UC Berkeley
Other Recent Work

Title
Optimal Contracts with Lock-In

Permalink
https://escholarship.org/uc/item/19f324hf

Authors
Farrell, Joseph
Shapiro, Carl

Publication Date
1987-10-12

Peer reviewed

UNIVERSITY OF CALIFORNIA, BERKELEY

Department of Economics

Berkeley, California 94720

OPTIMAL CONTRACTS WITH LOCK-IN

Joseph Farrell and Carl Shapiro

October 12, 1987

Abstract

We analyze long-term bilateral contracts when buyers incur
relationship-specific set-up costs and sellers choose some aspect
of product or service quality that is not verifiable to third
parties. If set-up costs are observable, the first-best outcome
can be achieved even though contracts cannot enforceably specify
quality; this does not even require long-term contracts. If set-up
costs are unobservable, however, then long-term price contracts
can strictly outperform short-term contracts, although they are
not constrained Pareto efficient. Equilibrium may involve either
inefficiently low quality with no buyer switching, or efficient
quality with inefficient switching. A policy of taxing switching,
or the availability of "budget-breaking" third parties, can improve
welfare.

Joseph Farrell, U.C. Berkeley
Carl Shapiro, Princeton University

Optimal Contracts with Lock-In

Joseph Farrell

Carl Shapiro

1. Introduction

In many markets, buyers and sellers invest in *relationship-specific assets*. A worker moves to take a new job; another invests in job-specific human capital. A purchasing agent establishes a working relationship with a particular supplier. Manufacturing employees learn to use pieces of equipment peculiar to a single upstream vendor. An author learns the idiosyncracies of a word processing program. "Software lock-in" arises when a computer user writes software or organizes her data in a way tailored to one operating system. Other examples of such *lock-in* through the sinking of set-up costs abound.

In such markets, if transactions must take place over time, a buyer realizes that selecting a particular seller puts her in that seller's power in the future. Even if competition is perfect *ex ante*, once a relationship is established there is some *ex post* bilateral monopoly, which, as Oliver Williamson (1985) has emphasized, can lead to problems of *opportunism*.

Opportunism is most clearly seen when a supplier who has built up a relationship with a customer chooses price or quality unilaterally. The supplier has an incentive to choose a higher price or lower quality than would be efficient. In the examples above, the worker worries about future wages, the manufacturer fears that its equipment supplier will set high fees for service and spare parts, the author is at risk that her word processing

1

system not be improved and made compatible with subsequent generations of computers and printers, etc. Once the buyer is locked in and can no longer readily switch to another supplier, the seller's *ex post* monopoly power can lead to inefficiency, in which case we call it opportunism. The inefficiencies associated with price gouging have been explored in the growing literature on repeated short-term interactions with lock-in. These papers assume that firms compete only in spot prices, with no commitments for the future. See Paul Klemperer (1987a,b) and ourselves (1987) and the references therein.

What institutional arrangements arise to avoid the inefficiencies associated with lock-in and opportunism? Long-term contracts are an obvious remedy, especially when buyers are large and sophisticated. With a long-term contract, the buyer protects herself from *ex post exploitation* by specifying seller performance prior to becoming locked in. *Complete* long-term contracts that could fully specify all relevant variables would obviously solve all opportunism problems and render lock-in unimportant. Clearly, however, such contracts are often impossible or prohibitively costly, and much scope is left for opportunism. In this paper, we study the efficacy of *incomplete* long-term contracts in mitigating the problems associated with lock-in. One of our central questions is whether realistically incomplete long-term contracts outperform short-term ones.

We study situations where long-term contracts are incomplete because some dimensions of performance by the seller are *noncontractible*. By this we mean that some aspects of performance are not *verifiable* to third parties, and thus cannot be included in a contract. We loosely refer to such variables as "quality." In the examples above, a court would find it difficult to determine whether the equipment supplier provides "good" service, whether the word processing company has shown good faith in retaining compatibility with new hardware, or whether the computer company has upgraded its system sufficiently to serve a growing customer's data processing needs. Custom or rush orders are also difficult to provide for in a long-term contract.

We begin with a general result that constitutes a warning about the dangers of *partial* long-term contracts. We consider bilateral situations in which *some* but not all variables can enforceably be agreed on in advance, and set at their efficient levels. For instance, although it is easy to specify prices, quality may be hard to define and enforceably specify

in advance. It might seem obvious that such partial protection is desirable: surely it is better to have one variable inefficiently set than two. But as we show below, this intuition is generally incorrect. Instead, we find that a second-best result applies: if the buyer cannot specify *all* variables, trying to specify an incomplete set of variables in the contract may easily be worse than specifying none at all.

The remainder of the paper develops a simple model of *ex ante* competition and *ex post* lock-in in the presence of noncontractible quality. Assuming that payoffs (in particular, the buyer's switching cost) are common knowledge, we show that short-term contracts are fully efficient, so long-term contracts cannot (strictly) help. Obviously, this is unrealistic, and we therefore add to our model the realistic assumption that buyers' switching costs are not observable to the seller. Then, without a long-term contract, the seller may be tempted *ex post* to raise his price or lower his quality so far that some buyers will inefficiently switch away: spot market interactions would not be fully efficient. Long-term price contracts, however, may be able to mitigate this problem.

A long-term contract that prevents the seller from raising the price *ex post* might tempt him to lower the quality instead. Because that would be *ex post* inefficient, such a long-term contract would be unattractive: buyers would prefer no long-term contract (or one that simply mimics the outcome without one). In other cases, however, signing a long-term price contract that prevents the seller from driving away low switch-cost buyers by raising price may induce him not to drive them away at all: a long-term price contract can be positively helpful.

The model we analyze, then, has sellers competing *ex ante* to serve buyers, who then become locked-in to an extent that differs across buyers and is unobservable to sellers. In each period, sellers choose prices if prices are not specified by long-term contracts, and choose qualities in any case. Each period each buyer either buys again from the seller who sold to her the previous period or else switches to another seller, in which case she must pay a new set-up cost. We show that long-term price contracts can sometimes help, sometimes not; and that although they can help, they are imperfect even compared to what an uninformed planner or regulator could achieve.

3

2. Beware Incomplete Contracts

Is an incomplete long-term contract better than none at all? Intuition might suggest that a partial contract provides some protection for a buyer about to become locked-in, but we show in this section that partial protection often is of negative value. We call this the *Principle of Negative Protection.*

Partial Protection with Separable Preferences

We begin by expounding the natural intuition that partial protection is better than none. We believe that this intuition stems from a special case where it is correct. As we shall see below, however, the intuition does not generalize beyond that case.

Suppose that there are two variables, x_1 and x_2, that matter to the seller, S, and the buyer, B. Each party's payoff depends separably on x_1 and x_2:

$$u^i(x_1, x_2) = u_1^i(x_1) + u_2^i(x_2), \quad i = B, S. \tag{1}$$

Suppose that x_1 can be specified in a contract, but that x_2 cannot, and must be chosen unilaterally by one person or the other, in such a way as to maximize his own utility. Finally, suppose that the unilateral choices of each variable are different from one another and from the efficient choices x_i^*.

It follows that (*i*) whatever value of x_1 is anticipated will not affect the choice of x_2; (*ii*) the welfare effects of different choices of x_1 are independent of the choice of x_2; and consequently (*iii*) it is desirable to sign a contract ensuring that $x_1 = x_1^*$.

The Principle of Negative Protection

When payoffs are not separable, specifying the level of one variable so as to avoid opportunism in that variable can adversely affect the choice of another variable. This effect can be powerful enough to spoil the efficiency gains from the more-efficient level of the first variable. As a consequence, partial protection may be worse than no protection at all.

An important context in which this principle of negative protection applies is that of competition with lock-in. Commonly, sellers compete to serve a buyer, with the seller

setting all of the variables *ex post*, subject to a constraint that the buyer get a certain "reservation" utility level (her utility from severing the relationship). In that case, the buyer does not benefit from any restrictions on the seller's ability to choose freely, while the seller surely cannot gain, and in general must lose. Another way to say this is that the seller will in any case exploit the buyer down to her reservation utility level, and if given a free hand will do so efficiently, while if constrained he will do so inefficiently. We formalize this in:

Proposition 1. *Consider a relationship between two agents, B and S. Suppose that S chooses a vector $(x_1, ..., x_n)$ of attributes of the relationship (prices, quantities, qualities, etc.). Suppose that any or all of $x_1, ..., x_m$ can be specified in a contract, where $m < n$, but that $x_{m+1}, ..., x_n$ cannot. Suppose further that S will set the variables he controls so as to give B an exogenous utility level. Then it is inefficient to specify any of the variables $x_1, ..., x_m$ in a contract, except trivially to "specify" them at the levels that S would choose if they were not specified.*

With *ex ante* competition, the buyer ultimately pays for the inefficiency caused by partial protection. Since the seller earns lower quasi-profits after the buyer is locked-in, he offers a less attractive introductory offer to the buyer at the outset of the relationship. We will see this principle in action in the model analyzed below.

3. A Model of Competition with Lock-In

We examine a market where buyers consume a good or service on a regular basis and there are relationship-specific investments. A good example to keep in mind is that of a buyer who establishes a relationship with a supplier and places regular orders. Another example is that of a customer who purchases a durable, learns how to operate it, and then requires service and spare parts on an ongoing basis. A third example is that of a buyer who installs a computer system and requires upgrades over time.

We model sellers competing *ex ante* to serve buyers, who then become locked-in. In each period, the seller chooses prices if prices are not specified by long-term contracts, and chooses quality in any case. Quality cannot be specified in advance, but is contemporaneously observable to the buyer. Each period, the buyer either purchases again from the

seller who sold to her the previous period or else switches to another seller, in which case she incurs a new set-up cost. This fact makes the set-up cost a relationship-specific sunk investment, which we denote by $s > 0$.[1] We refer to s interchangeably as a *set-up cost* and as a *switching cost*: the switching cost *from* one relationship is the set-up cost *with* the next. When s is common knowledge, it does not matter who nominally bears it, but for definiteness we assume that the buyer pays s.

We assume that a buyer inelastically requires one unit of the good each period. Extending our analysis to the case of elastic demands is easy, as we will show below, provided that two-part tariffs are available. A buyer who gets a unit of quality q at price p in a period enjoys a net benefit $q - p$ during that period, less s if she set up a new relationship.[2] Buyers discount benefits using the discount factor $\delta \in (0, 1)$. We assume for simplicity that switching from one seller to another, although costly, can be done instantaneously,[3] and that a buyer cannot be required to make any purchases before switching away from a seller.[4]

We assume that production takes place under constant returns to scale. This assumption allows us to consider each buyer in isolation. We assume that many identical sellers stand ready to serve each buyer; there is perfect *ex ante* competition.

Sellers maximize discounted profits using the discount factor δ. The seller chooses the quality q of his good or service each period.[5] Each seller can produce a unit of quality q at cost $c(q)$, where $c'(q) > 0$ and $c''(q) > 0$. The efficient quality is q^*, defined to maximize $q - c(q)$, so $c'(q^*) = 1$. We think of "quality" as capturing such decisions as whether the seller continues to improve his product, whether he puts his best employees on the buyer's account, how much customer service he provides, how he responds to unforeseen requests by the buyer, etc. It seems difficult for third parties to accurately assess the

---

[1] Set-up costs that are not sunk can be incorporated into the ordinary cost function $c(\cdot)$ defined below.

[2] Without loss of generality, we define the units of quality so that utility is linear in q.

[3] We report below in footnote 9 on the case in which it takes one period for a buyer to switch suppliers.

[4] This may seem unnatural, since quality is contemporaneously observable to the buyer: it is strange to think of signing a contract and never buying at all! In fact, however, it makes no qualitative difference to our results, and simplifies the algebra.

[5] The analysis is identical if instead sellers choose quality once-and-for-all (e.g., by making some non-depreciating investment) after the buyer incurs the set-up cost.

seller's performance along any of these dimensions, although the buyer will be able to do so as the relationship unfolds. Critically, therefore, we assume that quality is not verifiable to third parties, and hence not contractible, although it is observable to the buyer contemporaneously.[6]

4. Short-Term Contracts

Observable Switching Costs

In this section, we analyze short-term contracts, *i.e.*, spot-market relationships, as a preface to our analysis of long-term contracts below. We show that when the buyer's switching cost s is common knowledge, short-term contracts achieve first-best efficiency despite lock-in and despite the potential problem of quality. This result is closely related to work of Gary Becker (1962), Vincent Crawford (1983), Sanford Grossman and Oliver Hart (1986), Joseph Farrell and Nancy Gallini (1986) and Drew Fudenberg, Bengt Holmstrom and Paul Milgrom (1987), but it is instructive to see it in our particular model. We then allow the switching cost s to be unobservable to the seller, and show how this can lead to inefficiency.

To analyze behavior without long-term contracts, we must specify the "rules of the game". We assume that, in each period, the seller sets price and quality, and then the buyer chooses whether to buy or to switch. Clearly, other institutional arrangements are possible; for example, the buyer could set the price and the seller could decide whether to deliver or not. But our assumption seems to reflect standard practice in many markets. The absence of a long-term contract means that the seller chooses both p and q anew each period, and that either party can walk away with no responsibilities.

We do allow for the possibility that sellers will set different prices to new and old buyers: with large buyers, the seller plausibly can and will keep track of who has previously bought from him. The payment at the beginning of a relationship (when the buyer incurs s) is f, and the price each period is p. Equivalently, the price to new buyers is $p + f$ and the price to old buyers is p. We expect f to be negative.

---

[6] See Oliver Hart and John Moore (1985) for an extensive discussion of variables that are observable but not contractible.

Clearly, the seller will set quality at q^* each period, since he wants to produce efficiently any surplus he offers the buyer. So there is no quality-incentive problem with short-term contracts. Moreover, once the buyer is locked in, the seller never will set so high a price that she leaves, since *ex post* he can extract the rents from the relationship. So we have

Proposition 2. *When s is observable, repeated short-term contracts achieve first-best efficiency.*

Let us calculate the equilibrium prices. Define the first-best social surplus, gross of set-up costs, as

$$u^* \equiv \frac{q^* - c(q^*)}{1 - \delta}. \tag{2}$$

The buyer can get $u^* - s$ by leaving the current relationship.[7] The seller sets p as high as possible without losing the buyer. This indifference requires $(q^* - p)/(1 - \delta) = u^* - s$, or

$$p = c(q^*) + (1 - \delta)s. \tag{3}$$

The seller expects quasi-profits of $(p - c(q^*))/(1 - \delta) = s$ from the locked-in buyer. Sellers compete *ex ante* with *introductory offers* $(f < 0)$ for these quasi-profits. This drives profits to zero:

$$f + \frac{p - c(q^*)}{1 - \delta} = 0, \tag{4}$$

which reduces to $f = -s$, using equation (3). Since it is worth s to have a locked-in buyer with switching cost s, sellers compete by paying the buyer s to get her in the first place. In equilibrium the seller pays the buyer's set-up cost, in the form of an introductory offer,[8] and exploits her later to recoup this subsidy.[9]

In practice, an obvious difficulty with such an equilibrium is that s is not common knowledge. Introductory offers are therefore vulnerable to exploitation by buyers with low

---

[7] The surplus a buyer gets when she switches is generally endogenous, since she faces an endogenous contracting problem with her new partner. But, in an equilibrium supporting the first best, she gets exactly u^* (gross of set-up cost) from a new partner.

[8] This does not mean, of course, that the ultimate incidence of s is on the seller: since sellers make zero profits, the incidence is on the buyer.

[9] If it takes a buyer one period to find a new seller, the analysis goes through just as above, replacing s with $s + q^* - c(q^*)$.

switch costs, who will take the sweetener $-f$ and then switch away instead of providing the anticipated stream of quasi-profits. We next explore equilibrium in markets where buyers differ unobservably in their switching costs.

Unobservable Switch Costs

For many reasons, buyers' set-up costs may differ in ways that sellers cannot observe.[10] Often, switching costs are labor costs incurred by the buyer, e.g., retraining, software development, or opportunity costs that are hard for outsiders to measure. So, a buyer may well know her switching costs before entering a relationship, although the seller does not. An adverse selection problem arises.

We assume that there are two types of buyers, with set-up costs s_L and s_H, where $0 < s_L < s_H$. We denote by ϕ, $0 < \phi < 1$, the fraction of buyers in the population who have the high switching cost s_H.[11] We assume that a buyer's switching history is observable to the seller. Indeed, as we shall see, a buyer may choose to switch in order to prove her switch cost is low. This assumption seems to go naturally with our focus on large buyers: it is worth sellers' while to keep track of who has sold to whom and who has switched.

Consider the pricing strategy of a seller with a locked-in buyer.[12] One policy is to price low enough to retain the buyer even if she has a low switch cost. This policy guarantees the seller s_L (in present value) from this buyer. Alternatively, he could treat the buyer as having a high switch cost and set a higher price. This strategy extracts s_H from such buyers, but makes buyers with low switching costs depart. The expected payoff to this more aggressive strategy is ϕs_H. The optimal strategy, and hence the equilibrium outcome, depends upon the comparison of these two payoffs, s_L and ϕs_H.

If $\phi s_H \leq s_L$, then sellers are not tempted to exploit the s_H's fully, and the equilibrium outcome mimics that in a market where all buyers have low switching costs. Namely,

---

10   The analysis above is perfectly applicable if buyers differ in their switch costs, so long as those costs are observable, since sellers compete for clients on an individual basis. So, the key issue is observability, not homogeneity, of set-up costs.

11   Given our assumptions of constant returns to scale and risk-neutral sellers, it makes no difference whether there is a population with a mix of types or a single buyer of unknown type.

12   As above, the seller will set quality at q^* so as to produce efficiently any surplus he offers.

sellers set

$$\begin{cases} p = c(q^*) + (1-\delta)s_L, \\ f = -s_L, \end{cases}$$

and there is no switching. A seller who offered a more generous introductory offer would attract all the buyers but would lose money, while a seller who offered a smaller introductory offer would attract no buyers.

If, on the other hand, $\phi s_H > s_L$, then a seller with a representative customer base finds it profitable to drive away the s_L's and keep the s_H's to exploit more ruthlessly. In this case, equilibrium involves the low switch-cost buyers switching. In fact, anyone who switches is taken to be an s_L, and she therefore gets u^* once she has switched (as in the previous subsection). That being so, any buyer will switch if and only if she expects her current relationship to give her less than $u^* - s$ in present value.

A sellers who attracts both types of buyers but loses his s_L buyers sets a price p that makes the s_H's just indifferent to switching:

$$\frac{q^* - p}{1 - \delta} = u^* - s_H. \tag{5}$$

The zero-profit condition ties down the introductory subsidy of $-f$:

$$-f + \phi \frac{p - c(q^*)}{1 - \delta} = 0. \tag{6}$$

Solving for p and f using equations (5) and (6) yields

$$\begin{cases} p = c(q^*) + (1-\delta)s_H, \\ f = -\phi s_H. \end{cases}$$

These equations describe the unique equilibrium when $\phi s_H > s_L$.

In such a "switching equilibrium," each s_L is subsidized by the s_H's to the tune of the sweetener $-f = \phi s_H$. But an s_L buyer must incur her set-up cost in order to get this subsidy, so her net gain is $\phi s_H - s_L$. A buyer with high switching costs is herself indifferent to switching, so her total payoff is what she could get by switching, $(u^* - s_H) + (\phi s_H - s_H)$. We summarize these results in

Proposition 3. *With unobservable s and no long-term price contracts, quality choice is efficient. The equilibrium is the first best if $\phi s_H \leq s_L$, but if $\phi s_H > s_L$ then equilibrium involves the s_L's switching. In such a switching equilibrium, the buyers' payoffs are $(u^* - s_L) + (\phi s_H - s_L)$ to the s_L's and $(u^* - s_H) + (\phi s_H - s_H)$ to the s_H's.*

With unobservable s, equilibrium involves some costly and inefficient switching if and only if a seller prefers to exploit just the s_H's, rather than retaining all buyers and earning s_L. It is sellers' freedom to set prices *ex post*, along with the temptation to drive away the s_L's while exploiting just the s_H's that causes inefficiency. Efficiency is likely to be disrupted if there are many buyer with high switch costs (ϕ large), if the low switch-cost buyers find it very easy to switch (s_L low), or if the heterogeneity in switching costs is substantial ($s_H - s_L$ large).

5. Long Term Price Contracts with Observable Switching Costs

If quality were not an issue, or if it were contractible, then long-term contracts would obviously achieve the first best: simply specify price at $c(q^*)$ and quality at q^* in all periods with a large penalty for any quality reduction by the seller. The unobservability of s would not disrupt the efficiency of this contract.

But, as we have argued, there are often "quality" variables that are not contractible. So far, we have had no occasion to consider qualities q other than q^*, since any exploitation by the seller is more efficiently (and so more profitably) done in the form of price than of quality. This changes when we consider long-term price contracts.

Obviously, long-term contracts can mimic the equilibrium with short-term contracts (that is, contracts can promise the prices that would rule without long-term contracts). Since short-term contracts achieve the first best when s is observable, efficiency can be achieved *a fortiori* with long-term contracts. Furthermore, since no alternative contract can give buyers more utility than in the first best without losing money, the efficient allocation is the unique equilibrium when s is observable.

In this section, we discuss some alternative efficient long-term price contracts. Then in the next section we ask what happens when s is not observable. We show that efficiency

11

can sometimes be achieved even when $\phi s_H > s_L$. In this sense, long-term contracts are valuable.

What can long-term contracts specify when quality is not contractible? We assume that the party enforcing the contract, the "court," can only observe whether a transaction takes place in a given period and any payments made; it cannot observe the quality. Therefore, a contract can only specify the price paid if a transaction occurs, p, the price paid if no transaction occurs, d, and the up-front payment upon signing of the contract, f.[13]

By the revelation principle, we can think of sellers as offering contracts of the form (f, p, d, q), where f, p, and d are explicitly in the contract and q is a claim regarding the quality that will be supplied. In order to be acceptable, the quality claim must be *credible*, *i.e.*, once the buyer accepts the contract and bears her set-up cost, it must be in the seller's interest to supply the claimed quality.

For a contract promising q^* to be credible, it must be the case that the buyer is *just* willing to stay if the seller produces q^*: otherwise, the seller would (at least slightly) reduce q. This *credibility* or *ex post exploitation* condition is:

$$\frac{q^* - p}{1 - \delta} = u^* - s - d. \tag{7}$$

In equation (7), the left-hand side gives a locked-in buyer's (prospective) lifetime utility from staying and the right-hand side gives her utility from departing: she must pay the departure fee d and switching cost s, and can then contract with a new partner to enjoy surplus of u^*.

Substituting for u^*, we can rewrite (7) as

$$p = c(q^*) + (1 - \delta)(s + d). \tag{8}$$

In present value, the seller can exploit the buyer just to the extent of $s + d$, the buyer's full cost of leaving. Notice that, although $p > c(q^*)$, sellers cannot compete by lowering p,

---

[13] By convention, we measure all payments as going from buyer to seller. Given the stationarity of the problem, it is sufficient to look at contracts that make the prices d and p independent of the length of the relationship.

since buyers know that a seller who sold at a price below (8) would provide quality below q^*.[14]

Ex ante competition drives profits to zero, so we have

$$f + \frac{p - c(q*)}{1 - \delta} = 0. \tag{9}$$

In view of (8), (9) becomes

$$f + d = -s. \tag{10}$$

Since $f + d$ is the total payment from buyer to seller in the event of an immediate divorce, equation (10) states that the seller must pay the buyer's switching cost if the relationship is broken off.

Clearly, equations (8) and (10) are necessary conditions for an equilibrium supporting the first best. They also are sufficient. By construction, buyers have no incentive to switch, and competing sellers clearly cannot profitably offer buyers more utility. We must also check that the seller will indeed supply quality q^*. Equation (7) tells us that any reduction in quality will make buyers leave, so we need only check that driving away buyers is not profitable. *Ex post*, the seller collects the rents in the relationship, relative to switching and forming new relationships; he will not lower quality and lose these quasi-profits.[15] So any triple (f, p, d) satisfying equations (8) and (10) does indeed support first-best efficiency.

Two principles characterize optimal contracts. First, from equation (10), the seller must have paid the buyer's switching cost in the event of a divorce. Whether this payment consisted of an up-front introductory offer $(f < 0)$, or of a payment upon departure $(d < 0)$, or some mix, does not matter, so long as the total payment covers s. Second, the parties realize that higher severance payments d make switching less attractive to the buyer, make *ex post* exploitation greater, and thus must go along with a higher contract price p if quality incentives are to be maintained. Such is the content of equation (8).

Equations (8) and (10) give two equations in the three contract variables, p, f, and d. The contract parameters are underdetermined, reflecting the possibility of borrowing and lending between the buyer and the seller.

---

[14] Observe the similarity between equation (8) and markup equations in the literature on quality-guaranteeing prices, e.g., Shapiro (1983).

[15] The profits earned from driving away his customers are $f + d$, which by equation (10) are negative.

It is useful to look at two particular long-term pricing contracts that support the first best. First, consider the contract that mimics the short-term contract or no-contract outcome. Setting $d = 0$, equations (8) and (10) give us

$$\begin{cases} p = c(q^*) + s(1 - \delta), \\ f = -s, \\ d = 0. \end{cases}$$

Under this contract, the seller pays the buyer's set-up costs when the relationship begins, knowing that he will recoup this investment in the form of subsequent markups over cost.

A second efficient contract has no up-front payment. Setting $f = 0$ in equations (8) and (10) gives

$$\begin{cases} p = c(q^*) \\ d = -s, \\ f = 0. \end{cases}$$

The product is offered at cost, with the seller agreeing to pay the buyer's switching cost if she ever leaves. With this contract, it is transparent how the departure fee internalizes the social costs that the seller might impose by reducing quality and inducing switching. To put it another way, the buyer faces no net switching cost, and so she will leave in response to any shortfall in quality. This contract is equivalent to the previous one plus a loan of s from the buyer to the seller, with interest on the loan paid in the form of lower prices and default on the loan payments secured by the departure fee of $-s$.

The principle of negative protection warns us that a long-term contract specifying price $p = c(q^*)$ without any departure fee *cannot* support the first best. Such a contract would make the seller set quality below the efficient level, and as a result (despite the lower prices) the buyer would be worse off.

Dealing with Elastic Demand

In our analysis above, we assumed inelastic demand. For instance, we showed that an introductory price of $c(q^*) - \delta s$ followed by prices of $c(q^*) + (1 - \delta)s$ yields efficiency. It might seem that elastic demand would drastically change this result.

14

If nonlinear pricing is available, efficiency is preserved. The seller need only make a take-it-or-leave-it offer of the efficient quantity, and the elasticity of demand and returns to scale are irrelevant. So the question only becomes interesting when pricing must be uniform. In that case, spot markets do not achieve efficiency, for *ex post* exploitation inevitably involves distortions. The buyer buys too much in the first period when the attractive introductory price is offered, and too little subsequently.

It may be tempting to try to level out the prices by a long-term price contract, but as the principle of negative protection warns, that would be a mistake. Price and quality would then no longer be set *ex post* so as to maximize the seller's payoff subject to the buyer's, and the buyer's *ex post* payoff would not change. Hence, the outcome would be less efficient.

Long-term contracts can achieve efficiency with elastic demand and linear pricing, but they are no longer underdetermined for the purpose. Instead, it is necessary to set p equal to marginal cost (at the optimum). Then d is determined by the need to make the buyer indifferent to switching in the optimum, and the up-front payment f is tied down by the zero-profit condition. With constant returns, our contract $p = c(q^*)$, $f = 0$, $d = -s$ does the job.

6. Long Term Contracts with Unobservable Switching Costs

Now we ask what long-term contracts can achieve when switching costs are unobservable and $\phi s_H > s_L$ so spot-market relationships do not achieve efficiency. We show that the outcome without long-term contracts, the switching equilibrium, remains an equilibrium when such contracts are available. For some parameter values, however, even though $\phi s_H > s_L$, the first-best outcome may be an equilibrium with long-term price contracts. In other words, introducing long-term price contracts expands the set of parameter values for which the first best is an equilibrium. Moreover, even when the first-best is not an equilibrium, there may be another equilibrium in which there is no switching but in which the quality supplied is some $\hat{q} < q^*$. This equilibrium may be better than the switching equilibrium in terms of overall buyers' surplus.

We allow for the possibility that firms offer buyers different contracts depending upon their previous switching history. In particular, firms infer, as above, that buyers who have

switched are of type s_L, and offer contracts designed for these buyers. As we shall show, we need consider only two submarkets: one for buyers who have never switched, which we call the *virgin market*, and another for buyers who have switched at least once, which we call the *aftermarket*.

A *contract equilibrium* consists of the following. Sellers compete for buyers by offering contracts (f, p, d, q) to buyers contingent on the buyers' switching histories. Each buyer chooses the credible contract offering her the largest payoff, and keeps switching from one contract to another as may be optimal, taking account of the contracts that will be offered as a function of her switching history. In equilibrium, there exists no credible contract that would, given the presence and credibility of existing equilibrium contracts, attract any buyers and earn positive profits.

One subtle aspect of the analysis is that a contract's credibility (as well as its profitability) may depend upon the mix of buyers expected to accept it, *i.e.*, which of the two types of buyers will choose this contract. We assume that both buyers and sellers have rational expectations regarding the mix accepting any equilibrium contract. We further assume that, in evaluating the credibility of deviant contracts, both buyers and sellers assume that such a contract is credible if it would be credible, were it to attract exactly those buyers who would take it if they judged it credible. We do not contemplate the possible loss of credibility of equilibrium contracts (or their withdrawal) on account of the presence of deviant ones: in that sense, we give equilibrium contracts an advantage.

Conditions for Efficiency

We ask first when the first best can be supported as an equilibrium. Clearly, $\phi s_H \leq s_L$ is a sufficient condition. We now show that the first-best outcome is supported by long-term contracts under a weaker condition than this.

Consider a long-term price contract that would be efficient if all buyers had switch cost s_L; for definiteness, consider the contract that sets $p = c(q^*)$, $f = 0$, and $d = -s_L$. This contract is credible if the seller prefers to set q^* and retain all the buyers, rather than to reduce q below q^* to save costs, but thereby lose the s_L buyers.

If the seller could do this exploitation through price, then he would choose to do so precisely when $\phi s_H > s_L$. But if the contract prohibits him from price gouging, he may

refrain from quality gouging because his cost savings are less than the price increase would be – precisely because q^* is the efficient quality.

How far could the seller cut quality and still keep his s_H customers? To retain a buyer with switching cost s the seller can cut quality to the level q given by

$$\frac{q - p}{1 - \delta} = u^* - s - d,$$

assuming that after switching a buyer can get $u^* - s$. Substituting $d = -s_L$ and $p = c(q^*)$ from the first-best contract, this equation tells us that to make s_H indifferent the seller can reduce quality to $q^* - (1 - \delta)(s_H - s_L)$. Defining

$$\Delta q \equiv (1 - \delta)(s_H - s_L),$$

a seller exploiting just the s_H buyers would produce quality $q^* - \Delta q$.

Such a quality reduction *ex post* would generate quasi-profits of $\phi(p - c(q^* - \Delta q))/(1 - \delta)$ from the s_H buyers, less the departure fees $(1 - \phi)s_L$ paid to the s_L buyers who leave. In contrast, maintaining quality at q^* gives *ex post* quasi-profits of zero. Substituting $p = c(q^*)$, we see that quality q^* is credible, and the first-best outcome is an equilibrium, only if

$$\phi(c(q^*) - c(q^* - \Delta q)) \leq (1 - \delta)(1 - \phi)s_L. \tag{11}$$

Since $c(\cdot)$ is convex, (11) is weaker than $\phi s_H \leq s_L$.[16]

To show that (11) is actually sufficient for the first-best to be an equilibrium with long-term contracts, we need only check that alternative contracts do not break the equilibrium. But this is trivial: any deviant contract must offer more than u^* to all buyers who accept it, and clearly no contract can profitably do so. This gives us:

Proposition 4. *With long-term price contracts, the first best is an equilibrium if (11) holds. This condition is strictly weaker than $\phi s_H \leq s_L$, the condition for the first best to emerge under short-term (spot-market) contracts. As $c(\cdot)$ becomes more convex, the*

---

[16] Using the intermediate value theorem, (11) can be rewritten as $\phi s_H - s_L \leq s_L(1 - \phi)(k - 1)$, where $k = 1/\{1 - \frac{1}{2}\Delta q c''(q)\}$ for some $q \in (q^* - \Delta q, q^*)$. Since $c''(q) > 0$, $k > 1$, and this condition is weaker than $\phi s_H - s_L \leq 0$.

disparity grows: that is, the advantage of long-term contracts becomes greater the more efficiency gain there is from enforcing optimal quality choice.

Proposition 4 tells us that long-term contracts are especially useful when it is important to maintain quality near its efficient level – even though quality is not contractible! The reason is that "cheating" the buyers, by driving away the s_L buyers and exploiting the s_H buyers as much as possible, is less attractive to the seller if he must do it inefficiently than if he can do it efficiently. In the Principle of Negative Protection, we pointed out that *given* the extent to which the seller is going to exploit the buyer, it pays to let him do so efficiently. Here, however, we see that a contract that prevents the seller from exploiting buyers efficiently may deter him from doing so at all. The general principle is that *whatever exploitation the seller does in equilibrium, should be done as efficiently as possible; but whatever exploitation he is not meant to do in equilibrium, should be made as inefficient as possible*, since he is then most easily deterred from doing it.

Equilibrium if Efficiency is Not Attainable

When (11) fails, no contract designed to keep all buyers and to enforce quality q^* is credible: a seller prefers to reduce q and to drive away the s_L's. Two possible contracting responses to this problem come to mind, and it turns out that each can occur in equilibrium, and that no other equilibrium can exist.

First, we could accept that there will be switching in equilibrium; (11) is especially likely to fail if s_L is small, so the overall efficiency loss from having the s_L buyers switch is not too great. If we do that, then obviously it is better that they should be driven away not by low quality (which has efficiency costs for those who stay) but by high prices. So we might design a contract that will pay the low switch-cost buyers to leave, thus leaving the high switch-cost buyers, who then get efficient service. Obviously they would prefer not to have to pay the others to leave, but they may have to do so. This arrangement duplicates the switching equilibrium analyzed above; it always exists when $\phi s_H > s_L$.

A second response to the failure of (11) is to notice that it fails because the cost savings from reducing quality by Δq below q^* are too large, relative to the loss of customers. By planning for some lower quality $\hat{q} < q^*$ (and pricing accordingly), we could reduce the

prospective cost savings from a quality reduction of Δq, since $c(\cdot)$ is convex, and thus induce the seller *not* to reduce quality and drive away the s_L's. This arrangement is sometimes a second equilibrium, which we call the $\hat{q}$ *equilibrium*. Sometimes – but not always – it gives a higher aggregate payoff than the switching equilibrium.

7. Formal Analysis of Long-Term Contracts and Unobservable Switch Costs

In this section, we complete our analysis of long-term contracts when s is private information. The analysis rests on the following Proposition, proven in the Appendix:

Proposition 5. *Only two types of contract equilibria are possible. (1) Switching Equilibria: All buyers in the virgin market adopt the same contract; the s_L buyers switch away from this contract into the aftermarket, while the s_H buyers do not switch. The aftermarket contract is the efficient contract that would arise in a market with only s_L buyers. Quality is efficient. (2) No-Switching Equilibria: All buyers select the same contract and none switch from it. The quality supplied may or may not be q^*.*

Using Proposition 5, we now characterize further each type of equilibrium and determine when each exists.

Switching Equilibrium

We first find the prices and payoffs that must prevail in any switching equilibrium. We already know that the aftermarket contracts are exactly those from equations (8) and (10) with $s = s_L$. A buyer earns $u^* - s$ by switching into the aftermarket.

The virgin-market contracts are designed to (just) retain the s_H buyers and to lose the s_L buyers. By Lemma 1 in the Appendix, sellers provide quality q^* in these contracts, so as to produce surplus efficiently for the s_H buyers. Choosing $d = 0$ as a convenient normalization, the *ex post* exploitation and zero-profit equations are exactly equations (5) and (6) above. The equilibrium contract is therefore

$$\begin{cases} p = c(q^*) + s_H(1 - \delta), \\ f = -\phi s_H, \\ d = 0. \end{cases}$$

19

These necessary conditions show that this equilibrium exactly mimics the switching equilibrium without long-term contracts. In the Appendix, we derive some more necessary (and jointly sufficient) conditions for the contracts described actually to be an equilibrium, thereby proving

Proposition 6. *A switching equilibrium is an equilibrium with long-term contracts if and only if $\phi s_H \geq s_L$. Whenever a switching equilibrium exists it is identical to the equilibrium without long-term contracts.*

Proposition 6 might suggest that when $\phi s_H > s_L$ there is no gain from long-term contracts. That is not so, for two reasons. First, as we have seen, (11) may hold even though $\phi s_H > s_L$; $\phi s_H \leq s_L$ is sufficient but not necessary for the first best to be sustainable with long-term contracts. And second, even if (11) fails, the switching equilibrium need not be the unique equilibrium. Although in that case the first-best is not an equilibrium, there may be another equilibrium without switching. We now explore such equilibria.

No-Switching Equilibrium

In an equilibrium without switching, all buyers sign the same contract; see the proof of Lemma 2 in the Appendix. Our task here is to describe the equilibrium and to find the conditions under which it exists. Of course, we have already done so for the case in which sellers provide efficient quality q^*; we now examine the more general problem, looking for an equilibrium in which some quality $\hat{q}$ not necessarily equal to q^* is provided and nobody switches.

The *ex post* exploitation and zero profit conditions are by now familiar, except that quality need not be set efficiently:

$$\frac{q-p}{1-\delta} = u^* - s_L - d, \tag{12}$$

reflecting the fact that low switch-cost buyers are on the margin of leaving, and

$$f + \frac{p - c(q)}{1-\delta} = 0. \tag{13}$$

Choosing $f = 0$ without loss of generality, and calling $u \equiv \dfrac{q - c(q)}{1 - \delta}$, we have as necessary conditions

$$
\begin{cases}
p = c(q), \\
d = u^* - u - s_L, \\
f = 0.
\end{cases}
$$

Credibility requires that the seller not want to lower quality in order to exploit just the s_H buyers, abandoning the s_L buyers. To do so, a seller would lower quality by Δq. The credibility condition is therefore

$$
\frac{p - c(q)}{1 - \delta} \geq \phi \frac{p - c(q - \Delta q)}{1 - \delta} + (1 - \phi)d.
$$

Substituting for p and d this can be rewritten as

$$
\phi(c(q) - c(q - \Delta q)) \leq (1 - \delta)(1 - \phi)(s_L + u - u^*). \tag{14}
$$

It is perfectly possible that q^* satisfies inequality (14). In other words, we may have

$$
\phi(c(q^*) - c(q^* - \Delta q)) \leq (1 - \delta)(1 - \phi)s_L. \tag{15}
$$

If so, we are back to the first-best equilibrium discussed above; (15) is exactly (11), the necessary and sufficient condition for the first best to be an equilibrium.

If inequality (11) fails, the presence of some high switch cost buyers destroys the credibility of supplying quality q^* to all buyers. A lower quality will still be credible, provided that we can make $c'(q)$ as small as we like by reducing q. We define $\hat{q}$ as the largest q that satisfies equation (14), i.e.,

$$
\phi(c(\hat{q}) - c(\hat{q} - \Delta q) \equiv (1 - \delta)(1 - \phi)(s_L + \hat{u} - u^*), \tag{16}
$$

where $\hat{u} \equiv (\hat{q} - c(\hat{q}))/(1 - \delta)$. Below, we refer to the no-switching equilibrium quality as $\hat{q}$, using the convention that $\hat{q} \equiv q^*$ if (11) holds.

Now the no-switching equilibrium is easy to describe. Sellers offer the following contract:

$$
\begin{cases}
p = c(\hat{q}), \\
d = u^* - \hat{u} - s_L, \\
f = 0.
\end{cases}
$$

All buyers accept this contract and never switch. Each buyer gets $\hat{u}$ less her original set-up cost.

When does this contract in fact support a no-switching equilibrium? We have checked the credibility of $\hat{q}$. In the Appendix we check that there are no credible deviant contracts that could be offered by others and upset the equilibrium, proving (see Table 1)

Proposition 7. *The unique equilibrium without long-term price contracts – the first-best outcome when $\phi s_H \leq s_L$ and the switching equilibrium otherwise – remains an equilibrium when we allow long-term contracts. Indeed, the first-best remains the unique equilibrium when $\phi s_H < s_L$. But the first-best is an equilibrium if and only if the weaker condition (11) holds. If (11) holds and $\phi s_H > s_L$, then we have two equilibria: the first-best and the switching equilibrium. The s_L buyers prefer the switching equilibrium, and the s_H buyers prefer the first-best equilibrium. If $\phi s_H > s_L$ and (11) fails, a no-switching equilibrium with quality $\hat{q} < q^*$ exists if and only if $\hat{u} \geq u^* - s_H(1 - \phi)$; quality $\hat{q}$ is defined by equation (16). The s_L buyers prefer the switching equilibrium, but the s_H buyers prefer the $\hat{q}$ equilibrium if it exists. The overall payoff comparison can go either way. If (11) fails and there is no $\hat{q}$ equilibrium, then the switching equilibrium is unique.*

One surprising feature of this result is that equilibrium is not generally unique: the switching equilibrium can coexist with either the first-best or the $\hat{q}$ equilibrium. Why? When the switching equilibrium exists, $\phi s_H > s_L$, which means that the s_L buyers actually do better than in the first-best: they take the introductory offer of ϕs_H and switch away, thus getting some cash plus their efficient aftermarket contract. So no other contract could tempt them away. One might expect that a deviant contract could readily tempt away the s_H buyers, who are subsidizing the s_L's, and this would be true if it were not for the problem of credibility. Any contract designed to give the s_H buyers more than they get in the switching equilibrium would also be attractive to the s_L buyers, who would exploit it just as they do the equilibrium virgin-market contract.

At the same time, if the first-best or $\hat{q}$ equilibrium is established, there is no problem with deviant contracts trying to establish a switching equilibrium. For such contracts cannot attract the s_H buyers to be exploited, and are not profitable if they attract only the s_L buyers.

It is our assumption that buyers do not question the credibility of existing equilibrium contracts when deviant contracts appear that leads to multiple equilibria.[17] If, on the contrary, buyers believed in newly-offered contracts and re-evaluated the credibility of previous contracts, then in some cases no equilibrium would exist.

8. Constrained Efficiency

When (11) fails, equilibrium is not first-best efficient. A natural question is whether there is scope for intervention in the market as a result. Obviously, a regulator who could observe q or s would be able to enforce the first best. A more interesting question is whether intervention can improve things without such knowledge. The answer is yes, and in a somewhat surprising way.

Because the inefficiency is caused by switching costs, one might expect that subsidizing switching could help. But, as condition (11) shows, it would not. Subsidizing switching does not affect $\Delta q \equiv (1-\delta)(s_H - s_L)$, but reduces s_L, so it makes (11) *less* likely to hold. As we have seen, the reason for inefficiency is not switching costs, but their unobservability. While nothing can be done about the *absolute* unpredictability Δs of switch costs, its *relative* importance $(s_H - s_L)/s_L$ can be reduced, not by subsidizing switching, but by *taxing* it!

Not only (11) but also $\phi s_H \leq s_L$ is made to hold if we impose a large tax on switching. Of course, a larger tax is needed to make this latter condition rather than the weaker (11) hold. Thus we have:

Proposition 8. *A sufficiently large tax on switching will make the first-best the unique equilibrium with or without long-term price contracts. The tax required is smaller if such contracts are available.*

The first best also can be supported if a trustworthy third party is available. A third party allows the contracting parties to sever the link between what the buyer is paid when

---

[17] Non-uniqueness persists even if we allow for deviant *menus* of contracts. To see this, observe first that (by what we have just shown) a successful deviant menu must attract *both* buyer types. Now, if it would induce switching by any buyers, it would surely induce switching by the s_L buyers, and then it could not profitably attract both types away from the switching equilibrium. On the other hand, if it did not induce switching, then it would have to give *all* buyers at least $u^* + \phi s_H - s_L > u^*$, and so could not be profitable.

he switches and what the seller pays. The problem was that sometimes the payment $-d$ must be so large, in order to deter exploitation, that the buyer will switch just to collect it even though the quality is indeed q^*; this is the switching equilibrium. If in the event of a divorce the buyer receives s_L but the seller pays much more (the difference going to the third party), then the efficient outcome can be supported.[18]

Interestingly, the presence of honest third parties is not quite as good as the government's power of taxation. Although reliable third parties permit the first best to be supported as an equilibrium, they cannot always make this the *unique* equilibrium. If the switching equilibrium exists in their absence, it continues to exist when third parties are available, since the s_L's will not be attracted from the switching equilibrium by a payoff of u^*.

9. Conclusion

We have explored the implications of a particular incompleteness in long-term contracting – noncontractible quality - in markets where buyers and sellers make relationship-specific investments. We have identified a number of points that should carry over from our simple model to more general ones. First, the simplest model suggesting that short-term contracts are adequate is misleading: even so simple a modification as introducing two buyer types changes it dramatically. This is not surprising. Buyers and sellers clearly do worry about lock-in and consider long-term contracts important in the presence of significant relationship-specific assets.

Second, we identified the *principle of negative protection*, which states that if the seller will in any case exploit the buyer *ex post*, then it pays the buyer to let him do so efficiently, as he will do if his actions are unconstrained. If some but not all variables can be specified in a long-term contract, then it provides *negative* protection for the buyer to insist on specifying them. However, there are benefits from making the seller do his exploitation inefficiently, if it means that he will then refrain altogether. This is one thing a long-term partie' (in our model, price) contract can do. Thus long-term partial contracts can make the first-best an equilibrium when it was not before.

---

[18] As is typical with such schemes, the entire arrangement is susceptible to collusion between the buyer and the third party against the seller.

Interestingly, even though quality is not contractible, long-term contracts offer the greatest advantage when it is important that quality be kept near its efficient level. If efficiency falls off rapidly as quality departs from its efficient level, then a long-term price contract may deter the seller from some *ex post* exploitation and be of positive value. But it is exactly in such circumstances that an improperly-designed long-term pricing contract runs the greatest danger of inducing the seller to inefficiently lower quality rather than raise price to exploit locked-in buyers.

Appendix

Here we prove Propositions 5, 6, and 7. First we establish two preliminary results:

Lemma 1. *Any equilibrium contract selected by only a single type of buyer must call for quality q^*.*

Proof. Suppose that such a contract calls for quality $q \neq q^*$. Clearly a seller could design a deviant contract with $q = q^*$ offering buyers no less surplus, and earning positive profits, if such a contract is credible and still attracts only the specified type of buyers. Credibility can always be maintained by altering p so that equation (8) still holds. What about not attracting the other group of buyers (this may be important either to preserve positive profits or to preserve credibility)? If that group would not switch, were it to sign the contract, those buyers do not care about q or p (in view of *ex post* exploitation) and will find the deviant contract unattractive, just as it found the original one (by supposition). And if that group would switch, it pays $f + d$, so the deviant need only make sure that this sum does not fall. ∎

Lemma 2. *Both types of buyers select the same contract in the virgin market.*

Proof. First we prove that there cannot be self-selection if there is no switching in equilibrium. Given that neither type switches, the two types of buyers make exactly the same comparison between any pair of contracts. So both types must be indifferent between the two equilibrium contracts. By Lemma 1 these contracts must call for the same quality. Therefore, they must be identical (except possibility for a different mix between f and d): since sellers earn zero profits from either contract and have the same stream of costs, the present value of the payments made by buyers in either contract is the same.

Second, consider self-selection with some switching. Note that any contract that attracts only one type and induces switching is a null contract: everyone pays $f + d$ and no goods are produced, whence we know that $f + d$ must equal zero for that contract. Our candidate equilibrium therefore entails one group switching from the null contract simply to show that they have incurred their s in order to enter the aftermarket. Clearly, if the high switch cost buyers prefer doing this to taking the alternative contract, so do the other buyers, since the low switch cost buyers earn more from the switching/aftermarket strategy

than do the s_H types. We conclude that it must be the s_L buyers who switch from the null contract. Now consider what the contract for the s_H buyers must look like. In order not to attract the s_L buyers, we must have $f + d \geq 0$. But $f + d$ must also be nonpositive, or else firm would make money by cheating the s_H buyers. So $f + d = 0$ for the s_H contract. As usual, we can use our degree of freedom to set $f = 0$, from which we have $d = 0$ as well. Then zero profits tells us that $p = c(q^*)$, which is not credible. And $p > c(q^*)$ cannot be an equilibrium contract, since a deviant contract could set $f \in (-\phi(p - c(q^*)), 0)$, attract both types and earn positive profits. ∎

Proof of Proposition 5

If no switching occurs, then the virgin market is the whole market. By Lemma 2, all buyers purchase the same contract in the virgin market. The only buyers on the margin for leaving must be those with low switch costs, so the dormant aftermarket contains contracts designed for them. This is the No-Switching Equilibrium.

Suppose instead that switching does occur in equilibrium. We now establish that it must be the low switch cost buyers only who switch. Everyone cannot switch, since the virgin market contracts would then be null ones and a seller would have the incentive to offer the aftermarket contract in the virgin market, saving buyers their switch cost associated with signing the null contract in the virgin market. Nor could the high switch cost buyers switch. Since everyone takes the same virgin market contract (Lemma 2), and the s_H buyers earn a lower payoff from switching than do the s_L buyers, if the former groups wants to switch, so must the latter. If the s_L buyers do switch, the aftermarket is just for them, so the analysis from the previous section applies, with aftermarket contracts being first best for s_L buyers. The s_H buyers stay with the virgin market contract and the s_L buyers incur s_L to enter the aftermarket. This describes exactly the Switching Equilibrium. ∎

Proof of Proposition 6

The text establishes necessary conditions for equilibrium contracts in a switching equilibrium. Here, we derive further necessary and jointly sufficient conditions for these contracts to constitute an equilibrium.

First, we consider deviant virgin-market contracts. We begin with deviant contracts designed to attract only one type of buyer. If such a contract induced switching, it would be a null contract, and could not profitably attract either group, since a buyer would do better to exploit the equilibrium virgin-market contract for ϕs_H. This leaves the possibility of attracting a single group and keeping that group. It is not possible to attract just the low switch-cost buyers this way, since the deviant contract must offer at least $u^* + \phi s_H - 2s_L$ to attract the s_L buyers and thus must attract the s_H buyers as well.

Nor is it possible to attract just the s_H buyers, even though they are subsidizing the others in equilibrium. Avoiding attracting the s_L buyers requires that $f + d \geq -\phi s_H$, whereas credibly attracting the s_H buyers requires $f + d \leq -\phi s_H$. (We can restrict attention to deviant contracts offering quality q^*.) Together, these conditions constrain the deviant to offer a contract equivalent to the equilibrium virgin-market contract. Intuition suggests that it might be possible to keep the s_L buyers away by offering less of a sweetener while lowering p to attract the s_H buyers. But lowering p requires a lower d if credibility is to be maintained, and this will again attract the unwelcome s_L buyers.

Most important, consider deviant contracts designed to attract both types of buyers. Any such contract that then loses both types is a null contract, and cannot profitably attract any buyer. And if it attracts both types and loses just one, it must be the s_L buyers. But this is just the equilibrium contract, which is no threat to the equilibrium. This leaves only the possibility of a deviant contract designed to attract both types of buyers and induce no switching.

What is the most attractive credible contract, *given* that both types of buyers select it? This is the contract promising quality $\hat{q}$ as defined by equation (16). The $\hat{q}$ contract blocks the switching equilibrium if and only if it indeed attracts both types, *i.e.*, if and only if $\hat{u} > u^* + \phi s_H - s_L$. (The s_L buyers are the harder ones to attract with a contract that gives both types the same payoff.) This inequality gives us the key necessary condition for the switching equilibrium:

$$\hat{u} - u^* \leq \phi s_H - s_L. \tag{17}$$

Notice that if $\phi s_H > s_L$ (17) is automatic.

It remains to check is that the virgin-market contract is credible. As usual, the seller will not choose to lower quality and lose all buyers, since he can exploit remaining buyers to the extent of their switching costs, in this case s_H. But we must check that the seller will not *increase* quality in order to retain both types of buyers. By how much would he have to raise quality? Since he is just keeping the s_H buyers in equilibrium, to just keep the s_L buyers he would have to raise quality by Δq. By this deviation, he would earn $f + (p - c(q^* + \Delta q))/(1 - \delta)$. These profits must be no larger than his equilibrium profits of zero. Making this comparison gives us another necessary condition for a switching equilibrium,

$$c(q^* + \Delta q) - c(q^*) \geq (1 - \phi)(1 - \delta)s_H. \tag{18}$$

Using the intermediate value theorem again, it is not hard to show, using the convexity of $c(\cdot)$, that (18) holds whenever $\phi s_H > s_L$.

In summary, any switching equilibrium involves contracts equivalent to those described by equations (5) and (6). It is the same outcome as arises without long-term contracts when $\phi s_H > s_L$. Inefficiency arises not from quality choice, but rather from the costly switching behavior of the s_L buyers. Such an equilibrium in fact exists if and only if conditions (18) and (17) are satisfied. Both these conditions automatically hold if the first best cannot be supported in equilibrium, *i.e.*, if (11) fails.

However, (18) and (17) can hold, and thus a switching equilibrium can exist, even when (11) also holds. Provided that $\phi s_H > s_L$, the switching equilibrium is actually *better* than the first-best for the low switch-cost buyers, since they collect cash from the introductory offer before switching to their permanent supplier. Because the first-best contract cannot attract the s_L buyers when the switching-equilibrium contracts are in place, it would attract only the s_H buyers; but given that it gets only them, it is not credible! So switching equilibrium can persist even when the first-best equilibrium also exists. The key point is that the s_H buyers are unable to "split off" and cease to subsidize the s_L buyers.

To finish the proof, we must show that $\phi s_H \geq s_L$ is necessary and sufficient for a switching equilibrium. We have shown that it is sufficient: it clearly implies (18) and (17). But it is also necessary, since if $\phi s_H \leq s_L$ then the first best is an equilibrium: that

is, the first-best contract is credible if it attracts both buyer types. But if $\phi s_H \leq s_L$, then the first-best contract *does* attract both buyer types away from a putative switching equilibrium, so no switching equilibrium exists. ∎

Proof of Proposition 7

By the definition of $\hat{q}$, no deviant contract designed to attract and keep both types can earn positive profits. And of course no contract that will make both types switch away (from it) can be profitable. What about deviant contracts that attract both types and then induce switching by the s_L buyers? The most attractive remunerative such contract is the one that supports the switching equilibrium. Clearly, a necessary condition for the $\hat{q}$-equilibrium is that this deviant contract should not attract both types of buyers. (We need not worry about the prospect that it would attract only the s_L's, since it would then be unprofitable.) So we must check that the switching equilibrium contract is no more attractive to the s_H buyers than their equilibrium payoff. This requires $u^* + \phi s_H - s_H \leq \hat{u}$, or $\hat{u} - u^* \geq \phi s_H - s_H$.[19]

We also must consider deviant contracts designed to attract only one type of buyer. But attracting just one type and inducing switching is the null contract, which cannot profitably attract either type (either type could do better by switching away from the equilibrium contract). And if one type is attracted without switching, the other type also would be attracted. ∎

---

[19] Notice incidentally that when $\hat{q} = q^*$, $\hat{u} = u^*$ and the switching equilibrium contract cannot attract the s_H buyers.

References

Becker, Gary (1962) "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy*, **70**, S9–S49.

Crawford, Vincent (1983) "Long-Term Relationships Governed by Short-Term Contracts," Discussion Paper, University of California, San Diego; revised 1986.

Farrell, Joseph and Nancy Gallini (1986) "Second Sourcing as a Commitment," Working Paper, University of California, Berkeley.

Farrell, Joseph and Carl Shapiro (1987) "Dynamic Competition with Lock-In," Discussion Papers in Economics No. 121, Woodrow Wilson School, Princeton University, February.

Fudenberg, Drew, Bengt Holmstrom and Paul Milgrom (1987) "Short-Term Contracts and Long-Term Agency Relationships," mimeo.

Grossman, Sanford and Hart, Oliver (1986) "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, **94**, 691–719.

Hart, Oliver and John Moore (1985) "Contracts and Renegotiation," Working Paper No. 367, Department of Economics, M.I.T., Cambridge, MA.

Klemperer, Paul (1987a) "Markets with Consumer Switching Costs," *Quarterly Journal of Economics*, **102**, 375–394.

Klemperer, Paul (1987b) "The Competitiveness of Markets with Consumer Switching Costs," *Rand Journal of Economics*, **18**, 138–150.

Shapiro, Carl (1983) "Premiums for High-Quality Products as Returns to Reputations," *Quarterly Journal of Economics*, **98**, 659–679.

Williamson, Oliver (1985) *The Economic Institutions of Capitalism.* New York: The Free Press.

	$\hat{u} > u^* - s_H(1-\phi)$	$\hat{u} < u^* - s_H(1-\phi)$
$\phi s_H < s_L$	Short-Term Contracts Support First Best	Not
	Long-Term Contracts Support First Best	Possible
$\phi s_H < s_L$ and (11) Holds	Short-Term Contracts Give Switching Equilibrium	Not
	Long-Term Contracts Support First Best	Possible
(11) Fails	Short-Term Contracts Give Switching Equilibrium	Short-Term Contracts Give Switching Equilibrium
	Long-Term Contracts Give Switching Equilibrium & $\hat{q}$ Equilibrium	Long-Term Contracts Give Switching Equilibrium

Short-Term vs. Long-Term Contracts

Table 1

	s Observable Or $\phi s_H < s_L$	s Unobservable And $\phi s_H > s_L$
Quality Contractible	Short-Term Contracts Support First Best	Short-Term Contracts Give Switching Equilibrium
	Long-Term Contracts Support First Best	Long-Term Contracts Support First Best
Quality Not Contractible	Short-Term Contracts Support First Best	Short-Term Contracts Give Switching Equilibrium
	Long-Term Contracts Support First Best	Long-Term Contracts Give Switching Equilibrium & Possibly Superior $\hat{q}$ Equilibrium

Quality and Contract Length

Table 2

RECENT ISSUES OF THE WORKING PAPER SERIES
OF THE DEPARTMENT OF ECONOMICS
UNIVERSITY OF CALIFORNIA, BERKELEY

Copies may be obtained from the Institute of Business and Economic
Research. See the inside cover for further details.

8745   Leo K. Simon
       BASIC TIMING GAMES

       Jul-87.

8746   Leo K. Simon and Maxwell B. Stinchcombe
       EXTENSIVE FORM GAMES IN CONTINUOUS TIME:
       PURE STRATEGIES
       Jul-87.

8747   Joseph Farrell
       INFORMATION AND THE COASE THEOREM

       Jul-87.

8748   Philippe Aghion, Patrick Bolton and Bruno Jul
       LEARNING THROUGH PRICE EXPERIMENTATION BY A MONOPOLIST FACING UNKNOWN DEMAND

       Aug-87.

8749   Patrick Bolton
       THE PRINCIPLE OF MAXIMUM DETERRENCE REVISITED

       Aug-87.

8750   Jeffrey A. Frankel, Kenneth A. Froot
       CREDIBILITY, THE OPTIMAL SPEED OF TRADE LIBERALIZATION, REAL INTEREST RATES,
       AND THE LATIN AMERICAN DEBT
       Aug-87.

8751   Barry Eichengreen
       DID INTERNATIONAL ECONOMIC FORCES CAUSE THE GREAT DEPRESSION?

       Sep-87.

RECENT ISSUES OF THE WORKING PAPER SERIES
OF THE DEPARTMENT OF ECONOMICS
UNIVERSITY OF CALIFORNIA, BERKELEY

Copies may be obtained from the Institute of Business and Economic
Research. See the inside cover for further details.

8752 Jeffrey A. Frankel
      MONETARY AND PORTFOLIO-BALANCE MODELS OF EXCHANGE RATE DETERMINATION

      Sep-87.

8753 Albert Fishlow
      MARKET FORCES OR GROUP INTERESTS: INCONVERTIBLE CURRENCY
      IN PRE-1914 LATIN AMERICA
      Sep-87.

8754 Albert Fishlow
      SOME REFLECTIONS ON COMPARATIVE LATIN AMERICAN
      ECONOMIC PERFORMANCE AND POLICY
      Sep-87.

8755 Barry Eichengreen
      REAL EXCHANGE RATE BEHAVIOR UNDER ALTERNATIVE
      INTERNATIONAL MONETARY REGIMES: INTERWAR EVIDENCE
      Sep-87.

8756 Leo K. Simon and William R. Zame
      DISCONTINUOUS GAMES AND ENDOGENOUS SHARING RULES

      Oct-87.

8757 Leo K. Simon
      A MULTISTAGE DUEL IN CONTINUOUS TIME

      Oct-87.

8758 Joseph Farrell and Carl Shapiro
      OPTIMAL CONTRACTS WITH LOCK-IN

      Oct-87.

8759 Joseph Farrell and Eric Maskin
      RENEGOTIATION IN REPEATED GAMES

      Oct-87.