 0$ by Lemma 1 and $\partial D / \partial p < 0$ by assumption), and since $\pi^R > E[\bar{\pi}^R(c)] \geq 0$, R 's equilibrium markup must therefore be increasing in c . Given that the *overall* markup $p^*(c) - c$ is strictly decreasing in c by Lemma 1, this immediately yields the following result concerning the (inverse) co-movement of margins at the wholesale and retail level:

Corollary 2. *M 's markup $w - c$ is decreasing in c , whereas R 's markup $p^* - w$ is increasing in c .*

The differences in predictions between one-shot play and repeated play are illustrated in Fig. 3 which, for a specific (linear) demand function, shows prices and profits under the ERCP (for a certain level of $\bar{\pi}^R$) in panel (a), and prices and profits in the stage-game equilibrium in panel (b). Note, in particular, that R 's markup is increasing under the ERCP ($p^*(c) - w$ in panel (a)), whereas it is decreasing in the stage game equilibrium ($\bar{p}(c) - \bar{w}(c)$ in panel (b)).

The possibility of an ‘inverse’ relationship between wholesale and retail markups has previously been noted in the advertising literature. Steiner (1993, 717) emphasizes that the inverse association of wholesale and retail markups is a “prevalent although not ubiquitous phenomenon” that has largely gone unnoticed in the economics literature. Lal and Narasimhan (1996) show that manufacturer advertising may actually give rise to inversely related margins by squeezing the retail margin and increasing the producer margin.³²

Note finally that keeping π^R constant and strictly positive for all c is only possible if demand at the efficient price is always strictly positive, which is why Assumption 1 (that private information does *not* concern whether the market should be supplied or not) is crucial to our result.³³

³²Estimating a static successive monopoly model (similar to the stage game discussed in Section 2.1), Bresnahan and Reiss (1985) find proportional markups in the U.S. automobile industry.

³³One might, however, easily extend our analysis also to this case by allowing for a transfer from M to R for c such that $D(p^*(c)) = 0$. Such a payment may be thought of as a compensation for damages.

4 Extensions and Limitations

This section considers a number of relevant extensions and limitations of our analysis in Section 3.

4.1 Extensions

We begin with the case where RPRs communicate not only to the R , but also to consumers.³⁴ Letting the demand function depend both on the actual retail price and the RPR, we allow for *behavioral* aspects in consumer demand. Specifically, we examine the notion that demand may be stimulated by ‘moon pricing’ (i.e., the practice of setting ‘fictitiously’ high RPRs so as to fool consumers into thinking they are buying at bargain prices). We show that, if consumer demand is maximized at some discount from the RPR, the supply chain can maximize its surplus by means of a simple variation on the relational contract in which the RPR is systematically higher than the intended retail price. That is, R ’s downward deviation from the RPR is itself part of the implicit agreement.

Next, we consider a setting in which not only production costs, but also demand may vary across periods. This is natural if the periods in our analysis are interpreted as product cycles, where a new period represents the introduction of a new or improved product. We find that supply-chain efficiency is still attainable in the context of a relational agreement. Moreover, we argue that, if the retailer has better projections about consumer demand than the retailer, then the RPR becomes even more crucial as a communication device.

4.1.1 ‘Moon Pricing’

The rationale for RPRs in Section 3 was derived under the assumption that consumers’ willingness to pay for the final good is independent of the RPR \tilde{p} , so that consumer demand depends only on the retail price. Behavioral work suggests, however, that consumer demand may depend on both actual *and* recommended retail price. In particular, the direct demand effect of a RPR may derive from ‘mental accounting’ on consumers’ behalf (Thaler, 1985) or from loss aversion if the retail price exceeds the RPR (Puppe and Rosenkranz, 2006). In this section, we bring these two approaches together.

Suppose that consumer demand is of the form $D(p, \tilde{p})$, that is, demand is immediately affected by the RPR \tilde{p} . Facing this type of consumer demand, joint-surplus maximization now implies that, for any c , the supply chain must choose the surplus-maximizing *combination* of retail price p and RPR \tilde{p} . For any c , let $p^*(c), \tilde{p}^*(c)$

³⁴This obviously requires that RPRs are observable to consumers.

denote this surplus-maximizing combination. To fix ideas, let us assume that, for any p and c , consumer demand is maximal if consumers are offered the good at some discount $\alpha > 0$ from the *recommended* price \tilde{p} , so that $p^*(c) = (1 - \alpha) \cdot \tilde{p}^*(c)$ for any c .

It is straightforward to see that, in this case, the supply chain can maximize its total surplus with a simple *variation* on the ERCP where parties take the following actions on the equilibrium path: Instead of M recommending $\tilde{p} = p^*(c)$ and R setting $p = \tilde{p}$ for any c , M instead recommends $\tilde{p} = \tilde{p}^*(c)$ and R sets $p = (1 - \alpha) \cdot \tilde{p}$. In other words, M intentionally recommends a price *above* the intended final price (‘moon pricing’), and R marks the final price down correspondingly.³⁵

In this type of relational agreement, RPRs take on a double function: They extract the maximal willingness to pay from consumers while at the same time communicating the necessary information from M to R for implementation of the surplus-maximizing retail price. To relate this to our main analysis above, recall from our discussion in Section 3.6 that if consumer demand is *independent* of the RPR, communication of the optimal retail price $p^*(c)$ is only *one* of many ways to get the necessary information from M to R (communicating any monotone transformation of $p^*(c)$ will do as well). If, on the other hand, consumer demand *does* depend on *what* is communicated (i.e., on the RPR), then firms will pick the mode of communication which maximizes joint surplus.

Notice finally that, in this case, it is actually a part of the implicit agreement between M and R for the latter to *deviate* from the recommendation—but in a specific way which is itself part of the implicit agreement.

4.1.2 RPRs with Varying Demand Conditions

As suggested above, a natural way to interpret ‘periods’ in our model is in terms of product cycles, where a new period represents the introduction of a new variety (a new pharmaceutical drug, a new car model, etc.). Such a new variety may be associated with different consumer demand, as in the case of a car for a new market segment or a drug with new treatment characteristics. To allow for this, we extend our previous analysis to settings where not only the costs of production, but also demand changes across periods.

Particularly, we assume that the demand function in any period is given by $D(p, \theta)$, where the demand parameter θ is independently drawn (along with c) at the beginning of each stage, and where $\partial D / \partial \theta > 0$. Supply-chain efficiency now

³⁵Obviously, this type of relational agreement can be implemented more generally for any demand function $D(p, \tilde{p})$ such that $p^*(c)$ can be expressed as a function of $\tilde{p}^*(c)$.

requires the retail price to be tailored also to market demand in each period, so $p = p^*(c, \theta) \equiv \arg \max_p (p - c) \cdot D(p, \theta)$.

4.1.2.1 Commonly Known Demand Conditions

The introduction of variability in demand does not change our previous analysis significantly if θ is observed by both M and R at the beginning of each stage game. Given sufficient patience, supply-chain efficiency is again attainable—essentially by conditioning the above relational agreement on θ , in addition. The only caveat is that the variation in demand dilutes the clear-cut comparative predictions derived in Section 3.7 concerning the variation of margins at the manufacturing vs. the retail level.

More specifically, supply-chain efficiency can be attained by a simple extension to the ERCP which stipulates that, along the equilibrium path, (i) M recommends $\tilde{p} = p^*(c, \theta)$ in every period and sets w according to some schedule $w(c, \theta)$, which is again part of the implicit agreement, and (ii) that R sets $p = \tilde{p}$ in every period. By straightforward extension of Proposition 1, R 's equilibrium profit π^R must be independent of c for any given level of θ , and given sufficient patience, there always exist off-equilibrium strategies to enforce the extended ERCP.³⁶

Consequently, supply-chain efficiency is attainable through this extended ERCP (for δ large enough) if and only if (i) R 's equilibrium profit is independent of c (for any θ), and (ii) each party's expected profit from the relational contract exceeds its expected profit from the one-shot equilibrium.

While efficiency is thereby attainable, the variability in demand significantly relaxes our previous restrictions on how the wholesale price and, as a consequence, markups must vary across periods. Essentially, while π^R must be independent of c , it may vary arbitrarily in θ . Thus, while efficiency implies that M must bear the full variability in joint surplus due to variations in c , there is no similar restriction on how the variation in joint surplus due to demand fluctuations are to be shared.^{37,38}

³⁶Formally, the no-deviation constraints (1) and (2), conditioned on θ and with expectations on the right-hand side taken over c and θ , must hold for every θ . In analogy to our above analysis, it is easy to see that, so long as $E[\pi^R(\theta)] \geq E[\bar{\pi}^R(c, \theta)]$ and $E[\pi^M(c, \theta)] \geq E[\bar{\pi}^M(c, \theta)]$, the extended ERCP can be supported by reversion to the stage-game equilibrium in case of deviation (where $\bar{\pi}^i(c, \theta)$ denotes i 's equilibrium profits in the stage game).

³⁷Since $w(c, \theta) = p^*(c, \theta) - \pi^R(\theta)/D[p^*(c, \theta)]$ in straightforward extension to Corollary 1, and since $\pi^R(\theta)$ may vary arbitrarily in θ (the only restriction being a lower bound on its expected value), $w(c, \theta)$ may vary arbitrarily in θ .

³⁸As above, different extended ERCP will of course require different critical discount factors. It may therefore be possible to derive prediction on the $w(c, \theta)$ schedule under the additional assumption that, for any *expected* division of surplus, parties pick the $w(c, \theta)$ schedule so as to minimize the critical δ required (as in (Nocke and White, 2007)).

4.1.2.2 Superior Demand Information on Manufacturer's Behalf

If we interpret the periods in our models in terms of the introduction of new products, then it is easy to imagine situations in which, at the time retail prices must be determined, the manufacturer in fact has a much better idea of the size of demand to be expected: The car manufacturer is likely to have performed an extensive market analysis before developing the new model, and the manufacturer of pharmaceuticals will know more about the (demand-relevant) properties of his newly developed drug than his retail outlets. Therefore, at the time the retailer sets his price for the good, the manufacturer enjoys superior information not only on his costs of production, but also on projected demand. As we argue in this section, the RPR can then serve to communicate *both* these aspects to the retailer, thereby extending the role of RPRs as a communication device by a further dimension.³⁹

To this end, assume again that c and θ are drawn anew at the beginning of each period, but assume now that *both* are observed only by M . As we will argue in the following, supply-chain efficiency (i.e., achieving a retail price of $p^*(c, \theta)$) is again attainable by means of a simple extension to the above ERCP.

To understand the mechanics, note that there is an important difference between the nature of M 's private information on c and his private information on θ : While R does not know θ at the time he sets the price, given that $\partial D/\partial \theta > 0$, he can perfectly infer it by the end of the period through his observation of actual demand. Consequently, in terms of the categorization of possible deviations in Section 3, M deviating from the truthful (implicit) communication of θ constitutes a deviation which is eventually observable by R , and therefore can again be disciplined by adequate off-equilibrium strategies (i.e., the threat of terminating cooperation upon deviation). Enforcement of truthful revelation of θ from M to R is therefore analogous to the enforcement of R 's adherence to the RPR.

What makes a formal analysis slightly more involved is a second way in which M 's private information on θ differs from his information on c : In contrast to information on production costs, R has an immediate interest in information on demand in the sense that his best response in the stage game depends on his beliefs concerning θ . Consequently, R has an incentive to infer something about θ from M 's choice of w and \tilde{p} . For this reason, even in the stage game, there may be some (imperfect) signaling from M to R , so that characterizing the equilibrium of the one-shot game is more involved (and in contrast to our discussion in Section 3.5, the RPR may now

³⁹While the economics literature has traditionally focussed on the case in which it is the *retailer* who enjoys superior information on demand, the role of private information on behalf of the *manufacturer*—particularly when it comes to introducing new or improved goods to the market—has long been stressed in the management and marketing literature (cf. Chu, 1992; Desai and Srinivasan, 1995; Lariviere and Padmanabhan, 1997; Desai, 2000).

play a role for communication even in the stage game).⁴⁰

These technical challenges notwithstanding, for the purpose of arguing that an efficient relational contract exists, it suffices to observe that, whatever the precise nature of signaling, the one-shot equilibrium will still be inefficient from the point of view of maximizing the supply chain's surplus.⁴¹ Thus, given that (i) the short-run gains from deviating from the extended ERCP are bounded and (ii) given that both parties incur strictly positive future losses if cooperation terminates, the above trigger strategies (reversion to stage-game equilibrium strategies upon any observable deviation) can still implement the ERCP. Moreover, in this extended ERCP, the RPR now has the more complex function of communicating information on both c and θ .

4.2 Limitations

We now discuss some of the limitations of our approach, and how they may be alleviated. We begin with the requirement that parties are sufficiently patient to make the relational contract self-enforcing. Next, we explain that allowing for non-linear (rather than linear) tariff schedules would leave our key result on disciplining non-observable deviations unaffected and, if anything, enlarge the scope for vertical cooperation using relational contracts. Finally, we sketch the extent to which our analysis of bilateral vertical relationships generalizes to settings with multiple manufacturers and retailers.

4.2.1 The Role of 'Sufficient Patience'

To assure the existence of off-equilibrium strategies which support the relational contract, our analysis has assumed parties to be 'sufficiently patient'. Quite generally, we interpret this assumption as a useful polar counterpart to the case of parties being completely myopic (i.e., the non-repeated stage game), where linear transfer schedules necessarily lead to double marginalization.

Note, however, that we have not just relied on this assumption to guarantee the existence of a relational contract, but also to argue that the surplus from such a

⁴⁰Generally speaking, in the spirit of the usual analysis of cheap-talk games (cf. Crawford and Sobel, 1982), whether and how much signaling occurs in the stage-game equilibrium will depend on how strongly M and R 's interests diverge, which in turn depends on the specification of how demand depends on θ .

⁴¹Formally, it is easy to see that in any perfectly separating equilibrium of the stage game (i.e., whenever R can perfectly infer c and θ), M would have an incentive to deviate so as to strategically lower R 's price. Hence, there cannot be perfect separation in equilibrium, implying that the outcome cannot be efficient.

contract may be arbitrarily split (which in turn justifies the focus on efficient contracts). As argued in Section 3.4, the assumption of sufficient patience becomes particularly crucial if parties' bargaining power is very asymmetric at the time the terms of the relational contract are 'negotiated' (i.e., when parties coordinate on the relevant $w(c)$ schedule): More asymmetric splits of the surplus generated by the relational contract require higher discount factors. Therefore, our assumption that parties coordinate on an *efficient* ERCP becomes increasingly critical for more asymmetric bargaining positions in the sense that distribution and efficiency may no longer be separable. A straightforward way of dealing with this limitation is to allow for one-time fixed transfers at the time the relational contract is negotiated (see Footnote 21).

Alternatively, and to address the issue of limited patience more generally, one would need to explicitly characterize the set of feasible expected average payoffs for fixed δ .

4.2.2 Linear Transfer Schemes

We have deliberately focused on the use of *linear* transfer schemes to govern trade between manufacturer and retailer. This restriction concerns trade both on and off the equilibrium path, that is, within the relational contract and in the one-shot equilibrium (which parties revert to if the relational contract breaks down).

Our analysis easily extends to situations in which parties may use *non-linear* contracts on the equilibrium path. Particularly, the proof of Proposition 1 does not rely on any specific functional form of the transfer scheme, so that truthful revelation of cost information requires independence of R 's profits far more generally. The remaining observable deviations may then be disciplined as above (where non-linear transfers schemes may be used, however, to reduce parties' deviation profits, thereby reducing the critical discount factor).

Allowing for non-linear transfer schemes *off* the equilibrium path, in contrast, has more far-reaching consequences: Given the assumed sequence of moves in the stage game, non-linear contracts will ensure supply-chain efficiency, but also enable M to extract the *entire* surplus of the supply chain.⁴² Consequently, no ERCP other than the infinite repetition of this stage-game equilibrium will be supportable. That is, permitting non-linear transfer schemes off the equilibrium path implies that linear schemes can no longer be employed to achieve efficiency in the context of a relational contract. However, this seemingly threatening insight simply highlights another perhaps stronger assumption concerning the (standard) structure of the stage game: By letting M propose a take-it-or-leave-it offer, he is implicitly given absolute

⁴²It is easily seen that this is also true if M holds private information on production costs and/or demand.

bargaining power in negotiating the terms of exchange. Less asymmetric bargaining positions in the stage game will give both parties strictly positive profits in the stage game equilibrium, and thereby restore the possibility of sustaining less trivial ERCPs (with linear or non-linear transfer schemes).⁴³

4.2.3 Bilateral Relationships

For simplicity, we have focused on bilateral vertical relationships. Yet, our analysis generalizes naturally to a setting where a single manufacturer M sells through I retailers $R_i, i = 1, \dots, I$. It can be shown that, if retailers observe each other's price recommendations \tilde{p}_i , avoiding unobservable deviations by the manufacturer requires that the *sum* of R_1, \dots, R_I 's equilibrium profits is independent of the manufacturer's marginal cost c .⁴⁴

Moreover, our analysis is easily extended to a setting where K competing manufacturers M_1, \dots, M_K sell through distinct retailers R_1, \dots, R_K , facing asymmetric information about competing supply chains. In analogy to our above analysis, supply chains may now use their RPRs to facilitate internal vertical coordination so as to become more effective competitors in the market for the final product.

The extension to the case with *interlocking relationships* (Rey and Vergé, 2007) between manufacturers and retailers is more complex. To see this, note that the pattern of information exchange is far from obvious if two competing manufacturers $M_i, i = 1, 2$, sell through two retailers $R_j, j = 1, 2$, carrying *both* products. It is conceivable, for instance, that a vertical communication from M_1 to R_1 (e.g., a private RPR) is passed on across supply chains to M_2 (and vice versa), creating manifold options for strategic information exchange. Moreover, in such a setting, there are many conceivable coalitional constellations for collusion (such as collusion between manufacturers, between retailers, or between certain pairs of manufacturers and retailers). Pursuing this line of research is beyond the scope of this paper, but it might shed new light on the theory of information exchange in oligopoly (Raith, 1996). Extending the analysis along these lines might also be relevant from the perspective of antitrust policy. It would be interesting, for instance, to study whether vertical information exchange across supply chains can substitute for unlawful horizontal information exchange about retail prices, allowing supply chains to establish both vertical coordination and horizontal collusion.

⁴³As such, even under the assumption of linear contracts, the stage-game equilibrium payoffs $\bar{\pi}^M(c)$ and $\bar{\pi}^R(c)$ in our above analysis should more generally be reinterpreted as being dependent on the assumed distribution of bargaining power.

⁴⁴If retailers cannot observe each other's price recommendations (and thus cannot detect dissonant cost reports), M will recommend the surplus-maximizing retail prices $\mathbf{p}^*(c)$ if and only if the sum of R_1, \dots, R_I 's equilibrium profits is independent of *each* individual cost report \tilde{c}_i .

5 Conclusion

This paper has formalized the notion that non-binding RPRs serve as a communication device facilitating coordination within vertical supply chains. Specifically, assuming (i) repeated trade, and (ii) private information on behalf of the manufacturer, we have shown that RPRs may be part of a relational contract communicating private information from manufacturer to retailer that is indispensable for maximizing joint surplus. This relational contract has three desirable features: It is simple (using linear pricing schemes only), places minimal computational burden on the retailer (who must simply follow the RPR), and it is flexible in terms of profit distribution among manufacturer and retailer. The relational contract predicts that wholesale and retail margins are inversely related, which is inconsistent with standard (static) theory, but broadly accepted in the advertising literature. We have demonstrated that this relational vertical contract is self-enforcing provided that players are sufficiently patient.

This raises the question of why, in some industries, retailers appear to deviate systematically from RPRs. An extended version of our analysis where RPRs communicate not only from manufacturer to retailer, but also to consumers, suggests a possible answer. The relational contract can still implement the surplus-maximizing outcome, but RPRs now play a double role: In addition to communicating private information to retailers, they extract the maximal willingness to pay from consumers. As a result, actual retail prices may deviate systematically from RPRs along the equilibrium path. In particular, manufacturers may deliberately recommend a price above the intended retail price ('moon pricing').

Should non-binding RPRs be allowed or banned from an antitrust perspective? By our reasoning, RPRs serve as a communication device which facilitates the coordination of manufacturer and retailer, thereby improving supply-chain efficiency. Yet, vertical supply chains may also (perhaps imperfectly) coordinate in the absence of RPRs, so that a ban on RPRs will deteriorate (rather than eliminate) supply-chain coordination. This makes it unclear how consumers will be affected by a ban on RPRs. Moreover, we have argued that the communication via RPRs may also affect consumers more directly, with potentially undesirable effects. In particular, RPRs may serve to exploit behavioral consumers and even facilitate collusion among competing supply chains. Overall, it therefore seems fair to say that the implications for antitrust policy remain opaque. We hope to address this issue in future research.

Appendix A. Proofs

Proof of Lemma 1. To see that $\partial p^*/\partial c > 0$, observe that for any c , p^* maximizes $\pi(p, c) = (p - c)D(p)$. Since $\frac{\partial^2}{\partial p \partial c} \pi = -\frac{\partial}{\partial p} D(p) > 0$, the claim follows from elementary robust comparative statics.

To see that $\partial p^*/\partial c < 1$, define $m \equiv p - c$. The problem may then be reformulated as choosing m so as to maximize $m \cdot D(m + c)$. Since $\frac{\partial^2}{\partial m \partial c} [m \cdot D(m + c)] = \frac{\partial}{\partial p} D(p) < 0$, the claim is again implied by robust comparative statics.

Finally, by the envelope theorem, $\partial \pi^*/\partial c = -D(p^*) < 0$. \square

Proof of Proposition 1. Observe first that M picking a (\tilde{p}, w) such that there exists no $\tilde{c} \in [\underline{c}, \bar{c}]$ with $\tilde{p} \simeq p^*(\tilde{c})$ and $w = w(\tilde{c})$ would constitute an observable deviation on his behalf. Thus, M 's problem may be equivalently formulated as choosing a cost-level $\tilde{c} \in [\underline{c}, \bar{c}]$ and then announcing the equilibrium prices corresponding to this cost level. With slight abuse of notation, let $\pi^M(\tilde{c}, c) \equiv \pi(p^*(\tilde{c}), c) - \pi^R(\tilde{c})$ denote M 's profit for any such \tilde{c} and any true cost level c . Given that $\partial p^*/\partial c > 0$, M will then recommend $p^*(c)$ if and only if

$$c \in \arg \max_{\tilde{c}} \pi^M(\tilde{c}, c), \quad (\text{A.1})$$

which implies the first-order condition

$$\frac{\partial}{\partial p} \pi(p^*(c), c) \cdot \frac{\partial}{\partial c} p^*(c) - \frac{\partial}{\partial c} \pi^R(c) = 0.$$

Since $\frac{\partial}{\partial p} \pi(p^*(c), c) = 0$ by definition of $p^*(c)$, this establishes necessity of $\partial \pi^R/\partial c = 0$.

To establish sufficiency, note that π^R being independent of c and the fact that $p^*(c)$ maximizes $\pi(\cdot, c)$ implies

$$\begin{aligned} \pi^M(c, c) &= \pi[p^*(c), c] - \pi^R(c) = \pi[p^*(c), c] - \pi^R(\tilde{c}) \\ &\geq \pi[p^*(\tilde{c}), c] - \pi^R(\tilde{c}) = \pi^M(\tilde{c}, c) \end{aligned}$$

(where the inequality is strict for any c such that $p^*(c)$ is unique). \square

Proof of Lemma 2. In the context of the proof of Lemma 1, M 's profits for any true cost level c and any implicitly reported cost level \tilde{c} can be written as $\pi^M(\tilde{c}, c) = [w(\tilde{c}) - c] \cdot D[p^*(\tilde{c})]$. Using this, the first-order condition for (A.1) can be rewritten as

$$\frac{\partial}{\partial c} w = -[w(c) - c] \cdot \frac{\partial}{\partial p} D[p(c)] \cdot \frac{\partial}{\partial c} p^* / D[p(c)],$$

from which it follows that $\text{sgn}(\partial w/\partial c) = \text{sgn}(w(c) - c)$, which implies the claim. \square

References

- Abreu, Dilip (1986) 'Extremal equilibria of oligopolistic supergames.' *Journal of Economic Theory* 39(1), 191–225
- (1988) 'On the theory of infinitely repeated games with discounting.' *Econometrica* 56(6), 383–396
- Athey, Susan, and Kyle Bagwell (2001) 'Optimal collusion with private information.' *RAND Journal of Economics* 32(2), 428–465
- Baker, George, Robert Gibbons, and Kevin J. Murphy (2002) 'Relational contracts and the theory of the firm.' *Quarterly Journal of Economics* 117(1), 39–84
- Baron, David P., and Roger B. Myerson (1982) 'Regulating a monopolist with unknown costs.' *Econometrica* 50(4), 911–930
- Bernheim, B. Douglas, and Michael D. Whinston (1985) 'Common marketing agency as a device for facilitating collusion.' *RAND Journal of Economics* 16(2), 269–281
- Blair, Benjamin F., and Tracy Lewis (1994) 'Optimal retail contracts under asymmetric information.' *RAND Journal of Economics* 25(2), 284–296
- Bresnahan, Timothy F., and Peter C. Reiss (1985) 'Dealer and manufacturer margins.' *RAND Journal of Economics* 16(2), 253–268
- Chu, Wujin (1992) 'Demand signalling and screening in channels of distribution.' *Marketing Science* 11(4), 327–347
- Crawford, Vincent P., and Joel Sobel (1982) 'Strategic information transmission.' *Econometrica* 50(6), 1431–1452
- Desai, Preyas S. (2000) 'Multiple messages to retain retailers: Signaling new product demand.' *Marketing Science* 19(4), 381–389
- Desai, Preyas S., and Kannan Srinivasan (1995) 'Demand signalling under unobservable effort in franchising: Linear and nonlinear price contracts.' *Management Science* 41(10), 1608–1623
- Faber, Riemer P., and Martin C.W. Janssen (2008) 'On the effects of advisory prices in gasoline markets.' Mimeo, Erasmus University Rotterdam
- Fudenberg, Drew, and Eric Maskin (1986) 'The folk theorem in repeated games with discounting or with incomplete information.' *Econometrica* 54(3), 533–554

- Gal-Or, Esther (1991) 'Optimal franchising in oligopolistic markets with uncertain demand.' *International Journal of Industrial Organization* 9(3), 343–364
- Jullien, Bruno, and Patrick Rey (2007) 'Resale price maintenance and collusion.' *RAND Journal of Economics* 38(4), 983–1001
- Kahneman, David, Jack L. Knetsch, and Richard H. Thaler (1991) 'Anomalies: The endowment effect, loss aversion, and status quo bias.' *Journal of Economic Perspectives* 5(1), 193–206
- Kumar, Nirmalya (1996) 'The power of trust in manufacturer-retailer relationships.' *Harvard Business Review* 74(6), 92–106
- Laffont, Jean-Jacques, and Jean Tirole (1986) 'Using cost observation to regulate firms.' *Journal of Political Economy* 94(4), 614–641
- Lal, Rajiv, and Chakravarthi Narasimhan (1996) 'The inverse relationship between manufacturer and retailer margins: A theory.' *Marketing Science* 15(2), 132–151
- Lariviere, Martin A., and V. Padmanabhan (1997) 'Slotting allowances and new product introductions.' *Marketing Science* 16(2), 112–128
- Levin, Jonathan (2003) 'Relational incentive contracts.' *American Economic Review* 93(3), 835–857
- Loeb, Martin, and Wesley A. Magat (1979) 'A decentralized method for utility regulation.' *Journal of Law and Economics* 22(2), 399–404
- Mailath, George J. Volker Nocke, and Lucy White (2008) 'When the punishment must fit the crime: Remarks on the failure of simple penal codes in extensive-form games.' Mimeo, University of Oxford
- Martimort, David, and Salvatore Piccolo (2007) 'Resale price maintenance under asymmetric information.' *International Journal of Industrial Organization* 25(2), 315–339
- Mathewson, Frank, and Ralph Winter (1998) 'The law and economics of resale price maintenance.' *Review of Industrial Organization* 13(1–2), 57–84
- Nocke, Volker, and Lucy White (2007) 'Do vertical mergers facilitate upstream collusion?' *American Economic Review* 97(4), 1321–1339
- Puppe, Clemens, and Stephanie Rosenkranz (2006) 'Why suggest non-binding retail prices?' Working Papers 06-10, Utrecht School of Economics.

- Raith, Michael (1996) 'A general model of information sharing in oligopoly.' *Journal of Economic Theory* 71(1), 260–288
- Rey, Patrick, and Thibaud Vergé (2007) 'Economics of vertical restraints.' In *Handbook of Antitrust Economics*, ed. Paolo Buccirossi (Cambridge: MIT Press) pp. 353–390
- Rubinstein, Ariel, and Asher Wolinsky (1995) 'Remarks on infinitely repeated extensive-form games.' *Games and Economic Behavior* 9(1), 110–115
- Sorin, Sylvain (1995) 'A note on repeated extensive games.' *Games and Economic Behavior* 9(1), 116–123
- Spengler, Joseph J. (1950) 'Vertical integration and antitrust policy.' *Journal of Political Economy* 58(4), 347–352
- Steiner, Robert L. (1993) 'The inverse association between the margins of retailers and manufacturers.' *Review of Industrial Organization* 8(6), 717–740
- Thaler, Richard (1985) 'Mental accounting and consumer choice.' *Marketing Science* 4(3), 199–214
- Tirole, Jean (1988) *The Theory of Industrial Organization* (Cambridge, MA: MIT Press)
- Tversky, Amos, and David Kahneman (1991) 'Loss aversion in riskless choices: A reference-dependent model.' *Quarterly Journal of Economics* 106(4), 1039–1061
- Wen, Quan (2002) 'A folk theorem for repeated sequential games.' *Review of Economic Studies* 69(2), 493–512