

Too Much of A Good Thing?

Johannes Hörner and Anna Sanktjohanser*

January 17, 2022

Abstract

We consider a repeated game, in which due to private information and a lack of flexible transfers, cooperation cannot be sustained efficiently. In each round, the buyer either buys from the seller or takes an outside option. The fluctuating outside option may be public or private information. When the buyer visits, the seller chooses what quality to provide. We find that the buyer initially forgoes mutually beneficial trades before then visiting more often than he would like to, myopically. Under private information, the relationship recurrently undergoes gradual self-reinforcing downturns when trust is broken and instantaneous recoveries when loyalty is shown.

KEYWORDS: Trust, Loyalty, Imperfect Monitoring.

JEL CLASSIFICATION NUMBERS: C72, C73, C78.

*Hörner: Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, and CNRS (Toulouse-Capitole), 1 Esplanade de l'Université, 31080 Toulouse Cedex 06, France, joh.horner@gmail.com. Sanktjohanser: Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, and Toulouse School of Economics, 1 Esplanade de l'Université, 31080 Toulouse Cedex 06, France, anna.sanktjohanser@yale.edu. We would like to thank many seminar audiences, as well as (in alphabetical order) Simon Board, Alessandro Bonatti, Hector Chade, Drew Fudenberg, Hugo Hopenhayn, Jean Tirole and Alex Wolitzky. We would like to thank Stephan Waizmann and John Mori for excellent research assistance. Johannes Hörner acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program). Anna Sanktjohanser acknowledges that this project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 669217 - MARK-LIM).

Business success (for small and large businesses alike) is often tied to loyal relationships with trading partners. Virtually all large retail companies—from airlines to supermarket chains—run loyalty programs for their customers, recognizing that a good customer is a repeat customer. Traders in poor countries cite personal relationships as the most important factor for success (Fafchamps and Minten, 1999; Fafchamps, 2004). In other words, “to be an effective competitor (in the global economy) requires one to be a trusted cooperator (in some network)” (Morgan and Hunt, 1994). However, how are such relationships with “trusted cooperators” sustained, and how do they evolve over time when trading opportunities are not common knowledge?

This paper analyzes how a loyal relationship between two parties can be preserved over time in the face of adverse selection and moral hazard. More precisely, we consider a repeated game between a buyer and a seller, in which due to private information and a lack of flexible transfers, cooperation cannot be sustained efficiently. The buyer repeatedly faces a choice between buying from the seller (at exogenous terms) and taking an outside option. The value of the outside option (which may or may not be private information of the buyer) is i.i.d. and drawn from some exogenous distribution. When the buyer visits the seller, the seller chooses which quality to provide (at a cost). We are interested in understanding how this relationship is optimally managed, how it evolves, and how its evolution depends on the information available to the seller. With low discounting, a folk theorem holds. Instead, we assume that players are impatient and characterize the (exactly) optimal equilibria, with a particular focus on the buyer’s favorite equilibrium.

Our main results are as follows: (1) regardless of whether the buyer has private information, in the buyer-preferred equilibrium, (i) the buyer initially forgoes mutually beneficial trades (the *consideration stage*) to avoid having to pay rents, and (ii) once the relationship has started, the buyer will trade more often than he would myopically like to in order to preserve the relationship (the *loyalty loop*); (2) making information common knowledge reduces volatility in the relationship: in fact, in the buyer-preferred equilibrium when there is no private information, behavior is stationary after the first visit; (3) in the face of private information, the relationship

necessarily involves varying levels of loyalty and quality: the relationship experiences gradual downturns when trust is broken, but instantaneously recovers if loyalty is shown; (4) even in the face of private information, trust is eventually re-established: the relationship never breaks down.

The intuition for the consideration stage and the loyalty loop is as follows. Once the relationship has started, the buyer must motivate the seller to provide the desired quality and does so by coming more often than he would like (myopically). The buyer wants to delay starting this loyalty loop. The need to deliver rents to the seller after the first purchase means the buyer forgoes beneficial alternative trades; he therefore waits longer than he would like to in a myopic sense until he first visits. The seller is too much of a good thing for the buyer: to motivate the seller, the buyer must patronize her more often than he would like.

When there is no private information, the buyer-preferred equilibrium is stationary after the first purchase: the buyer uses a simple cut-off strategy, where he visits the seller if, and only if, the value of the outside option is sufficiently low, and the seller provides a constant quality. Given that there is no private information, deviations are observable, and no variation in payoffs or behaviors on the path is required to provide incentives: without loss, any deviation leads to a permanent breakdown of cooperation (players are “quick to anger and never forgive”).

When the value of the outside option is private information, the buyer-preferred equilibrium is no longer stationary after the first visit: because the buyer’s deviations are no longer observable, quality varies with the buyer’s loyalty. Consecutive failures to visit eventually lead to decreases in quality, which in turn make the buyer less willing to visit, by reducing the range of outside options for which he finds it worth his while. Continuing decreases in quality hurt the buyer, but continuing decreases in loyalty also hurt the seller. Hence, re-establishing trust when the buyer visits the seller boosts the seller’s payoff, and motivates the buyer to visit her. Therefore, when loyalty is shown (*i.e.*, when the buyer visits the seller), trust is immediately re-established: after just a single purchase from the seller, quality jumps back to the (appropriately defined) socially desirable level (players are “slow to anger and quick

to forgive”).

The relationship may slowly continue to sour for arbitrarily long stretches of time, with the buyer visiting less and less frequently, and the seller offering lower and lower quality. Fortunately, this downward spiral occurs sufficiently slowly that a visit always eventually reoccurs. Hence, the relationship never completely breaks down, even in the face of private information. Instead, this relationship alternates forever between phases of slow erosion in quality and loyalty, and immediate re-establishment of those.

The results speak to a variety of phenomena observed in long-term business relationships. For instance, the distinction between the consideration stage and the loyalty loop (reminiscent of the “insider bias” in Board (2011)) has been documented by the marketing literature in various contexts from the choice of fast-food restaurants (Nedungadi, 1990) to high-technology markets (Heide and Weiss, 1995). Our paper offers a simple explanation for why forgoing such mutually beneficial trades may be optimal for consumers.

Traders in developing countries cite loyal relationships as the most important factor in business success. Part of the reason for this is that repeated interactions facilitate the flow of information and ensure the regularity of trade (Fafchamps and Minten, 1999; Fafchamps, 2004). This is consistent with our findings: making information common knowledge reduces volatility in the relationship, thereby ensuring regular trade relative to a situation with private information.

Our paper also speaks to the structural asymmetry in many multitier loyalty programs (Bijmolt et al., 2018). The purpose of loyalty programs is to build a loyal customer base. Multitier loyalty programs categorize customers into hierarchical tiers based on behavior, rewarding customers with progressively preferential treatment and special privileges, among other things. Many such multitier loyalty programs display structural asymmetry in regard to promoting and demoting customers: airlines such as *Air France* have no limit on how quickly a customer can reach the highest tier (*Platinum* status) in its loyalty program—however, even if the customer has failed to be loyal over a given year, the airline demotes the customer by at most one tier. This structure is consistent with our result that downturns in the business relationship are

gradual, whereas recoveries are instantaneous.

Our work is related to two main strands of literature. The first is the literature on trading favors (see Athey and Bagwell, 2001; Möbius, 2001; Fuchs and Lippi, 2006; Hauser and Hopenhayn, 2008). In this literature, there is two-sided adverse selection: players have (privately known) opportunities to grant a favor at random times. Granting favors is efficient but costly. In our model, we have adverse selection on the buyer’s side but moral hazard on the seller’s side.

While there are no flexible transfers included in this analysis, the paper is also related to the literature on relational contracts with random opportunity costs (Board, 2011; Li and Matouschek, 2013). In this literature, the opportunity cost of paying a worker varies over time.¹ In Board (2011), the opportunity cost is publicly observable, but the principal can choose to assign a task between multiple agents. Board (2011) shows that there is insider bias, *i.e.*, the principal is more likely to assign the task to an agent with whom he has already established a relationship. While we do not model competition explicitly, the consideration stage (*i.e.*, the “hesitation” to start a relationship with a new seller) in our paper is reminiscent of the insider bias in Board (2011). In Li and Matouschek (2013), the relationship experiences gradual downturns as the principal is repeatedly hit by negative shocks, but trust can quickly be rebuilt. This is very similar to the dynamics in the unobservable case in our paper. However, in Li and Matouschek (2013), cycles are driven by the concavity of the production function, as they note, whereas we assume that costs are linear and make no assumption on the distribution of outside options.² Crucially, the availability of transfers in this literature means that a player can be rewarded or punished without

¹Also related is Halac (2012), where the opportunity cost is private information, but unlike in Board (2011), Li and Matouschek (2013) and this paper, the private information is perfectly persistent.

²To develop the analogy, both models involve two instruments to transfer utility between buyer and seller. In Li and Matouschek (2013), payments are allowed, a linear (but bounded) instrument. Here, varying the quality level provides a similar, linear instrument. In their paper, the second instrument, the technology, is assumed to be strictly concave. In ours, varying the continuation payoff vector is the second instrument. However, its strict concavity is a result, not an assumption. The analogy has its limitations, however: here, the asymmetry between gradual downturns and drastic recoveries is driven by the alignment between the seller’s payoff and the buyer’s *incentives* (in their paper, what matters is the alignment in payoffs as they emphasize).

surplus destruction. Here, on the other hand, the seller can only punish by decreasing the quality, but this is inefficient.³

The rest of the paper is structured as follows. The model is described in Section 2. Section 3 analyzes the complete information case (*i.e.*, the outside option is observable). Section 4 studies the incomplete information case (*i.e.*, the outside option is the private information of the buyer). Section 5 concludes the paper.

1 Set-Up

A buyer (he) and a seller (she) interact repeatedly. Time is discrete and infinite. In each round $n = 0, \dots$, the interaction unfolds as follows:

1. The buyer draws an outside option. This random variable \tilde{v} is i.i.d. across rounds, according to F , with support $\mathcal{V} = [\underline{v}, \bar{v}]$, with $\underline{v} \geq 0$.⁴ The distribution F admits a density f bounded away from zero on \mathcal{V} . The buyer observes this outside option. The seller may or may not observe this outside option: this distinction defines the observable vs. the unobservable case.
2. The buyer chooses whether to visit the seller (“*In*” or, simply, “*I*”) or to take the outside option (“*Out*” or, simply, “*O*”).
3. If the buyer chooses *In*, the seller picks a quality $q_n \in \mathcal{Q} := [0, 1]$.

If the buyer chooses *Out*, the reward vector to the players (buyer and seller) is $(v_n, 0)$, where v_n is the realized outside option in round n . If he chooses *In*, the vector is $(q_n - p, p - cq_n)$, where p, c are exogenous parameters such that $0 < c < p < 1$, and $\bar{v} > 1 - p$ (outside options can be attractive). That is, the seller receives a fixed, exogenous price p whenever the buyer chooses *In*, independent of the quality she

³In motivation, our paper is also related to the literature on relational contracts with random cost of effort (Levin, 2003; Chassang, 2010; Calzolari and Spagnolo, 2017). In Levin (2003) and Chassang (2010), the focus is (as in our paper) on the relationship between one principal and one particular agent. Calzolari and Spagnolo (2017) focus on the principal’s choice between multiple agents, and as in Board (2011) finds that there is an “insider bias.” In Levin (2003) and Calzolari and Spagnolo (2017), the cost of effort is private information, whereas in Chassang (2010), there is learning about the production possibilities of the agent over time.

⁴The assumption $\underline{v} \geq 0$ is for convenience and can be replaced by $\underline{v} > -p$: some outside options can be costly, as long as they are still preferable to paying the seller for a worthless good.

picks. Quality is measured in the buyer's utility and entails a linear cost to the seller. Given that the price is simply a transfer, surplus is v_n when the buyer chooses *Out* and $q_n(1-c)$ when he chooses *In*. Hence, the socially efficient decision is for the seller to choose $q_n = 1$ for all n (whenever the buyer chooses *In*) and for the buyer to go to the seller whenever $v_n \leq 1 - c$. However, in the one-shot game, because $p > c$, the buyer would only visit the seller if $v_n < 1 - p$ if he expects the seller to provide maximum quality. Of course, in the one-shot game, it is optimal for the seller to pocket the price and pick zero quality.

Players share a common discount factor $\delta \in [0, 1)$. Realized payoffs are then

$$B_0 = (1 - \delta) \sum_{n \geq 0} \delta^n (\mathbf{1}_{\{a_n=I\}}(q_n - p) + \mathbf{1}_{\{a_n=O\}}v_n)$$

for the buyer, and

$$S_0 = (1 - \delta) \sum_{n \geq 0} \delta^n \mathbf{1}_{\{a_n=I\}}(p - cq_n)$$

for the seller.

Throughout, the seller's quality choice is observable. For the moment, we further assume that information is complete: the seller observes the outside option before choosing quality.^{5,6}

A history $h^{n-1} \in H^{n-1}$ is a sequence $(v_0, q_0, \dots, v_{n-1}, q_{n-1}) \in (\mathcal{V} \times (\mathcal{Q} \cup \{\emptyset\}))^n$, with the convention that $q_m = \emptyset$ whenever the buyer chooses $a_m = O$ in that round, that is, if the buyer chooses not to visit the seller. A behavior strategy for the seller, then, is a sequence $\sigma^S := (\sigma_n^S)_n$, where σ_n^S is a probability transition from $H^{n-1} \times \mathcal{V} \times \{I, O\}$ to $\mathcal{Q} \cup \{\emptyset\}$, with the restriction that $\sigma_n^S(h^{n-1}, v_n, a) = \emptyset$ if, and only if, $a = O$.⁷ A behavior strategy for the buyer is a sequence $\sigma^B := (\sigma_n^B)_n$, where σ_n^B is a probability transition from $H^{n-1} \times \mathcal{V} \rightarrow \{I, O\}$. Note that the buyer's choice in past rounds is ignored in the histories, as it is encoded in the seller's quality choice. Given a strategy profile $\sigma = (\sigma^B, \sigma^S)$, expected payoffs are defined in the obvious way. When no risk

⁵For most results, one can equivalently assume that the seller only observes the outside option if the buyer takes it, rather than visiting her, an assumption more reasonable for some applications.

⁶Hence, we might equivalently assume that the outside option is fixed, but the seller's cost fluctuates.

⁷That is, for each $h^{n-1} \in H^{n-1}$, $\sigma_n^S(h^{n-1})$ is a distribution over $[0, 1] \cup \{\emptyset\}$, and the probability assigned to any Borel set $A \subset [0, 1] \cup \{\emptyset\}$ is a measurable function of h^{n-1} .

of confusion arises, we use B_0, S_0 to denote this (expected) payoff as well. We use subgame-perfect Nash equilibrium as a solution concept.⁸ For conciseness, we omit many statements that only hold “with probability one.” Throughout, we assume that a public randomization device is available, even if it is omitted from notations.⁹

We note that the seller can secure 0 by always choosing zero quality, and the buyer can secure $\mathbf{E}[\tilde{v}]$ by always choosing the outside option. These are also the minmax payoffs, and the vector $(\mathbf{E}[\tilde{v}], 0)$ is an equilibrium payoff: if the seller expects the buyer to never return, saving on the quality cost is optimal if the buyer unexpectedly shows up. Conversely, given that the buyer expects the seller to choose zero quality, it is best for the buyer to always pick the outside option. This strategy profile is referred to as *autarky*.

There are three curves tracing the boundary of the feasible payoff set \mathcal{F} . First, suppose the buyer uses a cut-off rule, according to which he visits the seller when the outside option is below some cut-off, and the seller chooses maximum quality. As the cut-off varies, the resulting payoff vector traces a curve in the space of buyer/seller payoff pairs (B, S) . Second, suppose the buyer always visits the seller, and the seller chooses the same quality whenever the buyer visits. As the quality varies, the resulting payoff vector traces another curve. Finally, suppose the buyer uses another cut-off rule, coming to the seller when the outside option is *above* some cut-off, and the seller chooses zero quality whenever the buyer visits. As the cut-off varies, the resulting payoff vector traces yet another curve. The convex hull of these curves is the feasible payoff set. The set of individually rational payoffs \mathcal{V} further requires that the seller obtain at least 0 and the buyer his outside option $\mathbf{E}[\tilde{v}]$.¹⁰ The sets \mathcal{F} and \mathcal{V} (shaded

⁸The solution with commitment on the buyer’s side is derived in Section 2.1. If the seller can commit, trade becomes uncomplicated, since by assumption $p > c$, and so the seller breaks even.

⁹It plays no role in the best equilibria we derive but facilitates the analysis and is used for simple equilibria; see Section 4.1.

¹⁰More formally, the feasible payoff is given by

$$\mathcal{F} := \{(B, S) \in \mathbf{R}^2 \mid B^L \leq B \leq B^U\},$$

where straightforward and omitted calculations give the bounds

$$B^U := \min \left\{ \frac{p(1-c)}{c} - \frac{S}{c}, \frac{1-p}{p-c}S + \left(1 - \frac{S}{p-c}\right) \mathbf{E} \left[\tilde{v} \mid \tilde{v} \geq F^{-1} \left(\frac{S}{p-c} \right) \right] \right\},$$




Figure 1: Feasible (and individually rational) payoff set when $(\delta, p, c) = \left(\frac{13}{20}, \frac{13}{25}, \frac{1}{4}\right)$.

area) are illustrated in Figure 1.

Plainly, a folk theorem holds in our setting regardless of whether outside options are observed: any feasible, strictly individually rational payoff can be supported provided the buyer and seller are sufficiently patient.¹¹ Our interest is in situations in which this is not the case.

Our goal is to understand the relationship between the buyer and seller and how it evolves. On the seller's side, a natural measure of the strength of the relationship is given by the quality she supplies; on the buyer's side, this measure is the probability with which he visits the seller.¹²

and

$$B^L := \left(1 - \frac{S}{p}\right) \mathbf{E} \left[\tilde{v} \mid \tilde{v} \leq F^{-1} \left(1 - \frac{S}{p}\right) \right] - S.$$

Hence, the feasible and individually rational payoff set is given by

$$\mathcal{V} := \{(B, S) \in \mathbf{R}^2 \mid (B, S) \in \mathcal{F}\} \cap [\mathbf{E}[\tilde{v}], \infty) \times \mathbf{R}_+.$$

¹¹The standard results under imperfect monitoring (*e.g.*, Fudenberg, Levine and Maskin, 1994) do not apply given that action and type sets are infinite, but we anticipate no difficulty in adapting Radner's (1985) review strategies to this setting.

¹²Because quality and outside options are drawn from intervals, we avoid the mixed strategies

To fix ideas, we focus mostly on the buyer-preferred equilibrium—that is, the equilibrium that maximizes the buyer’s *ex ante* payoff. To determine this equilibrium, it is convenient to solve for the Pareto frontier of the equilibrium payoff set, and the characterization of other equilibria (for instance, the seller-preferred equilibrium) follows.

Given that complete information is assumed for now, all deviations from a pure strategy profile are observable. Therefore, any equilibrium remains an equilibrium if we replace the specification of the continuation strategies after a deviation into autarky. This specification is assumed throughout, though it is not always stated explicitly. Hence, we focus on describing behavior as long as neither player has deviated.

The buyer’s favorite stage-game outcome involves visiting in any round if, and only if, his outside option is less than $q - p$, his payoff from visiting the seller. We refer to such a decision rule by the buyer as *myopic* or *opportunistic*, as it ignores the potential fallout from failing to visit the seller. Of course, the buyer would also like the seller to provide maximum quality. Independent of whether the seller does so, we rule out the possibility of myopic behavior in equilibrium (except for in the case of autarky) through the following assumption on the parameters.

Assumption A1: *For all $q > 0$, it holds that:*

$$(1 - \delta)cq > \delta F(q - p)(p - cq).$$

It immediately follows that

Lemma 1. *Under A1, there is no equilibrium in which the buyer visits if, and only if, $v_n \geq \mathbf{E}_{h^{n-1}, v_n}[q_n] - p$ for all n , h^{n-1}, v_n .*

Proof. Let

$$\lambda := \min_{q \in \mathcal{Q}} \{(1 - \delta)cq - \delta F(q - p)(p - cq)\},$$

which is strictly positive, given A1. Assume first that the buyer visits w.p.p. (with positive probability) in every round. Incentive compatibility requires that, for all n ,

 (and their counterintuitive properties) that would arise with discrete supports.

all h^{n-1}, v_n , and for all q_n in the support of $\sigma_n^S(h^{n-1}, v_n)$,

$$\begin{aligned} (1 - \delta)cq_n &\leq \mathbf{E}_{h^{n-1}, v_n, q_n} \left[\sum_{m=n+1}^{\infty} (1 - \delta)\delta^{m-n} F(\mathbf{E}_{h^{m-1}, v_m}[q_m] - p)(p - c\mathbf{E}_{h^{m-1}, v_m}[q_m]) \right] \\ &\leq (1 - \delta)c\mathbf{E}_{h^{n-1}, v_n, q_n} \left[\sum_{m=n+1}^{\infty} \delta^{m-n-1} ((1 - \delta)q_m - \lambda/c) \right], \end{aligned}$$

by the law of iterated expectations and **A1**. In particular, this requires

$$\mathbf{E}[q_n] \leq (1 - \delta) \sum_{m=n+1}^{\infty} \delta^{m-n-1} \mathbf{E}[q_m] - \lambda/c.$$

Plainly, this cannot hold for all n , given that $(\mathbf{E}[q_n])_n$ is a bounded sequence. If in some rounds, the buyer does not visit, some terms on the right-hand side must be replaced by zeros, and the contradiction follows along the same lines. \square

Hence, opportunistic behavior cannot be sustained as a persistent equilibrium phenomenon: the buyer must visit more often than he would like, at least occasionally.¹³ We impose Assumption **A1** to rule out such trivial equilibria.

In general, the condition in **A1** cannot be simplified by requiring it to hold for $q = 1$ only, though this is the case for many distributions, including the uniform distribution.¹⁴ It is perhaps easiest to view this as an assumption regarding the seller's impatience, as it is satisfied provided the interest rate $(1 - \delta)/\delta$ is sufficiently large.

Assumption **A1** is maintained throughout (including throughout the unobservable case) and omitted from all statements.

While there are parameters for which autarky is the only equilibrium, we focus on the case in which other equilibria exist, that is, the parties are impatient enough that the seller's incentives are not automatically satisfied but patient enough that she can be motivated.

¹³This does not imply that equilibria in which the buyer behaves myopically cannot exist; the buyer can also provide rents to the seller by accepting a lower quality good when his outside option is very low. Either way, the buyer must compromise on what he would like.

¹⁴Assumption **A1** certainly holds for $q = p$. When F is uniform, the difference between the left- and right-hand side is convex in q , so that it holds for all q if it holds for $q = 1$.

(I)initial Purchase (S)ubsequent Purchases




Figure 2: Buyer favorite equilibrium.

2 Observable Case: Main Results

Because of Assumption **A1**, the buyer cannot expect the seller to provide quality unless he makes a deliberate effort to visit her more often than he would like: he gets too much of a good thing. More formally, continuation strategies must deliver sufficient rents to the seller for her to be willing to provide the requisite quality.

Despite the “stationary” structure of the environment, the optimal decision rule for the buyer is not stationary. The following proposition describes the buyer-preferred equilibrium, which is illustrated by the two-state automaton in Figure 2. Recall that only on-path behavior is specified, as any deviation triggers autarky.

Proposition 1. *The buyer’s favorite equilibrium involves a constant quality q^* provided by the seller and two cut-off rules for the buyer. As long as the buyer chose Out in all prior rounds, he visits the seller in round n if, and only if, $v_n \leq v_I^*$. Otherwise, he visits the seller if, and only if, $v_n \leq v_S^*$. These cut-offs are such that*

$$v_I^* < q^* - p < v_S^*.$$

That is, the buyer uses two cut-offs: an *Initial* cut-off v_I^* that involves a lower probability of visiting the seller than myopic behavior would entail and a *Subsequent* cut-off v_S^* that leads to more visits to the seller than he would like.

In the marketing literature, these two phases are commonly referred to as the *consideration stage* and the *loyalty loop*. This dichotomy is widely documented – from the choice of fast-food restaurants (Nedungadi, 1990) to high-technology markets

(Heide and Weiss, 1995). It is usually ascribed to an initial lack of familiarity with the seller's product. Here instead, it is driven by the rents that the buyer owes the seller depending upon whether he has visited her yet. Using the higher cut-off reflects what is called commitment in that literature, "the implicit or explicit pledge of relational continuity between exchange partners," involving the willingness to make short-term sacrifices to realize longer-term benefits (Dwyer, Schurr and Oh, 1987).

As mentioned, what distinguishes these two phases is the promised utility to the seller. At the outset, the buyer owes the seller nothing. Once he visits her, his future behavior must provide the seller with a continuation payoff that compensates her for her cost. Both the initial payoff vector, and the subsequent payoff vector lie on the Pareto frontier of the equilibrium payoff set. The proof of the proposition follows from a general characterization of this Pareto frontier.

Given that the equilibrium behavior is not stationary, one might wonder why it does not involve more nuanced gradualism. After all, relationships that "start small" and grow as the relationship evolves are common (see Watson, 1999, 2002). Here, once the buyer crosses the seller's doorstep, there is no reason to delay engaging with the seller in the most profitable way possible, fast-forwarding to whichever continuation play is the most desirable. This is not a feature of the buyer-preferred equilibrium only but of any Pareto-efficient equilibrium - what is critical is that each party's preferences and opportunities are complete information.

To understand what determines these cut-offs, let us assume for now that maximum quality is both feasible (patronizing the seller often enough delivers sufficient rents) and desirable and that the buyer has no incentive to renege when his outside option calls for a seller visit. Further, let us assume that these values are drawn from the uniform distribution on the unit interval.

From the buyer's point of view, visiting the seller is equivalent to drawing the cut-off outside option at which he is indifferent between going or not, when in the consideration stage. That is, his payoff can be computed *as if* he never visits the seller, and remains forever in the consideration stage, but cashes in as outside option the critical cut-off whenever the actual outside option is below this cut-off. Hence, as

a function of this cut-off v_I^* , the buyer's payoff in state I , denoted B_I , is given by

$$B_I = F(v_I^*)v_I^* + (1 - F(v_I^*))\mathbf{E}[\tilde{v} \mid \tilde{v} \geq v_I^*] = \frac{1 + (v_I^*)^2}{2}. \quad (1)$$

On the other hand, the buyer's cut-off in the loyalty loop is determined by the seller's incentives, not the buyer's indifference. The cut-off in state S leaves no excess rents for the seller. She must be indifferent between supplying the maximum quality and saving the one-time cost if she does not; thus, the cut-off must solve

$$\delta F(v_S^*)(p - c) = (1 - \delta)c, \text{ or } v_S^* = \frac{1 - \delta}{\delta} \frac{c}{p - c}. \quad (2)$$

Hence,

$$B_S = F(v_S^*)(1 - p) + (1 - F(v_S^*))\mathbf{E}[\tilde{v} \mid \tilde{v} \geq v_S^*] = v_S^*(1 - p) + \frac{1 - (v_S^*)^2}{2}, \quad (3)$$

which is less than $(1 + (v_S^*)^2)/2$ precisely because the buyer must come more often than he likes: $v_S^* > 1 - p$.

The remaining unknown is the cut-off v_I^* . It is pinned down by the buyer's indifference in the consideration stage, when his outside option happens to be equal to this cut-off, namely,

$$(1 - \delta)v_I^* + \delta B_I = (1 - \delta)(1 - p) + \delta B_S,$$

which immediately implies that $v_I^* < 1 - p$, since $B_I > B_S$. All formulas for the payoffs and cut-offs immediately follow. A higher cost, or a lower discount factor, increases both cut-offs and decreases the buyer's payoff. A higher price decreases v_S^* , and this may be good (or not, depending on the parameters) for the buyer, as he needs to forfeit fewer outside options to compensate the seller. However, a higher price can increase v_I^* , as it makes initiating a relationship with the seller less costly and hence more attractive.

Proposition 1 does not explicitly specify the quality that the seller provides. (The values of q^* are discussed below; see Proposition 2.) Even when equilibria exist in which maximum quality is provided, the buyer might prefer a lower quality level. Indeed, to the extent that he needs to patronize the seller often enough to repay

his debt, he internalizes the future cost in terms of the foregone opportunities that a higher quality level calls for, and this future cost might deter him from seeking a higher quality.

In the absence of transfers, the relationship involves what anthropologists refer to as gift giving (see for instance Mauss, 1925). By giving a gift, they argue, one establishes a relationship by placing the recipient in debt. From the buyer's point of view, it is preferable that the gift that the seller bestows when she provides high quality does not make him too indebted, as he will have to reciprocate with more visits than he might find ideal.

There is another, more subtle, reason why quality might not be as high as possible which we assumed away in the simple formulas above. Even if the buyer would like to patronize the seller sufficiently often to compensate her for providing maximum quality, such visits must be credible: that is, the requisite cut-off might be so high that if the outside option is at this cut-off, the buyer might prefer to renege and take the outside option, even if it implies that autarky prevails thereafter. That is, the cut-off might not satisfy *interim* incentive compatibility on the buyer's side, and this credibility constraint might force the buyer to fall back on a lower cut-off; hence, on lower seller's rents and thus, on a lower quality.

Here, our focus has been on the buyer-preferred equilibrium. A byproduct of the proof of Proposition 1 is a "near" characterization of the boundary of the entire equilibrium payoff set.¹⁵ Here, we provide a brief informal description of the entire equilibrium payoff set.¹⁶ The payoffs on the boundary of the equilibrium payoff set which deliver less than the seller receives after the first visit (as described in Proposition 1) are delivered in a fashion similar to that of the buyer-preferred equilibrium: the initial cut-off is calibrated to deliver the desired payoff to the seller, but the continuation after the first visit matches the buyer's favorite equilibrium. When the seller's payoff is higher than what she receives after the first visit in the buyer-preferred equi-

¹⁵More precisely, Proposition 2 does not provide information about the boundary of the equilibrium payoff set when the seller receives less than the buyer would grant her in his favorite equilibrium and ignores the buyer's *interim* incentive compatibility.

¹⁶See the additional appendix for further details, and see Figure 3 for a graphical illustration of the equilibrium payoff set.

librium, the seller's incentive compatibility no longer binds. In this case, the buyer's incentive compatibility might pin down the cut-off that determines the visits by the buyer. For sufficiently high seller payoffs, the seller provides suboptimal quality when the outside option is low, that is, when the buyer is not tempted to take it.




Figure 3: Equilibrium payoff set (blue) vs. feasible & IR set (red) when $(\delta, p, c) = (\frac{13}{20}, \frac{13}{25}, \frac{1}{4})$.

2.1 A Sketch of the Proof and a More General Result

This section can be skipped without loss for the reader with little time for proofs. The proof borrows techniques introduced by Spear and Srivastava (1987) and Thomas and Worrall (1990). To establish Proposition 1, we solve for the bounded solution to the following optimization program:

$$\mathcal{B}: [0, p] \rightarrow \mathbf{R} \tag{\mathcal{P}}$$

$$S \mapsto \sup (1 - \delta) \left[\int_A (q(v) - p) F(dv) + \int_{A^c} v F(dv) \right] + \delta \int_{\mathcal{V}} \mathcal{B}(S(v)) F(dv),$$

over $A \subset \mathcal{V}$, $S: \mathcal{V} \rightarrow \mathbf{R}_+$, and $q: \mathcal{V} \rightarrow \mathcal{Q} \cup \{\emptyset\}$ (all functions of S) such that $\forall v \in A$,

$$(1 - \delta)cq(v) \leq \delta S(v), \tag{4}$$

and

$$S \leq (1 - \delta) \int_A (p - cq(v)) F(dv) + \delta \int_{\mathcal{V}} S(v) F(dv). \tag{5}$$

The set A is the set of outside options for which the buyer visits the seller, $q(\cdot)$ is the quality expected from the seller, and $S(\cdot)$ is her (on-path) minimum continuation

utility given the buyer's outside option v .¹⁷ Equation (4) has an obvious interpretation as an incentive constraint for the seller and (5) as a promise-keeping constraint. Note that promise-keeping involves an inequality: that is, we only require that the seller receive no less than S . The program \mathcal{P} delivers the maximum payoff $\mathcal{B}(S)$ to the buyer given that the seller's payoff is at least S and that he does not anticipate more than $\mathcal{B}(S(v))$ from the next round onward given that the seller must receive at least $S(v)$ going forward.

Program \mathcal{P} assumes that the buyer never randomizes his decision: either he visits the seller, or he does not. This is made for notational convenience and entails no loss of generality. One might expect that the seller's promised utility from the next round onward need not depend on the current outside option provided that it lies in A (or in A^c). Indeed, one of the steps involved in solving \mathcal{P} consists of showing that, given S , it is sufficient to consider a pair of promises $(S^I, S^O) \in \mathbf{R}_+^2$, depending on whether the buyer picks *In* or *Out*. Unless the promise S is very high, one can also pick a quality that is independent of the outside option, that is, a scalar $q \in \mathcal{Q}$. Finally, one might expect that A can be taken to be a lower interval $[\underline{v}, v^*]$: after all, the seller only cares about the probability that the buyer visits, not about the circumstances under which he does. This simplification is also warranted, as we show.

Finally, as noted above (see ft. 15), the program omits the buyer's incentive to visit, depending on the realized outside option. It assumes commitment on the buyer's side; namely, for all $v' \in A$,

$$(1 - \delta)v' + \delta\mathbf{E}[\tilde{v}] \geq (1 - \delta)(q(v') - p) + \delta\mathcal{B}(S(v')). \quad (6)$$

This *does* entail a loss of generality, as it rules out a set of parameters for which (6) affects the quality q and cut-off v^* that characterize the buyer-preferred equilibrium. As mentioned, the reader is referred to the additional appendix for a proof that takes this additional constraint into account.

The uniqueness of the bounded solution of \mathcal{P} is a consequence of the principle of optimality, given discounting. We now give the arguments of the maximum. The proof

¹⁷The use of S for the (scalar) minimum promise and of $S(\cdot)$ as the continuation promise, given v and S , should cause no confusion because the problem will immediately simplify.

(in Appendix A) consists of solving for $\mathcal{B}(\cdot)$ given the conjectured transitions (that is, the values S^I, S^O) and then applying a verification argument, namely, showing that the resulting function solves the optimality equation.

Let $\kappa := \min\{p(1-c)/c, \bar{v}\}$. It holds that $\kappa > \underline{v}$, for otherwise $1-p < \underline{v}$, and autarky is the unique equilibrium.

How the promised utility evolves depends on how it compares to a cut-off. Let

$$\tilde{S} := \min \left\{ \frac{1-\delta}{\delta}c, \frac{p}{\frac{1}{F(\kappa)} + \frac{\delta}{1-\delta}} \right\}. \quad (7)$$

Proposition 2. *For all $S \in [0, pF(\kappa)]$, the following holds for the solution of \mathcal{P} .*

Promises are $S^O = S$,¹⁸ and $S^I = \max\{S, \tilde{S}\}$. Quality is equal to

$$q^*(S) = \min \left\{ 1, \frac{1}{c} \left(p - \frac{\max\{S, \tilde{S}\}}{F(\kappa)} \right) \right\},$$

provided the minimum is positive, and 0 otherwise. Finally, for some (unique) $\underline{S} \in [0, \tilde{S}]$,

$$v^*(S) = \begin{cases} F^{-1} \left(\frac{(1-\delta)\max\{\underline{S}, S\}}{(1-\delta)p - \delta \max\{\underline{S}, S\}} \right) & \text{for } S < \tilde{S}, \\ \min \left\{ F^{-1} \left(\frac{S}{p-c} \right), \kappa \right\} & \text{for } S \geq \tilde{S}, \end{cases}$$

whenever $q^ \geq 0$ ($v^*(S) = F^{-1}(S/p)$ otherwise).*

The seller's incentive constraint binds if, and only if, $S \leq \tilde{S}$. The threshold \underline{S} is the minimum payoff that the buyer delivers to the seller: because giving the seller less than \tilde{S} hurts the buyer as well, he always gives at least \tilde{S} to the seller, even if he does not owe her that much.¹⁹ Any higher promise is exactly delivered, as the buyer's payoff then decreases in S . That is, \underline{S} is the seller's payoff at which the buyer's payoff is maximized and therefore also the payoff that he would choose at the outset; \underline{S} is the seller's payoff in the consideration stage. In contrast, \tilde{S} is the payoff she receives when the loyalty loop starts. Promises above \tilde{S} require frequent visits, and the buyer

¹⁸If $S \leq \underline{S}$, we may equivalently specify that $S^O = \underline{S}$, since this is the payoff actually delivered to the seller in the continuation. Note that the domain of the seller's promised utility is $[0, p]$, which is larger than $[0, pF(\kappa)]$, when $\kappa < 1$. For $S > (pF(\kappa), p]$, lower quality is provided for low values of the outside option. We omit the details.

¹⁹The equation that characterizes \underline{S} is given in the proof of Proposition 2 (See (10), Appendix A).




Figure 4: Quantile and Quality.

no longer gains from delivering distinct continuation payoffs according to his choice to visit.

For $S \leq \tilde{S}$, the quality that the seller chooses if the buyer chooses In is constant: what matters to the seller is the continuation payoff S^I , which is constant over this range (and equal to \tilde{S}). Higher promises call for lower quality choices (though not necessarily strictly lower choices, in case the boundary condition $q \leq 1$ binds). Higher promises also call for more frequent visits, so that the cut-off v^* strictly increases in $S \geq \underline{S}$, the lowest seller's payoff actually delivered.

Figure 4 illustrates Proposition 2.

2.2 Discussion

Persistence Here, it has been assumed that the outside option is drawn afresh in every round, independent of past realizations. This reduces the relationship to a repeated game (though with a nontrivial extensive form in each round) and so simplifies the analysis considerably.²⁰ In practice, however, alternative trading opportunities

²⁰The i.i.d. assumption prevents the analysis from being conducted in continuous time. With persistence, however, equilibrium analysis would require the current outside option to be added to the promised utility as a state variable. Betting that this is intractable might be highly unfortunate

are uncertain yet persistent: the current buyer's market, say, foreshadows the buyer's market in the near future.

For concreteness, assume that, given an outside option worth v in round n , the outside option in round $n + 1$ is equal to v once again with probability λ , independent of all other aspects of the history. With complementary probability (and in the initial round), this outside option is drawn afresh from F . Therefore, the outside option follows a renewal process – a fairly specific, but common, way to model persistence, that includes the i.i.d. case when $\lambda = 0$ and the perfectly persistent case, when $\lambda = 1$.

The seller is only concerned about the frequency with which the buyer patronizes her and not about the circumstances in which he does. Thus, persistence only matters to the extent that it affects the buyer's willingness to honor his promise to visit the seller again. When $\lambda = 1$, the buyer is only willing to come if $q - p \geq v$, and so, if it is myopically optimal to visit the seller.²¹ If, say, $v > 1 - p$, autarky must prevail.

This extreme case brings to the foreground a constraint mentioned earlier: the buyer's *interim* incentive compatibility (see (6)). Given the outside option at hand, is it worth honoring the promise to the seller, if failing to do so leads in autarky?

This constraint curtails the future sales volume the buyer can credibly promise to the seller. Sales volume, however, is not the only channel through which to transfer surplus. An alternative, if less efficient, way to reward the seller is to let her pick a lower quality when the buyer's outside option is unattractive to save on cost. More precisely, for some parameters, the buyer-preferred equilibrium can involve a cut-off v^* such that

$$(1 - \delta)v^* + \delta(1 - \lambda)\mathbf{E}[\tilde{v}] = (1 - \delta)(1 - p) + \delta(1 - \lambda)B,^{22}$$

but equally safe.

²¹Of course, quality could vary over time, in which case the inequality must be adjusted accordingly.

²²To see this, first note that the value $O(v^*)$ from autarky given v^* solves

$$O(v^*) = (1 - \delta)v^* + \delta\lambda O(v^*) + \delta(1 - \lambda)\mathbf{E}[\tilde{v}],$$

and, given outside option v^* , the buyer is indifferent between autarky and obtaining quality $q = 1$ if

$$O(v^*) = (1 - \delta)(1 - p) + \delta\lambda O(v^*) + \delta(1 - \lambda)B.$$

Equating these two right-hand side terms gives the expression in the text.

making the buyer indifferent between honoring the promise and getting the *ex ante* payoff B as the continuation payoff and reneging when his outside option is equal to the cut-off, with $q(v^*) = 1$.²³ However, we may pick $q(v) < q(v^*)$ for $v < v^*$ (with $q(v)$ still high enough to satisfy the version of (6) that applies given commitment) to reduce the seller’s cost and at the same time boost her expected payoff, so as to ensure that, for all $v \leq v^*$,

$$(1 - \delta)cq(v) \leq \delta S(v),$$

where $S(v)$ is the seller’s continuation payoff when the buyer visits the seller given outside option v .²⁴

Flexible prices It is not the purpose of this paper to discuss why contractual prices do not always arise in buyer-seller relationships and when they fail to do so (see Macaulay, 1963). Plainly, here, the dynamics in the buyer-seller relationship arise because money is not available as a flexible tool to settle debts. If the buyer could commit to a bonus, a “tip” that is commensurate to the seller’s quality choice, first-best would be readily achievable, and the relationship would reduce to a repetition of identical transactions. Absent commitment, the size of the bonus is constrained by the loss that reversal to autarky entails, but the logic remains similar: money would be used first and foremost as a superior way of settling liabilities.

Note that the seller does not necessarily prefer a higher price, or the buyer a lower price. This is because the buyer, say, internalizes the cost of the future foregone opportunities to compensate the seller for the quality she provides when the price is low. To be obligated to someone is to be his slave, as the saying goes; paying a higher price up front is to free oneself.

One-time welcome coupons to lure first-time customers is a widespread practice that is often described as a way of acquainting the buyer with the seller’s product. Our model provides an alternative explanation for such discounts, as they mitigate

²³This is only one possibility, as depending on the parameters, $q(v^*) = 1$ might not be what the buyer prefers, and the surplus transfer via lower quality for lower outside options might not satisfy both the seller’s and the buyer’s incentive compatibility constraints.

²⁴Given persistence, $S(v) = \frac{1-\delta}{1-\delta\lambda}(p - cq(v)) + \frac{\delta(1-\lambda)}{1-\delta\lambda}F(v^*)(p - c\mathbf{E}[q])$.

the delay that comes with the consideration stage. However, such a discount is not the only way to deal with this delay. For instance, offering a discount to returning customers, which loyalty programs often do, reduces the future opportunity cost that makes the buyer reluctant to start a relationship, encouraging him to come sooner.

Incomplete Information Loyalty programs are meant to address not only moral hazard but also adverse selection: some customers are intrinsically more valuable and loyal than others, and calibrating quality with the status of the relationship allows sellers to screen them. Such incomplete information can be modeled similarly to persistence: with some probability, a customer is potentially a repeat customer, or “local,” whose characteristics are as described in the model; with complementary probability, the customer is a one-time customer that has no interest in purchasing again: a “tourist.” In the initial round, quality is constrained by the risk the seller faces that her customer is a tourist. In turn, this exacerbates the reluctance to kick off the relationship. Interesting dynamics might arise with a richer set of customer types that might help explain why quality would slowly grow over time as the seller learns that her customer is recurrent.

Competition Our model takes a drastic shortcut in modeling competition as an exogenous outside option. We note that after the consideration stage, the buyer’s visits occur in an i.i.d. fashion, which is at least consistent with the assumption that we have made on the outside option—an i.i.d. draw. Hence, it would not be difficult to model this outside option as an offering by an alternative firm, or a competitive fringe of such firms, with the customer’s preference for one over the other product being modeled as a random shock. However, doing so explicitly raises modeling choices (what do firms observe about each others’ offerings?) that open the door for phenomena (collusion among firms, in particular) that our model abstracts away. Board (2011) highlights the insider bias that arises in such interactions; such a bias would arise in our setting as well.

3 Incomplete Information

Henceforth, the buyer's outside option is private information. This does not affect realized payoffs, so their specification is not repeated.

A seller's history $h^{n-1} \in H^{n-1}$ is a sequence $(q_0, \dots, q_{n-1}) \in (\mathcal{Q} \cup \{\emptyset\})^n$, with the convention that $q_m = \emptyset$ whenever the buyer chooses $a_m = O$ in that round, that is, when the buyer does not visit the seller. A behavior strategy for the seller, then, is a sequence $\sigma^S := (\sigma_n^S)_n$ where σ_n^S is a probability transition from $H^{n-1} \times \{I, O\}$ into $\mathcal{Q} \cup \{\emptyset\}$, with the restriction that $\sigma_n^S(h^{n-1}, a) = \emptyset$ if, and only if, $a = O$.²⁵ A buyer's history $\hat{h}^{n-1} \in \hat{H}^{n-1}$ is a sequence $(v_0, q_0, \dots, v_{n-1}, q_{n-1}) \in (\mathcal{V} \times \mathcal{Q} \cup \{\emptyset\})^n$; that is, it specifies the outside option and the seller's quality whenever applicable. A behavior strategy for the buyer is a sequence $\sigma^B := (\sigma_n^B)_n$, where σ_n^B is a probability transition from $\hat{H}^{n-1} \times \mathcal{V} \rightarrow \{I, O\}$. Note that the buyer's choice in past rounds is omitted from the histories, as it is implied by the seller's quality choice. We use perfect Bayesian equilibrium as a solution concept. In addition to Assumption **A1**, we assume that always visiting the seller, independent of the outside option, is worse than autarky for the buyer, even when quality is maximum. This ensures that both visiting the seller and not visiting the seller are on the equilibrium path (hence, monitoring has "full support"). This ensures that a buyer cannot hope for his first-best payoff and that he chooses *Out* whenever $v = \bar{v}$.

Three remarks are in order. First, without loss of generality (as far as equilibrium payoffs are concerned, as well as equilibrium outcomes, in terms of public histories) attention can be restricted to buyer's strategies that are independent of past outside options, that is, that are measurable with respect to (the σ -algebra on) $H^{n-1} \times \mathcal{V}$. This follows from the Markovian (indeed, i.i.d.) structure on the process of the outside option (and the product structure of monitoring, see Fudenberg and Levine, 1994).

Second, because the outside option is unobservable, the cut-off structure of the buyer's strategy, which is a property of the extremal equilibria of interest under complete information, but not of all equilibria, is now a feature of any best reply

²⁵That is, for each $h^{n-1} \in H^{n-1}$, $\sigma_n^S(h^{n-1})$ is a distribution over $\mathcal{Q} \cup \{\emptyset\}$, and the probability assigned to any Borel set $A \subset \mathcal{Q} \cup \{\emptyset\}$ is a measurable function of h^{n-1} .




Figure 5: Comparison of equilibrium payoff sets when $(\delta, p, c) = (\frac{13}{20}, \frac{13}{25}, \frac{1}{4})$. The equilibrium payoff set for the observable (unobservable) case is shown in blue (black); the feasible and IR payoff set is shown in red. (B^*, S^*) denotes the payoffs in the buyer's first-best (on the boundary of the feasible payoff set), and (\bar{B}, \bar{S}) denotes the payoffs in the buyer-preferred equilibrium in the observable case.

of the buyer. That is, if $\sigma_n^B(h^{n-1}, v)$ assigns positive probability to *Out* for some v , then it assigns probability zero (resp., one) to *In* (*Out*) for all $v' < v$ ($v' > v$). For definiteness, we follow the convention that an indifferent buyer chooses *In*. Furthermore, at the cut-off, the buyer is indifferent between going to the seller or not.²⁶

Third, because the seller's chosen quality is public, it is still without loss to assume that any seller deviation triggers autarky. Hence, we focus on histories within which the seller has not deviated.

For concreteness, we still focus on the buyer's favorite equilibrium, though this now requires us to also solve for many other extremal equilibria that serve as possible continuations. Figure 5 compares the equilibrium payoff set in the complete information case to that in the incomplete information case.

3.1 Simple Equilibria

Assumption **A1** implies that the seller's quality cannot be independent of the history, except under autarky. If it were, then the buyer would act myopically, and the seller would not recoup her cost. More precisely, if the buyer is supposed to visit (w.p.p.) after h^{n-1} , and the seller is expected to pick $q > 0$, then the seller's

²⁶This property is definitely not shared by the buyer's favorite equilibrium under complete information, except in the initial stage.

continuation strategy $\sigma^S |_{(h^{n-1}, \cdot)}$ must depend on the buyer's decision in round n .

To gain some intuition of the trade-offs involved, it is useful to consider equilibria that can be represented by a finite-state automaton—even if those are suboptimal. Let us start with two-state automata. Our goal is not to conduct an exhaustive analysis of such equilibria but rather to focus on those parameters for which the (buyer) optimal two-state automaton shares interesting properties with the best (perfect Bayesian) equilibrium.²⁷

Call these states “ H ” and “ L ,” with the convention that the buyer's payoff B^H in state H is the higher of the two. Since this is the buyer's favorite two-state automaton, the seller's incentive compatibility must bind in state H . As is clear from the previous discussion, $q^H \neq q^L$. Since we focus on specific parameters, let us further assume that $S^H \geq S^L$, $q^H = 1$ and $q^L > 0$. That is, we consider parameters for which the higher quality is maximum, the lower is not zero, and the buyer's favorite state is also the seller's favorite; after all, while the cost is higher, the buyer is also, unsurprisingly, going to the seller more often in that state: $v^H > v^L$.²⁸

It is not hard to see that when in state H , play remains in state H whenever the buyer visits; similarly, when in state L , play remains in state L whenever the buyer fails to visit. More interesting are the switching probabilities r^H and r^L in states H and L when the buyer makes the other choice.

Perhaps surprisingly, it must be that $r^L = 1$: if the buyer visits the seller in state L , play transits to state H . The seller forgives the buyer regardless of his past conduct. To see this, suppose $r^L < 1$. Note that the buyer's payoff in that state can be written as

$$B^L = \mathbf{E}[\max\{v^L, \tilde{v}\}],$$

as a buyer that goes to the seller reaps the same payoff as a buyer whose outside

²⁷Transitions across states are implemented with the public randomization device. Hence, the prevailing state is common knowledge. In the best equilibrium (see Section 3.2), transitions are deterministic, so such a device is not necessary.

²⁸Specifically, the assumption is that $q^H = 1, q^L > 0$ and $S^H > S^L$ are satisfied provided $\mathcal{V} = [0, 1]$ and $F = \mathcal{U}[0, 1]$,

$$\frac{(\sqrt{5-4c}+3)c}{2(c+1)} < p < \frac{1}{2} \left(c + \sqrt{(4-3c)c} \right),$$

and δ is in the range described below; see (8) and the discussion that follows.




Figure 6: Best two-state automaton.

option makes him indifferent between visiting or not. Hence, by decreasing q^L and increasing r^L to keep v^L constant (the first change making him less likely to visit the seller, the second, more likely), we keep both the buyer's payoff B^L and so also B^H constant, yet we unambiguously increase S^L given that the seller supplies a lower quality, the buyer visits as often as before, and visits lead to more frequent transitions to the desirable state H ($S^H \geq S^L$). Since S^H also depends on the continuation payoff in state L , it increases as well. This implies that the seller's incentive compatibility condition in state H becomes slack, a contradiction. Loosely speaking, the seller's payoffs are aligned with the buyer's incentives.

This reasoning does not apply to the transitions that apply in state H : increasing r^H , the probability of switching to state L if the buyer does not come to the seller, while decreasing q^H to keep v^H constant does not affect the buyer's payoff B^H , as before, but it might no longer benefit the seller: switching more often to state L could be costly, since $S^L \leq S^H$. There is a trade-off between a lower quality cost and a lower continuation payoff, so r^H might well be interior.

Figure 6 illustrates the best two-state automaton.

To obtain a better sense of the choices involved, we illustrate this discussion with the following parameters. First, let the outside option be drawn from the uniform

distribution on the unit interval (that is, $\mathcal{V} = [0, 1]$, $F = \mathcal{U}[0, 1]$). Second, define

$$\varepsilon := \frac{\delta}{1 - \delta} - \frac{c}{(1 - p)(p - c)}, \quad (8)$$

The parameter ε (taken to be nonnegative) measures the extent to which the seller must be given rents to supply high quality. If $\varepsilon = 0$, the discount factor solves $\delta(1 - p)(p - c) = (1 - \delta)c$; then, the seller is barely willing to provide maximum quality to the buyer when he behaves myopically, that is, when he visits with probability $F(1 - p) = 1 - p$. When $\varepsilon > 0$, the buyer's cut-off must be raised accordingly. We can solve for the best two-state automaton in terms of this parameter, assuming it is small enough that higher-order terms are negligible.²⁹

Given that the myopic cut-off is insufficient repeat business for the seller, the best equilibrium features a slightly higher cut-off, namely,

$$v^H = 1 - p + \nu^H \varepsilon,$$

for some constant $\nu^H \in \mathbf{R}_+$.³⁰ To induce the buyer to visit with that frequency, there are two instruments at the seller's disposal: how often we transit to state L when the buyer fails to visit, and how low quality is in that state. As it turns out, they are used to an equal degree in the best equilibrium, in the sense that

$$r^H = \rho^H \sqrt{\varepsilon},$$

and

$$1 - q^L = \gamma^L \sqrt{\varepsilon}.$$

Taken together (“multiplicatively”), these values shift the buyer's cut-off away from the myopic cut-off by the desired amount. Of course, lower quality in state L also means a commensurately lower cut-off in state L , despite the future benefits that visiting the seller provide:

$$v^L = 1 - p - \nu^L \sqrt{\varepsilon}.$$

The buyer's first-best payoff is $\mathbf{E}[\max\{1 - p, \tilde{v}\}] = (1 + (1 - p)^2)/2$, leaving the seller with payoff $(1 - p)(p - c)$. Relative to this benchmark, some of the payoff is transferred

²⁹Even in this simple case, in which at the optimum $q^H = 1$ and $r^L = 1$, the exact values of r^H and q^L involve the root of an uninspiring polynomial of degree 14—hence the use of expansions.

³⁰These positive constants and those that follow are specified in Appendix B.

from the buyer to the seller in state H , which both parties nonetheless prefer to state L . Indeed, the buyer gets³¹

$$B^H = \frac{1 + (1 - p)^2}{2} - \beta^H \varepsilon, \quad B^L = \frac{1 + (1 - p)^2}{2} - \beta^L \sqrt{\varepsilon}.$$

while the seller reaps

$$S^H = (1 - p)(p - c) + c\varepsilon, \quad S^L = (1 - p)(p - c) - \sigma^L \sqrt{\varepsilon}.$$

While this specification is optimal (in an asymptotic sense) for a two-state automaton, the focus on two states only is restrictive. The first observation, inspired by the analysis of the observable case, is that a consideration stage benefits the buyer: at the outset, he does not need to make good on the seller's rents entailed by the high state. Hence, the buyer profits from the initial state I in which the relationship begins but that is left for good once he chooses to visit the seller. Unsurprisingly, this consideration stage involves a cut-off v^I below the buyer's first-best cut-off,

$$v^I = 1 - p - \nu^I \varepsilon,$$

and a payoff that exceeds that reaped in the good state,

$$B^I = \frac{1 + (1 - p)^2}{2} - \beta^I \varepsilon,$$

for some constants $\nu^I, \beta^I \in \mathbf{R}_+$.³² See Figure 7.

The abovementioned approach is not the only way the two-state automaton can be improved. Given that visits to the low state are both costly and inescapable, they should be chosen wisely. Waiting before passing judgment, that is, before moving to the low state, makes it possible to improve the statistical power of the test that such a transition involves, thus reducing inefficiency without lessening incentives. “Splitting”

³¹The buyer's shortfall is of order ε . In the observable case, the shortfall is of order ε^2 only: providing additional patronage of order ε involves giving up outside options exceeding the flow payoff the seller delivers by at most ε . Instead, when the outside option is unobservable, “punishments” cannot be tailored to the outside option and are sometimes carried out even when the outside option is well above the cut-off.

³²Note that the best such three-state automaton uses the optimal two-state automaton as a continuation equilibrium because a larger buyer's payoff in the high state implies a larger buyer's payoff in the initial state.




Figure 7: Automaton with consideration stage.

the high state into two states in which quality remains maximum naturally leads to an adjustment in the buyer's optimal unobserved cut-off, enabling a utility transfer from the buyer to the seller that does not directly involve the inefficiency that lower quality entails. Similarly, splitting the low state allows us to fine tune the quality drop and to maintain a balance between rewards and punishments.³³ Figure 8 illustrates such an automaton.

There is no reason to expect that improvements are confined to one, two or finitely many splittings. To describe the buyer's favorite equilibrium, it is necessary to move beyond finite automata and describe the equilibrium payoff set with greater generality.

3.2 The Buyer's Favorite Equilibrium

This section provides an answer to the following two questions: how does equilibrium behavior develop? Must the relationship end in autarky?

Our focus remains on the buyer-preferred equilibrium, but the answers to these questions do not depend on it: the same answers also hold for the seller-preferred

³³More formally, as will be clear in the next section, optimal incentive schemes require the buyer's marginal payoff to be a martingale, if possible, and this can only be achieved if both penalties and rewards are in the cards.




Figure 8: Automaton with consideration stage and information aggregation.

equilibrium and, indeed, for any equilibrium whose payoff vector lies on the Pareto frontier of the equilibrium payoff set.

The technical reason for this is that, owing to the one-sided structure of imperfect monitoring, the upper boundary of this equilibrium payoff set is self-generating (in the sense of Abreu, Pearce and Stacchetti, 1990). More precisely, let

$$E_\delta := \{(B, S) \in \mathbf{R}_+^2 \mid \text{there exists an equilibrium } \sigma \text{ s.t. } (B, S) \text{ is the payoff given } \sigma\}$$

denote the (closed, compact, convex) equilibrium payoff set given δ , and

$$\mathcal{D} = \{(B, S) \in E_\delta \mid (B, S') \in E_\delta \Rightarrow S' \leq S\}$$

be its upper boundary, that is, those vectors that cannot be improved for the seller, holding the buyer's payoff fixed. The Pareto frontier of E_δ is a subset of \mathcal{D} , but the closed curve \mathcal{D} also includes Pareto-inferior vectors to the extent that, here as in the observable case, the parties' interests are not entirely misaligned: if the buyer obtains his autarky payoff only, the seller cannot obtain more than hers.

Figure 9 makes clear why \mathcal{D} is self-generating. Here, $\mathcal{S}(B)$ is the highest equilibrium payoff for the seller consistent with the buyer obtaining payoff B ; that is, $(B, \mathcal{S}(B))$ indicates the point of the curve \mathcal{D} with abscissa B . Holding fixed the qual-




Figure 9: The curve \mathcal{D} is self-generating.

ity that is expected from the seller and the buyer's continuation payoffs B^O, B^I as a function of the buyer's choice of *Out* or *In*, an increase in the seller's continuation payoffs S^O, S^I increases her expected payoff while relaxing her incentive constraints. Doing so to the greatest extent possible drives the seller's payoff to its highest level, $S(B)$. However, because it is generally necessary to punish the buyer when he fails to visit the seller and because such punishments may hurt the seller, Pareto-inferior payoff vectors arise on this path. As we shall see, a subset of \mathcal{D} is self-generating, but it is not the Pareto frontier.³⁴

The same figure exhibits some features of the equilibrium payoff set that can be shown more generally (see Appendix C.1). The autarky payoff vector $(\mathbf{E}[\tilde{v}], 0)$ is an extreme point of E_δ and of the (strictly convex) plane curve \mathcal{D} . At the other end of \mathcal{D} lies the buyer's favorite equilibrium vector, (\bar{B}, \bar{S}) . The slope of \mathcal{D} is infinite at both extremities.

³⁴The lower boundary of the equilibrium payoff set is also self-generating for the same reasons. The lower boundary can be studied along the same lines as the higher boundary. Such a study is omitted here, although this boundary is represented in Figures 10 and 13 below.




Figure 10: Dynamics along the curve \mathcal{D} .

The infinite slope of \mathcal{D} at the autarky payoff plays such an important role in what follows that it deserves some comments. For the seller to receive a payoff of order $\varepsilon > 0$, the buyer must visit with commensurate probability, and so only when his outside option is within (an order) ε of his lowest possible outside option, \underline{v} . In other words, his outside option cannot be more than ε below the level at which he would be indifferent between visiting or not, and so he gains at most ε from visiting. To recap, he visits with a probability of (an order) ε for a benefit of at most ε ; so, he obtains only (at best) ε^2 more than from autarky, which is much less than the seller's ε . This implies that

$$\lim_{B \downarrow \mathbf{E}[\underline{v}]} \frac{S}{B - \mathbf{E}[\underline{v}]} = +\infty. \quad (9)$$

How does behavior evolve along the curve \mathcal{D} ? Figure 10 schematically illustrates the dynamics, with red arrows indicating the change in payoff following the action *In* and blue arrows indicating the change following *Out*. Figure 11 illustrates a sample path over time: it shows how quality and the buyer's payoff change as a function of the buyer's decision to go to the seller (red) or take the outside option (blue) when $(\delta, p, c) = (\frac{13}{20}, \frac{13}{25}, \frac{1}{4})$. As can be seen from both Figure 10 and 11, the relationship

experiences gradual downturns (in terms of payoffs to both parties and the quality provided by the seller) when the buyer breaks the seller’s trust by not visiting (indicated by blue arrows in Figure 10 and blue dots in Figure 11). However, the relationship instantaneously recovers (in terms of payoffs to both parties, and also the quality provided by the seller) if loyalty is shown. This can be seen from the jumps in payoffs and quality following a visit by the buyer indicated by red arrows in Figure 10 and red dots in Figure 11.

This difference in how downturns and upturns in the relationship play out is consistent with the structural asymmetry observed in many multitier loyalty programs (Bijmolt et al., 2018). Multitier loyalty programs, designed to build a loyal customer base, categorize customers into tiers based on past behavior. In particular, such programs offer better perks, the more frequent a customer visits. Many such loyalty programs such as those of airlines like *Air France* or hotel chains like *Marriott* display a structural asymmetry when it comes to promoting and demoting customers: if customers fail to be loyal over a given year, these programs offer a “soft landing” - i.e., the customer is demoted by at most one tier. On the other hand, there is no limit on how quickly a customer can reach the highest tier - in fact, promotion is typically applied instantaneously (rather than at the end of the evaluation period).

Whenever the buyer fails to visit the seller, his payoff drops. This drop need not be large: indeed, $B^O = B$ when $B = \bar{B}$: that is, failing to come to the seller leads to a strictly lower payoff only when $B \in (\mathbf{E}[\tilde{v}], \bar{B})$ (see Appendix C, Claim 12). Figure 12 illustrates the buyer’s continuation payoff as a function of B .

Here, we recognize the equivalent of the consideration stage: when $B = \bar{B}$, the buyer owes nothing to the seller, and the seller’s payoff is only positive because of the expected rents that are delivered once the relationship starts in earnest. This initial stage prevails until the buyer visits the seller.

The buyer is not necessarily rewarded for visiting the seller. Indeed, there is an upper bound, $\tilde{B} < \bar{B}$, on the value of B^I . Starting from $B \in (\tilde{B}, \bar{B})$, the buyer’s payoff drops even when he patronizes the seller. When he does, his payoff is stuck below \tilde{B} from that time onward: failing to visit leads to a drop, and visiting the seller




Figure 11: Sample path for $(\delta, p, c) = (\frac{13}{20}, \frac{13}{25}, \frac{1}{4})$.

leads to \tilde{B} at best. Therefore, the payoffs in the range $(\tilde{B}, \bar{B}]$ are transient. The continuation payoff from going to the seller, B^I , is bounded not only above but also below. Let \hat{B} denote the buyer's payoff at which the weighted sum of payoffs $B + S/c$ is maximized along \mathcal{D} . Then, unless a corner solution prevails (whereby even zero quality and $B^O = \mathbf{E}[\tilde{v}]$ do not suffice to drive down the buyer's payoff arbitrarily close to $\mathbf{E}[\tilde{v}]$, if $B^I = \hat{B}$), B^I is always at least \hat{B} (see Appendix C, Claim 11). A singular reappearance by the buyer leads to a jump in his payoff. As in the two-state example, the seller's interests are aligned with the buyer's incentives: considering any candidate lower value of B^I , an increase in B^I benefits the seller and motivates the buyer so that concomitantly decreasing the quality (if possible) leads to an improvement in the seller's payoff, keeping the buyer's payoff fixed.^{35,36}




Figure 12: Continuation payoffs, B^O (left) and B^I (right), as a function of B .

How does the seller's quality and the buyer's cut-off vary with the buyer's payoff B ? Unfortunately, we have been unable to formally establish the robust pattern

³⁵This is somewhat imprecise, as \hat{B} is not the buyer's payoff such that the corresponding seller's payoff is maximized. What matters is the rate of substitution between the buyer's and seller's payoffs, relative to the rate of substitution in quality, $1/c$ —hence, the definition of \hat{B} .

³⁶As we hinted, this lower bound can be lower than \hat{B} in the case that decreasing quality is no longer feasible given B . At that point, promise-keeping determines the lower bound. The buyer's payoff jumps in this case as well, albeit to a level determined by feasibility. Additionally, while Figure 13 suggests that $\hat{B} < \bar{B}$, this need not be the case. If, instead, $\hat{B} = \bar{B}$, then any buyer visit leads to the same “resetting” of the buyer's payoff. The case $\hat{B} < \bar{B}$ provides the “richest” dynamics; hence, we focus on it here. The analysis in Appendix C provides a systematic taxonomy of the equilibrium structure.




Figure 13: Range of continuation payoffs following *In*.

that emerges from numerical simulations.³⁷ Figure 14 provides an illustration of this pattern. The lower the buyer's payoff is, the lower the seller's quality. There is a cut-off, labeled $B^q > \mathbf{E}[\tilde{v}]$ here, such that $q < 1$ if, and only if, $B < B^q$.³⁸ For the parameters chosen here (the same as those used in Figure 12), the lower bound on quality, 0, is not binding; as the buyer's payoff approaches that obtained from autarky, quality becomes sufficiently low that, given the price p that the buyer must pay, the buyer finds it worthwhile to visit the seller only given very low outside options, despite the jump in continuation value that such a visit entails.

When quality is maximum, the cut-off decreases with the buyer's payoff. This is consistent with our intuition, gleaned from our four-state example, that utility is transferred from the buyer to the seller via a higher cut-off as the buyer's payoff decreases and that this transfer is incentivized by the more plausible and pressing

³⁷We were able to establish some elements of that pattern, but not all, even in the special case of $F \sim \mathcal{U}[0, 1]$, and so we refrain from stating fragmentary results beyond those that immediately follow from the results in Appendix C. These properties depend on the rate of change in concavity (the third derivative of $B \mapsto \mathcal{S}(B)$), so we cannot rule out that they depend on F .

³⁸Of course, we do not rule out that $q < 1$ for all $B \in [\mathbf{E}[\tilde{v}], \bar{B}]$.




Figure 14: Cutoff and quality.

threat of a decrease in the seller's quality. For buyer's payoffs such that the seller's quality is no longer maximum, the deleterious effect of lower quality on the buyer's incentive to visit the seller more than offsets the stronger incentives provided by a higher continuation payoff difference (between visiting and not), and the cut-off decreases all the way to its minimum value, \underline{v} , as the buyer's payoff approaches the autarky payoff.

Figures 11 and 12 show how the relationship evolves over time. Specifically, Figure 11 provides a representative sample path of visits, quality levels, and continuation buyer's payoff, as the round n progresses. Figure 12 shows how the buyer's continuation payoff varies with the buyer's decision to visit. The next proposition addresses the asymptotic properties of this random payoff process.

Proposition 3. *The buyer's payoff moves arbitrarily close to his autarky payoff and arbitrarily close to $\tilde{B} \geq \hat{B}$ infinitely often.*

That is, the relationship never settles or dissolves.³⁹ It should be clear that,

³⁹This volatility in the face of private information is consistent with traders in developing countries focusing on relationships that facilitate the flow of information thereby ensuring the regularity of trade (Fafchamps and Minten, 1999; Fafchamps, 2004).

conditional on not dissolving, buyer (and seller) values must keep changing given that the buyer must be incentivized to visit. Hence, the main step in establishing Proposition 3 involves showing that the relationship does not eventually crumble for reasons related to our earlier observation that very low equilibrium payoffs harm the buyer more than the seller relative to autarky payoffs (see (9)). This insight implies that it is always preferable to have continuation payoffs that exceed the payoffs from autarky, even if by only a minute amount. However, this does not suffice to establish the result: given that the buyer is less and less likely to come as his payoff approaches the payoff from autarky, further degradation only becomes more likely as the relationship deteriorates. What matters is the rate at which this degradation occurs. Formally, because the marginal utility of the buyer as a function of the promised utility to the seller is a martingale (that is, along the boundary, the slope of B as a function of S is the expectation of the slope in the next round), it must converge, and the infinite slope at the autarky payoff guarantees that the probability that autarky ultimately prevails must be zero.⁴⁰

4 Concluding Comments

This paper shows that the buyer-seller relationship exhibits two asymmetries. The first occurs over time and captures the dichotomy between the buyer's consideration stage and the loyalty loop stressed by the marketing literature. The second is between drastic improvements and gradual deteriorations in the quality supplied by the seller at any moment in time. While these asymmetries can be explained by a host of factors (learning about the seller's product and the buyer's attributes in the first case; in the second case, undesirable behavioral responses to sudden deteriorations or imperceptible improvements, strict convexities in the production technology as in

⁴⁰The proof is slightly more involved because the buyer's marginal utility is not a martingale over the entire range $[\mathbf{E}[\tilde{v}], \bar{B}]$. The martingale property is a standard result in dynamic agency models, but it relies on the constraint set of the optimization programming having a nonempty interior (so that some first-order conditions must hold with equality, which is equivalent to the martingale property). This is not necessarily the case here, as we have three variables to choose (the two continuation payoffs for the buyer and the seller's quality) and three constraints that might bind (the seller's incentive compatibility, promise-keeping, and quality being no larger than 1). Fortunately, this turns out not to be an issue for values of B close enough to the autarky payoff. See Appendix C.3 for details.

Li and Matouschek, 2013), they arise naturally under moral hazard and adverse selection. The first asymmetry is readily understood. At the start of the relationship, the buyer is not bound to the seller just yet and thus feels less of a need to patronize her. The second involves the trade-off between the two instruments to discipline the buyer, current quality and future continuation payoff. Whereas quality and punishment/reward are substitutes for the buyer’s incentives (a higher punishment/reward makes up for a lower quality), quality and reward are not substitutes for the seller’s payoff when their interests are aligned (both lower quality and higher rewards are desirable); hence, the reward is as high (and the quality as low) as is consistent with the alignment of interests and the range of feasible qualities.

We have not discussed whether more or less information is actually desirable for trade, an important issue given prevalent privacy concerns. As Figure 5 illustrates, the buyer is better off when the seller knows the outside options provided that we focus attention on the equilibrium preferred by the buyer. This is because costly punishments no longer need to be carried out under complete information. However, this ranking does not extend to all equilibria. In particular, in the equilibrium preferred by the seller, the buyer is better off when the seller has less information: the seller is loath to carry out such punishments and so suffers as well when incomplete information forces her to do so. This suggests that bargaining power plays a key role in this debate.

References

- Abreu, D., D. Pearce, and E. Stacchetti.** 1990. “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring.” *Econometrica*, 58: 1041–1063.
- Athey, S., and K. Bagwell.** 2001. “Optimal Collusion With Private Information.” *RAND Journal of Economics*, 32: 428–465.
- Bijmolt, T. H. A., M. Krafft, F. J. Sese, and V. Viswanathan.** 2018. “Multi-tier Loyalty Programs to Stimulate Customer Engagement.” In *Customer Engagement Marketing*, edited by R. Palmatier, V. Kumar and C. Harmeling, Springer Nature, Cham, Switzerland.
- Board, S.** 2011. “Relational Contracts and the Value of Loyalty.” *American Economic Review*, 101: 3349–3367.
- Calzolari, G., and G. Spagnolo.** 2017. *Relational Contracts, Procurement Competition, and Supplier Collusion*. working paper, DIW Berlin.
- Chassang, S.** 2010. “Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts.” *American Economic Review*, 100: 448–465.

- Dwyer, F. R., P. H. Schurr, and S. Oh.** 1987. "Developing Buyer-Seller Relationships." *Journal of Marketing*, 51: 11–27.
- Fafchamps, M.** 2004. *Market institutions in Sub-Saharan Africa - Theory and Evidence*. Cambridge, MA: MIT Press.
- Fafchamps, M., and B. Minten.** 1999. "Relationships and traders in Madagascar." *Journal of Development Studies*, 35: 1–35.
- Fuchs, W., and F. Lippi.** 2006. "Monetary Union with Voluntary Participation." *Review of Economic Studies*, 73: 437–457.
- Fudenberg, D., and D. K. Levine.** 1994. "Efficiency and Observability with Long-Run and Short-Run Players." *Journal of Economic Theory*, 62: 103–135.
- Fudenberg, D., Levine, D. K., and E. Maskin.** 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica*, 5: 997–1039.
- Halac, M.** 2012. "Relational Contracts and the Value of Relationships." *American Economic Review*, 102: 750–779.
- Hauser, C., and H. Hopenhayn.** 2008. *Trading Favors: Optimal Exchange and Forgiveness*. working paper, Collegio Carlo Alberto.
- Heide, J. B., and A. M. Weiss.** 1995. "Vendor consideration and switching behavior for buyers in high technology markets." *Journal of Marketing*, 59: 30–43.
- Levin, J.** 2003. "Relational Incentive Contracts." *American Economic Review*, 93: 835–857.
- Li, J., and N. Matouschek.** 2013. "Managing Conflicts in Relational Contracts." *American Economic Review*, 103: 2328–2351.
- Mauss, M.** 1925. *Essai sur le don: forme et raison de l'échange dans les sociétés archaïques*. L'année Sociologique, seconde série, 30–186
- Möbius, M.** 2001. *Trading Favors*. working paper, Harvard.
- Morgan, R. M., and S. D. Hunt.** 1994. "The Commitment-Trust Theory of Relationship Marketing." *Journal of Marketing*, 58: 20–38.
- Nedungadi, P.** 1990. "Recall and Consumer Consideration Sets: Influencing Choice without Altering Brand Evaluations." *Journal of Consumer Research*, 17: 263–276.
- Radner, R.** 1985. "Repeated Principal-Agent Games with Discounting." *Econometrica*, 53: 1173–1198.
- Spear, S. E., and S. Srivastava.** 1987. "On Repeated Moral Hazard with Discounting." *Review of Economic Studies*, 54: 599–617.
- Thomas, J., and T. Worrall.** 1990. "Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem." *Journal of Economic Theory*, 51: 367–390.
- Watson, J.** 1999. "Starting Small and Renegotiation." *Journal of Economic Theory*, 85: 52–90.
- Watson, J.** 2002. "Starting Small and Commitment." *Games and Economic Behavior*, 38: 176–199.

APPENDIX: FOR ONLINE PUBLICATION

A Proof of Propositions 1 and 2

Plainly, Proposition 1 is an immediate corollary of Proposition 2 (ignoring the issue of buyer incentive compatibility, as explained after Proposition 2), so that it suffices to establish the latter.

Before doing so, we introduce the last unspecified threshold, namely the minimum payoff that the seller actually delivers. To do so, let \tilde{v} be the solution to

$$\frac{\tilde{S}}{c} + \tilde{B} - \frac{1-\delta}{\delta}p = \int_{\tilde{v} \geq v} \tilde{v}F(d\tilde{v}) + \left(\frac{1-\delta}{\delta} + F(v) \right) v, \quad (10)$$

where

$$\tilde{B} := F(v^*(\tilde{S}))(q^*(\tilde{S}) - p) + \int_{\tilde{v} \geq v^*(\tilde{S})} \tilde{v} dF(\tilde{v}), \quad (11)$$

where v^* , q^* and \tilde{S} are defined in Proposition 2. \tilde{B} is simply the buyer's payoff $\mathcal{B}(\tilde{S})$. As the right-hand side of (10) is increasing in v , the solution is unique, if it exists. As will be clear, $\tilde{v} \in (\underline{v}, \bar{v})$. The lowest payoff the seller actually receives, independent of the minimum promise S , is

$$\underline{S} := \frac{(1-\delta)F(\tilde{v})}{1-\delta + \delta F(\tilde{v})}p. \quad (12)$$

Why \tilde{v} and \underline{S} are defined this way will become clear from the proof.

Our proof consists in two steps. First, we explicitly describe the value functions, given the policy function (*e.g.*, the functions q^* , v^* , S^I , S^O with domain $[0, p]$) specified in Proposition 2. Second, we verify the optimality equation.

It is instructive to derive the optimal q^* , v^* , taking only S^I , S^O as given. We note that, for $S \in [0, \tilde{S}]$, the buyer and seller's payoffs solve, for $q = q^*(S)$, $v = v^*(S)$, $B = \mathcal{B}(S)$,

$$\begin{aligned} S &= (1-\delta)F(v)(p - cq) + \delta F(v)\tilde{S} + \delta \bar{F}(v)S, \\ B &= (1-\delta) \left[F(v)(q - p) + \int_{\tilde{v} \geq v} \tilde{v}F(d\tilde{v}) \right] + \delta F(v)\tilde{B} + \delta \bar{F}(v)B, \end{aligned} \quad (13)$$

assuming that the buyer delivers S and no more. For $S \in [\tilde{S}, p]$, this system becomes

$$\begin{aligned} S &= (1-\delta)F(v)(p - cq) + \delta S, \\ B &= (1-\delta) \left[F(v)(q - p) + \int_{\tilde{v} \geq v} \tilde{v}F(d\tilde{v}) \right] + \delta B. \end{aligned} \quad (14)$$

Because the seller's incentive constraint is taken to bind for $S \leq \tilde{S}$, $(1-\delta)cq = \delta\tilde{S}$, and so (13) can be rewritten as

$$\begin{aligned} S &= (1-\delta)F(v)p + \delta \bar{F}(v)S, \\ B &= (1-\delta) \left[F(v)(q - p) + \int_{\tilde{v} \geq v} \tilde{v}F(d\tilde{v}) \right] + \delta F(v)\tilde{B} + \delta \bar{F}(v)B. \end{aligned} \quad (13')$$

(The seller's payoff is "as if" she shirked if the buyer came.) As is clear from these formulas, it is often convenient to work with quantiles $t := F(v)$ rather than outside options v .

From the first equation of (13'), it is immediate that

$$t = \frac{(1-\delta)S}{(1-\delta)p - \delta\tilde{S}}, \quad (15)$$

which confirms the formula for v^* in Proposition 2 (with the caveat that S must be replaced with $\max\{S, \tilde{S}\}$ in case some value $\tilde{S} \geq S$ proves more advantageous to the buyer, see below).

From the second equation of (13'), and eliminating q using $(1-\delta)cq = \delta\tilde{S}$,

$$\begin{aligned} B &= \frac{(1-\delta) \left[t \left(\frac{\delta\tilde{S}}{(1-\delta)c} - p \right) + \int_{\tilde{v} \geq F^{-1}(t)} \tilde{v}F(d\tilde{v}) \right] + \delta t\tilde{B}}{1 - \delta(1-t)} \\ &= \frac{S}{p} \left[\frac{\delta}{1-\delta} (\tilde{S}/c + \tilde{B}) - p \right] + \left(1 - \frac{\delta S}{(1-\delta)p} \right) \int_{\tilde{v} \geq F^{-1}\left(\frac{(1-\delta)S}{(1-\delta)p - \delta\tilde{S}}\right)} \tilde{v}F(d\tilde{v}), \end{aligned} \quad (16)$$

using (15) to eliminate t . The derivative of (16) w.r.t. S equals

$$\frac{\delta}{(1-\delta)p} \left[\tilde{S}/c + \tilde{B} - \int_{\tilde{v} \geq v} \tilde{v}F(d\tilde{v}) - \left(\frac{1-\delta}{\delta} + F(v) \right) v \right] - 1. \quad (17)$$

In turn, this expression is decreasing in $v = F(t)$,⁴¹ hence decreasing in t and so in S (cf. (15)). Therefore, the right-hand side (R.H.S.) of (16) admits a unique maximum in S , obtained by setting (17) to 0.⁴² Hence, the expression for \tilde{v} given by (10), and thus, given (15), the seller receives always at least

$$S \geq \underline{S},$$

as defined in (12). Hence, (16) with S replaced with $\max\{S, \underline{S}\}$ gives the buyer's payoff given that the seller is promised at least S , for $S \leq \tilde{S}$.

As will be clear, $\underline{S} < \tilde{S}$, given that the buyer's favorite promise to the seller if he is bound to give \tilde{S} once he comes is less than \tilde{S} . Alongside (15) and the incentive constraint that gives $q = \frac{\delta}{1-\delta} \frac{\tilde{S}}{c}$, this completes the description of the policy and payoffs for $S \leq \tilde{S}$. We note that, given the formula for \tilde{S} given in (7), q is either one (when $\tilde{S} = \frac{1-\delta}{\delta}c$); else, it is a constant strictly lower than one.

Turning our attention to $S \in [\tilde{S}, p]$, (14) immediately simplifies to

$$\begin{aligned} S &= F(v)(p - cq), \\ B &= F(v)(q - p) + \int_{\tilde{v} \geq v} \tilde{v}F(d\tilde{v}). \end{aligned} \quad (18)$$

⁴¹Its derivative w.r.t. v is $-\frac{\delta}{(1-\delta)p} \left(\frac{1-\delta}{\delta} + F(v) \right) < 0$.

⁴²That is, provided a maximum exists. For this, note from (13') that, if $S = 0$, $F(v) = 0$ and so $B = \mathbf{E}[\tilde{v}]$. Yet, it is readily verified that $\mathcal{B}(\tilde{S}) > \mathbf{E}[\tilde{v}]$, and so it must be that $\mathcal{B}(S)$ is increasing for S small enough. At \tilde{S} , the derivative equals $\frac{1-p-F^{-1}\left(\frac{1-\delta}{\delta}\frac{c}{p-c}\right)}{p-c}$ when $\tilde{S} = \frac{1-\delta}{\delta}c$. But $1-p-F^{-1}\left(\frac{1-\delta}{\delta}\frac{c}{p-c}\right) < 0 \Leftrightarrow \delta F(1-p)(p-c) < (1-\delta)c$, which follows from **A1**. Similar computations yield a negative derivative at \tilde{S} when $\tilde{S} = \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$.

The first equation gives

$$q = \frac{1}{c} \left[p - \frac{S}{F(v)} \right], \quad (19)$$

and so, using as before $t = F(v)$, the second equation becomes

$$B = p \frac{1-c}{c} t - \frac{S}{c} + \int_{\tilde{v} \geq F^{-1}(t)} \tilde{v} F(d\tilde{v}), \quad (20)$$

Assuming for now that q , as given by (19) is in $(0, 1)$, we can then derive the optimal value of v . The derivative w.r.t. t of (20) is

$$p \frac{1-c}{c} - F^{-1}(t),$$

a function decreasing in t , and so B , as given by the R.H.S. of (20) is concave in t , with maximum given by

$$t = \min \left\{ 1, F \left(p \frac{1-c}{c} \right) \right\} =: F(\kappa),$$

where we recall that $\kappa := \min \{ \bar{v}, p \frac{1-c}{c} \}$. Hence,

$$v = \kappa. \quad (21)$$

In turn, (19) gives

$$q = \frac{1}{c} \left[p - \frac{S}{F(\kappa)} \right], \quad (22)$$

and (20) yields

$$B = p \frac{1-c}{c} F(\kappa) - \frac{S}{c} + \int_{\tilde{v} \geq \kappa} \tilde{v} F(d\tilde{v}), \quad (23)$$

We note that the threshold is constant, but quality and payoff decrease linearly in S . Recall, however, that we assumed $q \in (0, 1)$ (see (19)), or to put it differently, the value of S might be so small that the R.H.S. of (19) exceeds 1 when evaluated at $t = F(\kappa)$. In fact,

$$\frac{1-\delta}{\delta} c < \frac{p}{\frac{1}{F(\kappa)} + \frac{\delta}{1-\delta}} \Leftrightarrow \frac{1}{c} \left[p - \frac{S}{F(\kappa)} \right] \Big|_{S=\frac{1-\delta}{\delta} c} > 1,$$

implying that (22) does not hold at values of S right above \tilde{S} if $\tilde{S} = \frac{1-\delta}{\delta} c$. If instead $\tilde{S} = \frac{p}{\frac{1}{F(\kappa)} + \frac{\delta}{1-\delta}}$ then (22) satisfies $q \leq 1$.

If $\tilde{S} = \frac{1-\delta}{\delta} c$, given the concavity of B in t established in (20), it follows that it is then optimal to set t as high as is consistent with $q \leq 1$, that is, $q = 1$ and so v solves $\frac{1}{c} \left[p - \frac{S}{F(v)} \right] = 1$, or

$$v = F^{-1} \left(\frac{S}{p-c} \right), \quad (24)$$

with (20) giving

$$B = \frac{1-p}{p-c} S + \int_{\tilde{v} \geq F^{-1}(\frac{S}{p-c})} \tilde{v} F(d\tilde{v}). \quad (25)$$

In that case, q is constant, and v and B decrease in S .⁴³

Finally, as is clear from (19), the constraint $q \geq 0$ might bind if S is too large, given $v = \kappa$. This is the case, given this equation, if $S \geq pF(\kappa)$. In that case, $q = 0$, and we immediately get $v = F^{-1}(S/p)$ from that equation; from the second line of (18), it follows that

$$B = \int_{\tilde{v} \geq F^{-1}(S/p)} \tilde{v} dF(\tilde{v}) - S \quad (26)$$

in that range.

As is clear from this discussion, if $\tilde{S} = \frac{1-\delta}{\delta}c$, then there is a range $[\tilde{S}, F(\kappa)(p-c))$ over which $q = 1$, and v, B are given by (24)–(25). For $S \in [F(\kappa)(p-c), F(\kappa)p]$, v, q and B are given by (21)–(23). If instead $\tilde{S} = \frac{p}{\frac{1}{F(\kappa)} + \frac{\delta}{1-\delta}}$, the first range does not exist, and for all $S \in [\tilde{S}, F(\kappa)p]$, v, q and B are given by (21)–(23).

Finally, for $S \in [F(\kappa)p, p]$, the solution is $q = 0$, $v = F^{-1}(S/p)$, with payoff B given by (26).

It is readily verified that, in each of the four cases (according to the two possible values of \tilde{S} and κ), the resulting function \mathcal{B} is globally concave on $[0, p]$.

This completes the first step, namely, the derivation of v, q and B , as described in Proposition 2, taking as given that $S^O = S$ and $S^I = \max\{S, \tilde{S}\}$. Using the candidate $\mathcal{B}(\cdot)$ that we have derived, we now verify the principle of optimality. Consider then maximizing

$$\begin{aligned} \hat{\mathcal{B}}: [0, p] &\rightarrow \mathbf{R} & (\hat{\mathcal{P}}) \\ S &\mapsto \sup (1 - \delta) \left[\int_A (q(v) - p)F(dv) + \int_{A^c} vF(dv) \right] + \delta \int_{\mathcal{V}} \mathcal{B}(S(v))F(dv), \end{aligned}$$

over $A \subset \mathcal{V}$, $S: \mathcal{V} \rightarrow \mathbf{R}_+$, and $q: \mathcal{V} \rightarrow \mathcal{Q} \cup \{\emptyset\}$ (all functions of S) subject to (4) and (5).

Our goal is to show that $\hat{\mathcal{B}} = \mathcal{B}$. The first observation is that the candidate payoff functions $\mathcal{B}(\cdot)$ that we have derived (cf. (16), (23), (25) and (26)) are concave in S on their domain; further, it is easy to see that the function $\mathcal{B}(\cdot): [0, p] \rightarrow \mathbf{R}$ defined by patching them together according to their domains is concave as well (in fact, $\mathcal{B}(\cdot)$ is continuously differentiable on $[0, p]$), while the constraints (4) and (5) are linear in $S(\cdot)$. Hence, $\hat{\mathcal{P}}$ is a concave programme, and we can take $S(\cdot)$ to be constant on the sets A, A^c . Similarly, replacing $q(\cdot)$ on A by $\mathbf{E}[q(\tilde{v}) \mid \tilde{v} \in A]$ does not affect the objective or (5), and certainly satisfies (4), because (4) must hold for all $v \in A$. Finally, replacing A with a set $[\underline{v}, v]$ such that $\int_{\underline{v}}^v F(d\tilde{v}) = \int_A F(d\tilde{v})$ does not affect (4) and (5), and cannot decrease the objective. Hence, the problem reduces to finding, for every $S \in [0, p]$, a pair $S^I, S^O \in [0, p]$, $q \in \mathcal{Q}$, and $v \in \mathcal{V}$. The constraints being equalities or weak inequalities, and being continuous in v, q and S^I, S^O , as is the objective, the maximum is achieved. Finally, it is also clear that $\hat{\mathcal{B}}$ must be non-decreasing in S , since (5) is an inequality. Let $\hat{S} := \max\{S \mid \hat{\mathcal{B}}(S) = \mathcal{B}'(0)\}$ denote the highest promise that entails no loss to the buyer (well-defined by the theorem of the maximum). Plainly, either $S^O = \hat{S}$, or (5) binds with equality (decrease S^O otherwise).

⁴³The latter property isn't apparent from (25): but differentiating the R.H.S. gives $\frac{1-p}{p-c} - \frac{1}{p-c}F^{-1}\left(\frac{S}{p-c}\right)$, which in light of (24) is negative provided $v > 1 - p$; since $1 - p$ is the myopic threshold given that $q = 1$, this is necessarily the case given **A1**.

By concavity of \mathcal{B} also, either S^I is such that (a) (4) binds, $q > 0$, $v > \underline{v}$, and $S^I > S^O$,⁴⁴ or (b), we can take $S^I = S^O$.

On an interval of values such that case (a) applies,⁴⁵

$$\hat{\mathcal{B}}(S) = \max_{S^O, S^I, q, t} \left\{ (1 - \delta) \left(t(q - p) + \int_t^1 F^{-1}(t') dt' \right) + \delta t \mathcal{B}(S^I) + \delta(1 - t) \mathcal{B}(S^O) \right\},$$

such that

$$\frac{1 - \delta}{\delta} cq = S^I, \quad (27)$$

and

$$S \leq (1 - \delta)tp + \delta(1 - t)S^O, \text{ with equality if } S^O > \hat{S}. \quad (28)$$

where, as before, $t = F(v) \in [0, 1]$, and the last equation (promise-keeping) uses the preceding equality to eliminate quality. On an interval of values such that (b) applies,

$$\hat{\mathcal{B}}(S) = \max_{S', q, t} \left\{ (1 - \delta) \left(t(q - p) + \int_t^1 F^{-1}(t') dt' \right) + \delta \mathcal{B}(S') \right\}, \quad (29)$$

such that

$$\frac{1 - \delta}{\delta} cq \leq S', \text{ or } t = 0, \quad (30)$$

and

$$S \leq (1 - \delta)t(p - cq) + \delta S'. \quad (31)$$

We note that either (30) or (31) binds, without loss (decrease S' otherwise).

First, let us consider values, if any, for which incentive compatibility does not bind in the determination of $\hat{\mathcal{B}}$. Then, case (b) must apply, and (31) must bind. Inserting this constraint in the objective, by eliminating q , yields

$$\hat{\mathcal{B}}(S) + \frac{S}{c} = \max_{S', t} \left\{ (1 - \delta) \frac{1 - c}{c} pt + (1 - \delta) \int_t^1 F^{-1}(t') dt' + \delta \left(\mathcal{B}(S') + \frac{S'}{c} \right) \right\}. \quad (32)$$

Note that the R.H.S. is concave in S' , since \mathcal{B} is, and strictly concave unless $S' \in [\tilde{S}, F(\kappa)p]$ when $\tilde{S} = \frac{p}{\frac{1}{F(\kappa)} + \frac{\delta}{1 - \delta}}$; ($S' \in [F(\kappa)(p - c), F(\kappa)p]$ if $\tilde{S} = \frac{1 - \delta}{\delta} c$). If $q \in (0, 1)$, so that first-order conditions with respect to S' must apply, it holds that

$$\mathcal{B}'(S') + \frac{1}{c} = 0. \quad (33)$$

and so indeed S' must belong to the interval in which \mathcal{B} is affine. Hence, for S in this interval, choosing $S' = S$ satisfies the first-order conditions with respect to S' , and the optimality of t and $q \in (0, 1)$ as described in step 1 follows. Since this is a relaxed programme, yet $\hat{\mathcal{B}}(S) \leq \mathcal{B}(S)$ in this

⁴⁴Otherwise, increase S^I and decrease S^O , keeping the R.H.S. of (5) constant; if $v = \underline{v}$, S^I is irrelevant to $\hat{\mathcal{P}}$.

⁴⁵The maximum theorem ensures that intervals are without loss of generality.

⁴⁶Unless $S' = 0, p$, but the derivatives of $\mathcal{B}(\cdot)$ at these extreme values imply the inequality that is opposite to the first-order condition.

interval, it follows that $\hat{\mathcal{B}}(S) = \mathcal{B}(S)$ over this interval.

Still ignoring the incentive compatibility constraint, if the solution to $\hat{\mathcal{B}}$ involves $q = 1$, we eliminate t using (31), and get as objective

$$\hat{\mathcal{B}}(S) = \max_{S'} \left\{ \frac{1-p}{p-c} (S - \delta S') + (1-\delta) \int_{\frac{S-\delta S'}{(1-\delta)(p-c)}}^1 F^{-1}(t') dt' + \delta \mathcal{B}(S') \right\},$$

with derivative w.r.t. S' equal to

$$-\delta \frac{1-p}{p-c} + \frac{\delta}{p-c} F^{-1} \left(\frac{S - \delta S'}{(1-\delta)(p-c)} \right) + \delta \mathcal{B}'(S'), \quad (34)$$

a decreasing function in S' , and so the objective is strictly concave in S' , so that, if a maximizer exists, it is unique. Yet, evaluating the first-order condition at $S' = S$ gives⁴⁷

$$\mathcal{B}'(S) = \frac{1-p}{p-c} - \frac{1}{p-c} F^{-1} \left(\frac{S}{p-c} \right),$$

which indeed holds for $S \in [\tilde{S}, F(\kappa)(p-c)]$ if $\tilde{S} = \frac{1-\delta}{\delta}c$ (compare (25)). Finally, if $q = 0$, eliminating t the same yields

$$\hat{\mathcal{B}}(S) = \max_{S'} \left\{ \delta S' - S + (1-\delta) \int_{\frac{S-\delta S'}{(1-\delta)p}}^1 F^{-1}(t') dt' + \delta \mathcal{B}(S') \right\},$$

with derivative

$$\delta + \frac{\delta}{p} F^{-1} \left(\frac{S - \delta S'}{(1-\delta)p} \right) + \delta \mathcal{B}'(S').$$

This is decreasing in S' . Hence, $\hat{\mathcal{B}}$ is concave in S' in this case as well. Setting this derivative to 0 and evaluating at $S' = S$ gives

$$\mathcal{B}'(S) = -1 - \frac{1}{p} F^{-1} \left(\frac{S}{p} \right).$$

Comparing to (26), it follows that $\hat{\mathcal{B}} = \mathcal{B}$ on the range $[F(\kappa)p, p]$. Hence, we have shown that our policy is optimal for $S \geq \tilde{S}$.

Consider now $S \leq \tilde{S}$. First, we note that, upper bounding \mathcal{B} by (23) on the range $[F(\kappa), p]$ (that is, ignoring the constraint $q \geq 0$ that binds on this range, so that \mathcal{B} remains affine with slope $-1/c$ for these values too), it follows from the previous analysis that, in case (b), and still ignoring the incentive compatibility constraint, we may take $S' \in [0, \tilde{S}]$ in case $\tilde{S} = \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$, and $S' \in [0, F(\kappa)(p-c)]$, with $q = 1$ if $S' \in (\tilde{S}, F(\kappa)(p-c))$ in case $\tilde{S} = \frac{1-\delta}{\delta}c$. Indeed, if $q < 1$, (33) shows that S' belongs to the affine range of \mathcal{B} , and we can then lower S' without loss. If $q = 1$ and $S' \in (\tilde{S}, F(\kappa)(p-c))$ when $\tilde{S} = \frac{1-\delta}{\delta}c$, (34) must hold, but given the formula for \mathcal{B} in that range (cf.

⁴⁷Again, the cases $S' = 0, p$ are readily dismissed.

(25)), it follows that

$$F^{-1}\left(\frac{S - \delta S'}{(1 - \delta)(p - c)}\right) = F^{-1}\left(\frac{S'}{p - c}\right) \text{ or } S = S',$$

which isn't possible given $S \leq \tilde{S} < S'$. Hence, in all cases, we may take $S' \in [0, \tilde{S}]$. Given that this is true when case (b) is relaxed (ignoring incentive compatibility, upper bounding the continuation payoff for $S' \geq F(\kappa)$), the same conclusion holds for the original problem.

If $\tilde{S} = \frac{1-\delta}{\delta}c$, then if $q < 1$, $S' \geq F(\kappa)(p - c)$, yet, as we saw, without loss, $S' \geq \frac{1-\delta}{\delta}c$. Hence, in that case, $q = 1$. But then incentive compatibility in case (b) requires $S' = \frac{1-\delta}{\delta}c$. This gives us one candidate value $\mathcal{B}(S)$ to compute in that case. This value is readily shown to be less than $\mathcal{B}(S)$ given by (16).

If $\tilde{S} = \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$, and $q = 1$ for $S < \tilde{S}$, given that the solution to the relaxed problem gives $S' \leq \frac{1-\delta}{\delta}c$ without loss, we then have $(1 - \delta)c = \delta \frac{1-\delta}{\delta}c > \delta \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$, and so incentive compatibility would fail. Hence, either $q < 1$, but then necessarily $S' = \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$ in the relaxed problem (because we must have $\mathcal{B}'(S') = -1/c$), or ignoring incentive compatibility in case (b) is unwarranted. If $q < 1$ and $S' = \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$, then first-order conditions in (32) yields $t = F(\kappa)$, and so (31) gives, for the relaxed problem,

$$q = \frac{1}{c} \left(p - \frac{S - \delta \tilde{S}}{(1 - \delta)F(\kappa)} \right),$$

so that the seller's cost of quality is equal to

$$(1 - \delta)cq = (1 - \delta)p - \frac{S - \delta \tilde{S}}{F(\kappa)} \leq \delta \tilde{S},$$

with strict inequality unless $S = \tilde{S}$. Hence, incentive compatibility binds in this case as well. To conclude, in case (b), when $\tilde{S} = \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$, incentive compatibility must bind for $S < \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$, and either $S' = \frac{p}{\frac{1}{F(\kappa)} + 1 - \delta}$, $t = F(\kappa)$, and $q = \left(\frac{p}{\frac{1}{F(\kappa)} + 1 - \delta} \right) / \left(\frac{1-\delta}{\delta}c \right) < 1$, or $q = 1$ and $S' = \frac{1-\delta}{\delta}c$. This gives us two candidate values $\mathcal{B}(S)$ to compute in that case. These two values are readily shown to be less than $\mathcal{B}(S)$ given by (16).

Hence, when $S < \tilde{S}$, case (a) must apply. We start by considering values of S such that promise-keeping holds, if any. Using (27) and (28) to eliminate S^O and q yields

$$\begin{aligned} \hat{\mathcal{B}}(S) &= \max_{t, S^I} \left\{ \delta t \frac{S^I}{c} - (1 - \delta)tp + (1 - \delta) \int_t^1 F^{-1}(t') dt' \right. \\ &\quad \left. + \delta(1 - t)\mathcal{B}\left(\frac{S - (1 - \delta)tp}{\delta(1 - t)}\right) + \delta t \mathcal{B}(S^I) \right\}. \end{aligned} \quad (35)$$

Differentiating with respect to S^I , and recalling that $q \leq 1 \Leftrightarrow S^I \leq \frac{1-\delta}{\delta}c$ yields that S^I is the smaller of $\frac{1-\delta}{\delta}c$ and the solution to $1/c + \delta \mathcal{B}'(S^I) = 0$,⁴⁸ that is, $S^I = \tilde{S}$. Note also that (35) is

⁴⁸Recall that in case (a), $v > \underline{v}$, that is, $t > 0$.

concave in t , and thus admits at most a unique interior critical point, which is then a maximum. Equivalently, given promise-keeping, there is at most a unique S^O satisfying the first-order condition w.r.t. S^O , a maximum as well. Replacing t with S^O given promise-keeping yields

$$\begin{aligned} \hat{B}(S) = & \max_{S^O} \left\{ \frac{(\delta S^O - S) (\delta \tilde{S} - (1 - \delta)cp)}{\delta c S^O - (1 - \delta)cp} + (1 - \delta) \int_{\frac{S - \delta S^O}{(1 - \delta)p - \delta S^O}}^1 F^{-1}(t') dt' \right. \\ & \left. + \delta \frac{S - \delta S^O}{(1 - \delta)p - \delta S^O} \mathcal{B}(S^I) + \delta \frac{(1 - \delta)p - S}{(1 - \delta)p - \delta S^O} \mathcal{B}(S^O) \right\}. \end{aligned}$$

Taking first-order conditions w.r.t. S^O , and using (17) to evaluate $\mathcal{B}'(\cdot)$, establishes that $S^O = S$ is a critical point of the R.H.S., and thus the maximizer of our problem. This concludes the proof.

B Finite-State Automata: Missing Details

Here, we provide the specific values for the constants mentioned in Section 3.1. Details for their derivation are available upon request.

$$B^H = \frac{1 + (1 - p)^2}{2} - p\varepsilon/\gamma, \quad B^L = \frac{1 + (1 - p)^2}{2} - (1 - p)\tau\sqrt{\varepsilon/\gamma},$$

and

$$v^H = 1 - p + \varepsilon/\gamma, \quad v^L = 1 - p - \tau\sqrt{\varepsilon/\gamma},$$

as well as

$$S^H = (1 - p)(p - c) + c\varepsilon, \quad S^L = (1 - p)(p - c) - \frac{(p - c)^2 - p(1 - p)c}{p} \tau\sqrt{\varepsilon/\gamma},$$

where

$$\tau := \sqrt{\frac{2(p - c)p(1 + c - p)}{p^2 + (2 - p)pc - 2c^2}}, \quad \gamma := \frac{p(1 - p) - (p - c)^2 - (1 - p)^2c}{(1 - p)c}.$$

Finally,

$$q^L = 1 - \frac{p}{p - c} \tau\sqrt{\varepsilon/\gamma}, \quad r^H = (1 - p)(p - c)\tau\sqrt{\varepsilon/\gamma}.$$

The third state, corresponding to the consideration stage, involves payoff and threshold given by

$$v^I = 1 - p - \frac{c}{1 - p} \varepsilon/\gamma, \quad B^I = \frac{1 + (1 - p)^2}{2} - c\varepsilon/\gamma.$$

C Incomplete Information: Proofs for Section 3.2

C.1 The Equilibrium Payoff Set E_δ : Preliminaries

Throughout, we renormalize the buyer's payoff to be $B - \mathbf{E}[\tilde{v}]$, so that its minimum value is 0. Rather than study the maximum seller's payoff as a function of the buyer's payoff, we study

$$\begin{aligned} \phi: [0, \infty) &\rightarrow \mathbf{R} \\ B &\mapsto \sup_{\{(B', S') \in E_\delta \mid B' = B\}} \{B' + S'/c\}, \end{aligned}$$

with $\phi(B) = -\infty$ if no equilibrium exists that gives B to the seller. We let \bar{B} denote the maximum over B such that $\phi(B) > -\infty$.

Fixing an equilibrium, given the buyer's equilibrium payoff B after some history, we denote by B^I, B^O the buyer's continuation payoff according to whether the buyer chooses *In* or *Out*, and $S(B)$ (resp., $S(B^I), S(B^O)$) for the corresponding seller's payoff.

We start by studying the behavior of $\phi(\cdot)$ for values of B that are close to 0. Recall that the buyer's outside option is distributed according to F , with support $[0, 1]$, and that the distribution F is assumed to be twice continuously differentiable, with density f bounded away from 0.

Claim 1. *It holds that $\phi(0) = 0$.*

Proof. Suppose instead that $\phi(0) > 0$, i.e., there exists an equilibrium in which the seller gets more than 0 yet the buyer gets just $\mathbf{E}[\tilde{v}]$. Consider the first round n in which the buyer comes to the seller with positive probability (“wpp”). Note that $n < \infty$, for otherwise $\phi(0) = 0$. Second, the payoff of the buyer from round n onward is $\mathbf{E}[\tilde{v}]$, for his reward is $\mathbf{E}[\tilde{v}]$ in each round until round n . Hence, without loss, take $n = 1$.

Next, suppose that the buyer comes with probability one (“wp1”) in round 1. Then his payoff from coming is $\mathbf{E}[\tilde{v}]$. Consider instead the strategy of the buyer that comes to the seller in round n if, and only if, $\tilde{v} > v$ for some arbitrary $v \in (0, 1)$; if he comes, he follows the equilibrium strategy; if he does not come, he never does in the future either. Note that conditional on coming, he gets $\mathbf{E}[\tilde{v}]$, as explained (the current value is irrelevant, since he then comes). Conditional on not coming, he gets strictly more (since he gets a strictly higher (conditional expectation of his) reward today, and $\mathbf{E}[\tilde{v}]$ from tomorrow onward). Hence, this is a profitable deviation.

Hence, the buyer comes with probability less than 1 in round 1. This implies that there exists some threshold v such that he comes if, and only if, $\tilde{v} < v$. Note that knowing his outside option is strictly valuable to the buyer (since his reward from coming is strictly increasing in v , and his continuation payoff conditional on not coming is independent of it). Yet, conditional on not knowing his outside option, he can get $\mathbf{E}[\tilde{v}]$ by never coming. A contradiction. \square

Claim 2. *It holds that, for all $B \in (0, \bar{B}]$, $\phi(B) > 0$.*

Proof. Immediate, as if the buyer gets more than 0, he must go to the seller on a set of histories with positive probability. A strategy available to the seller is to always choose zero quality, which then secures a strictly positive payoff. \square

Note that ϕ is weakly concave (being the upper boundary of a convex set), hence continuous—in fact, with left- and right-derivatives everywhere.

Claim 3. *It holds that*

$$\lim_{B \uparrow \bar{B}} \frac{\phi(\bar{B}) - \phi(B)}{\bar{B} - B} = -\infty.$$

Further, for $B = \bar{B}$, we may set $B^O = \bar{B}$.

Proof. The fact that $B^O = \bar{B}$ whenever $B = \bar{B}$ is immediate: if not, increase B^O , keeping B^I and q constant. Because B^O does not affect the seller's incentives, this change must benefit the buyer (as he can still come to the seller if he'd like to, and if he doesn't, his continuation pay off is higher).

To show that the slope of ϕ is infinite at \bar{B} , consider setting $B^O = (1 - \alpha)\bar{B}$, for a range of (small) $\alpha > 0$. Indifference of the buyer at the critical cut-off v is equivalent to⁴⁹

$$(1 - \delta)(q - p - v) = \delta((1 - \alpha)B - B^I).$$

This implies that $(1 - \delta)dv = \delta\bar{B}d\alpha$ at $\alpha = 0$. The buyer's payoff changes by

$$dB = -\delta\bar{F}(v)\bar{B}d\alpha = -(1 - \delta)\bar{F}(v)dv,$$

while the seller's payoff change by

$$dS = F(v)((1 - \delta)(p - cq) + \delta S^I)dv - \delta\bar{F}(v)\bar{B}S'(\bar{B})d\alpha,^{50}$$

or, combining and evaluating at $\alpha = 0$,

$$dS = -\frac{F(v)}{(1 - \delta)\bar{F}(v)}((1 - \delta)(p - cq) + \delta S^I)dB + S'(\bar{B})dB,$$

or

$$S'(\bar{B}) = \left. \frac{dS}{dB} \right|_{B=\bar{B}} = -\frac{F(v)}{(1 - \delta)\bar{F}(v)}((1 - \delta)(p - cq) + \delta S^I) + S'(\bar{B}),$$

which implies that $\lim_{B \rightarrow \bar{B}} S'(B) = \pm\infty$ (it cannot be that $v_n \rightarrow \underline{v}$ for a sequence $B_n \rightarrow \bar{B}$, as otherwise \bar{B} is the autarky payoff). By concavity, $S'(\bar{B}) = -\infty$. Since $\phi(B) = B + S/c$, it follows that also

$$\lim_{B \uparrow \bar{B}} \frac{\phi(\bar{B}) - \phi(B)}{\bar{B} - B} = -\infty.$$

□

⁴⁹For α small, the buyer must come to the seller for $v \sim \underline{v}$, while he does not when $v \sim \bar{v}$ because $\bar{v} > 1 - p$ ($B^O = \bar{B}$ without loss if the buyer comes w.p.1, and so if he were to come when $v = \bar{v}$ it would have to be that $B = (1 - \delta)(q - p) + \delta B^I \geq (1 - \delta)\bar{v} + \delta\bar{B}$, which is impossible since $1 - p < \bar{v}$ and $B^I \leq \bar{B}$). So the buyer is indifferent for some $v \in \mathcal{V}$.

⁵⁰Here, we interpret $S'(\bar{B})$ as $\lim_{B \uparrow \bar{B}} S'(B)$, for any sequence of (B_n) , with $B_n \rightarrow \bar{B}$, at which S is differentiable.

Claim 4. *It holds that*

$$\lim_{B \downarrow 0} \frac{\phi(B)}{B} = +\infty.$$

Further, $\lim_{B \downarrow 0} B^O$ exists and is equal to 0.⁵¹ Similarly, $\lim_{B \downarrow 0} v = \underline{v}$.

Proof. Fix some $\varepsilon > 0$ small. Let $B^O = \frac{1-\delta}{\delta}\varepsilon$. We distinguish two cases.

- (a) $(1-\delta)p > \delta\hat{B}$: Let $B^I = \hat{B} - \frac{1-\delta}{\delta}\varepsilon(1+\varepsilon)$, $q = p - \frac{\delta}{1-\delta}\hat{B}$.
- (b) $(1-\delta)p \leq \delta\hat{B}$: Let $q = 0$, $B^I = \frac{1-\delta}{\delta}(\varepsilon^2 + \varepsilon + p)$, which is in $(0, \hat{B})$ for ε small enough.

In both cases, $v = \underline{v} + \varepsilon$ solves

$$(1-\delta)(\underline{v} + \varepsilon) + \delta B^O = (1-\delta)(q-p) + \delta B^I,$$

so that the buyer comes to the seller if, and only if, $v \geq \underline{v} + \varepsilon$. If the seller is indifferent for some $v > \underline{v}$ between coming and not, adding the seller's incentive constraint for $q > 0$ to the indifference condition for the buyer with type v gives, as in the previous claim,

$$\delta\phi(B^I) > (1-\delta)(p+v) + \delta B^O,$$

which is satisfied in case (a) as $\underline{v} + \varepsilon \rightarrow \underline{v}$, by the assumption (36). In case (b), the seller's incentive compatibility is trivial, since $q = 0$.

In either case, the buyer gets (net of the payoff of always going to the outside option)

$$(1-\delta)(\varepsilon(\varepsilon + F(\underline{v} + \varepsilon))),$$

while the seller gets $\delta F(\underline{v} + \varepsilon)S(\hat{B})$. Recall that, since the density of F is bounded away from 0, there exist constants $m, M > 0$ such that $M\varepsilon \geq F(\underline{v} + \varepsilon) \geq m\varepsilon$. Taking the ratio and $\underline{v} + \varepsilon \rightarrow \underline{v}$ yields the desired result in that case.

The claim that $\lim_{B \downarrow 0} B_0 = 0$ is immediate, as the buyer could otherwise secure a payoff bounded above $\mathbf{E}[\tilde{v}]$ by simply not coming in the first round. Similarly, if $v \rightarrow \underline{v}$ along some subsequence of B tending to 0, the seller would secure a payoff bounded above $\mathbf{E}[\tilde{v}]$, given that his preference to come to the seller for, say, all $\tilde{v} \leq (v + \underline{v})/2$ is bounded away from \underline{v} . \square

Note that this implies that E_δ has nonempty interior (whenever it properly includes $\mathbf{0}$) as for any $(B, \phi(B))$ vector, the line segment $[\mathbf{0}, (B, \phi(B))]$ is included in $\{(B, B + S/c) \mid (B, S) \in E_\delta\}$.

Note that, since ϕ is continuous, it admits a maximum. Fix the largest maximizer \hat{B} of ϕ , and let $\hat{\phi}$ denote the maximum.

Claim 5. *If $E_\delta \neq \{\emptyset\}$, then*

$$\delta\hat{\phi} > (1-\delta)p. \tag{36}$$

⁵¹More precisely, fix any sequence (B_n) , with $B_n \downarrow 0$, and any sequence of equilibria with buyer's payoff B_n ; then the corresponding sequence B_n^O converges to 0.

Proof. If some equilibrium quality q is s.t. $q > 0$, then incentive compatibility for the seller requires

$$\delta \frac{S(B^I)}{c} \geq (1 - \delta)q,$$

If the buyer is willing to come to the seller with positive probability in that round, then

$$(1 - \delta)(q - p) + \delta B^I > \delta B^O.$$

Adding up the inequalities gives

$$\delta \phi(B^I) > (1 - \delta)p + \delta B^O \geq (1 - \delta)p,$$

hence the conclusion. \square

From here on, we maintain as an assumption that (36) holds, as the equilibrium payoff set is simply the null payoff vector otherwise.

Claim 6. *If $B > 0$, then in any equilibrium achieving payoff vector $(B, \phi(B))$, the buyer comes wpp to the seller in the first round.*

Proof. Suppose not. Then (for at least some realization of the randomization device) the interim payoff vector is a convex combination of the vector $\mathbf{0}$, and some continuation payoff vector in E_δ . By the previous claim, such a convex combination must lie in the strict hypograph of ϕ , a contradiction, given that all continuation payoffs must lie on the graph. \square

Note that the previous proof establishes a somewhat stronger claim: for all realizations of the randomization device, the buyer comes to the seller wpp.

Let v denote the supremum of the types of the buyer who go to the seller, given some vector $(B, \phi(B))$ and corresponding equilibrium. By the previous claim, $v > \underline{v}$. In fact:

Claim 7. *Fix a payoff vector $(B, \phi(B))$ and corresponding equilibrium. Then without loss, there exists $v \in (\underline{v}, \bar{v}]$ s.t.*

$$(1 - \delta)v + \delta B^O = (1 - \delta)(q - p) + \delta B^I.$$

Proof. By the previous claim, the highest buyer's type that comes to the seller satisfies $v > \underline{v}$. If $v = \bar{v}$ strictly prefers to come to the seller, then we may increase B^O without loss. Hence, we may increase it either to the point at which indifference obtains when $v = \bar{v}$, and the claim follows, or $B^O = \bar{B}$. In the latter case, it would have to hold that $B = (1 - \delta)(q - p) + \delta B^I \geq (1 - \delta)\bar{v} + \delta \bar{B}$, for q the quality that the seller provides. This is impossible, since $B^I \leq \bar{B}$, and $q - p \leq 1 - p < \bar{v}$. \square

Suppose that $(B, \phi(B))$ is achieved without initial lottery. Then we have

$$B = (1 - \delta) \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) + \delta B^O, \tag{37}$$

as well as

$$\frac{S(B)}{c} = (1 - \delta)F(v) \left(\frac{p}{c} - q \right) + \delta F(v) \frac{S(B^I)}{c} + \delta \bar{F}(v) \frac{S(B^O)}{c}, \quad (38)$$

and, rearranging the indifference condition that defines v ,

$$0 = (1 - \delta)(q - p) + \delta B^I - ((1 - \delta)v + \delta B^O). \quad (39)$$

Multiply (39) by $F(v)$, add the resulting equations to (37) and (38), and we get

$$\phi(B) = (1 - \delta) \int_{\underline{v}}^v (\kappa - \tilde{v}) dF(\tilde{v}) + \delta F(v) \phi(B^I) + \delta \bar{F}(v) \phi(B^O), \quad (40)$$

where $B = (1 - \delta) \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) + \delta B^O$, $q \in [0, 1]$, and the seller's incentive constraint holds, namely either $q = 0$ or

$$\delta \phi(B^I) \geq (1 - \delta)(p + v) + \delta B^O,$$

combining the actual seller's IC constraint with the buyer of type v 's indifference, as before. The constraint $q \in [0, 1]$ can be rewritten as, using (39)

$$(1 - \delta)(p + v) + \delta(B^O - B^I) \in [0, 1 - \delta].$$

These last two constraints can be rewritten as (using the indifference condition)

$$\min\{\delta \phi(B^I), 1 - \delta + \delta B^I\} \geq (1 - \delta) \left(p + v - \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right) + B \geq \delta B^I. \quad (41)$$

We can rewrite (40) as

$$\phi(B) = (1 - \delta) \int_{\underline{v}}^v (\kappa - \tilde{v}) dF(\tilde{v}) + \delta F(v) \phi(B^I) + \delta \bar{F}(v) \phi \left(\frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right). \quad (42)$$

Recall that \bar{B} is the highest equilibrium payoff of the buyer.

Claim 8. *It holds that*

$$\delta \phi(\bar{B}) < 1 - \delta + \delta \bar{B}.$$

Proof. Note that, if $\delta \phi(\bar{B}) \geq 1 - \delta + \delta \bar{B}$, the seller's incentive compatibility is satisfied when $B^I = \bar{B}$ and thus can be ignored. This implies that the best equilibrium payoff for the buyer coincides with $q = 1$ always and myopic behavior by the buyer, contradicting Assumption **A1**. \square

C.2 Equilibrium Behavior for Payoffs (B, S) , $S = \phi(B)$

Given Claim 9, the vector $(\bar{B}, \phi(\bar{B}))$ lies in the hypograph of the map $B \mapsto 1 - \delta + \delta B$. Hence, either the graph of this map crosses the graph of $B \mapsto \phi(B)$, or it lies above it. Let \tilde{B} be where the largest intersection occurs if any, that is, the largest solution of $\delta \phi(B) = 1 - \delta + \delta B$, *provided* this solution is larger than \hat{B} (recall that this is the argmax of ϕ), and set $\tilde{B} = \hat{B}$ otherwise.

Motivated by the previous subsection, consider the programme $\mathcal{P}(B)$:

$$\max_{v, B^I} \left\{ (1 - \delta) \int_{\underline{v}}^v (\kappa - \tilde{v}) dF(\tilde{v}) + \delta F(v) \phi(B^I) + \delta \bar{F}(v) \phi \left(\frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right) \right\}$$

such that

$$\min\{\delta\phi(B^I), 1 - \delta + \delta B^I\} \geq (1 - \delta) \left(p + v - \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right) + B \geq \delta B^I, \quad (43)$$

over v , s.t. $B \geq (1 - \delta) \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v})$ (a nonempty set as this lower bound is 0 for $v = \underline{v}$). We note that, given that ϕ is increasing in B over $[0, \hat{B}]$, and given that the left-hand inequality of (43) is slack for $B \leq \hat{B}$, we always have $B^I \geq \hat{B}$ (unless the right-hand inequality of (43) binds). Similarly, because ϕ is decreasing in B over $[\hat{B}, \bar{B}]$, and because decreasing B relaxes the left-hand inequality for $B > \bar{B}$, it holds that $B \leq \bar{B}$. Accounting for the right-hand inequality of (43), we have established:

Claim 9. *It holds that*

$$B^I \in \left[\min \left\{ (1 - \delta)(p + \underline{v})/\delta, \hat{B} \right\}, \bar{B} \right].$$

Depending on the ordering between the bounds that appear in Claim 9, up to four regions can occur.⁵²

Claim 10. *For all $B \in [0, \bar{B}]$, $B^O \leq B$, and the inequality is strict if, and only if, $B \notin \{0, \bar{B}\}$.*

Hence, B^O may be as low as 0 (when $B = 0$) or as high as \bar{B} (when $B = \bar{B}$). To establish this claim, we distinguish the four regions of values of B , and prove the result region by region (Claims 11, 12, 13 and 14).

Definition 1. *Let:*

\mathcal{B}^L denote the subset of $\mathcal{B} := [0, \bar{B}]$ such that, in the solution to $\mathcal{P}(B)$,

$$\min \{ \delta\phi(B^I), 1 - \delta + \delta B^I \} > (1 - \delta) \left(p + v - \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right) + B,$$

i.e., $B^I = \hat{B}$;

\mathcal{B}^M denote the subset of \mathcal{B} such that

$$\delta\phi(B^I) > 1 - \delta + \delta B^I = (1 - \delta) \left(p + v - \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right) + B,$$

i.e., $B^I \in (\hat{B}, \bar{B})$;

⁵²Plainly, if the left-hand side of (43) binds, the R.H.S. does not, and vice-versa.

\mathcal{B}^H denote the subset of \mathcal{B} such that

$$\delta\phi(B^I) = 1 - \delta + \delta B^I = (1 - \delta) \left(p + v - \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right) + B,$$

i.e., $B^I = \tilde{B}$.

\mathcal{B}^{LL} denote the subset of \mathcal{B} such that

$$\delta B^I = (1 - \delta) \left(p + v - \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right) + B,$$

i.e., $q = 0$ and $B^I = \frac{B}{\delta} + \frac{1-\delta}{\delta} \left(p + v - \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right)$.

Because ϕ is continuous on its domain (as a concave function), the maximum theorem applies, and so the sets $\mathcal{B}^L, \mathcal{B}^M, \mathcal{B}^H$ must be unions of intervals (with \mathcal{B}^H being closed).

Claim 11. *The set \mathcal{B}^L is a subset of $[0, \hat{B}]$ (hence $B^I \geq B$ for any $B \in \mathcal{B}^L$), and $B^O < B$ whenever $B > 0$. The function ϕ is differentiable and strictly concave over any interval in \mathcal{B}^L . The cutoff v and the continuation payoff B^O are increasing in B over any such interval.*

Proof. Consider an interval of values in \mathcal{B}^L . We first prove that ϕ cannot be affine on an interval that contains B^O for some B in that interval. Suppose otherwise, i.e., $\phi(B) = aB + b$ on this interval. By the envelope, we then have $\phi'(B) = a = \bar{F}(v)\phi'(B^O) = \bar{F}(v)a$, and so either $v = \underline{v}$ or $a = 0$. The former is impossible given Claim 6.

If $a = 0$, then \hat{B} must be in this interval, and the objective rewrites $\phi(\hat{B}) = \max_v \int_{\underline{v}}^v (\kappa - \tilde{v}) dF(\tilde{v})$, implying

$$\phi(\hat{B}) = \int_{\underline{v}}^{\kappa \wedge \bar{v}} (\kappa - \tilde{v}) dF(\tilde{v}),$$

which is on the boundary of the feasible payoff set (this is the maximum feasible value of ϕ), requiring a constant value, namely $B^O = B^I = B$, which means that the buyer behaves myopically, a contradiction. So, B^O and B never lie on a segment of ϕ that would be affine, for $B \in \mathcal{B}^L$. That is, if $B \in \mathcal{B}^L$, and $B > B^O$, then $\partial\phi(B) > \partial\phi(B^O)$ (where $\partial\phi$ is the set of subgradients of ϕ). This implies that ϕ is strictly concave over any interval contained in \mathcal{V}^L .

Next, we show that ϕ is differentiable over any such interval. Fix $B \in \mathcal{B}^L$, $\bar{\epsilon} > 0$ such that $[B - \bar{\epsilon}, B + \bar{\epsilon}] \in \mathcal{B}^L$. Let v refer to the maximizer of $\mathcal{P}(B)$. For $\epsilon \in (0, \bar{\epsilon})$, let also v^ϵ solve

$$\frac{B + \epsilon}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^{v^\epsilon} F(\tilde{v}) d\tilde{v} = \frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v}.$$

Plainly, v^ϵ is a differentiable function of ϵ , with

$$\frac{dv^\epsilon}{d\epsilon} = \frac{1}{(1 - \delta)F(v^\epsilon)}.$$

Since v^ϵ is feasible at $B + \epsilon$, for small enough ϵ , we have

$$\begin{aligned} & \phi(B + \epsilon) - \phi(B) \\ \geq & (1 - \delta) \int_v^{v^\epsilon} (\kappa - \tilde{v}) dF(\tilde{v}) + \delta(F(v^\epsilon) - F(v)) \left(\phi(\hat{B}) - \phi\left(\frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v}\right) \right), \end{aligned}$$

and so

$$\begin{aligned} & \liminf_{\epsilon \downarrow 0} \frac{\phi(B + \epsilon) - \phi(B)}{\epsilon} \geq \\ & \left((1 - \delta)(\kappa - v) + \delta \left(\phi(\hat{B}) - \phi\left(\frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v}\right) \right) \right) \frac{f(v)}{(1 - \delta)F(v)}. \end{aligned}$$

By a similar argument,

$$\begin{aligned} & \limsup_{\epsilon \downarrow 0} \frac{\phi(B) - \phi(B - \epsilon)}{\epsilon} \leq \\ & \left((1 - \delta)(\kappa - v) + \delta \left(\phi(\hat{B}) - \phi\left(\frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v}\right) \right) \right) \frac{f(v)}{(1 - \delta)F(v)}. \end{aligned}$$

By concavity, $\liminf_{\epsilon \downarrow 0} \frac{\phi(B + \epsilon) - \phi(B)}{\epsilon} \leq \limsup_{\epsilon \downarrow 0} \frac{\phi(B) - \phi(B - \epsilon)}{\epsilon}$, and so these inequalities must be equalities, yielding that ϕ is differentiable at B .

Note that the maximality at the optimal v of the objective implies

$$0 = (1 - \delta)(\kappa - v)f(v) + \delta f(v)(\phi(\hat{B}) - \phi(B^O)) - (1 - \delta)\bar{F}(v)F(v)\phi'(B^O).$$

Hence, if $v < \kappa$, it must be that $B^O < \hat{B}$, since $\phi(\hat{B}) - \phi(B^O) \geq 0$. Note that, by the envelope,

$$\phi'(B) = \bar{F}(v)\phi' \left(\frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right). \quad (44)$$

Hence, $B \geq \hat{B}$ implies $B^O > B$, and so $B^O > B^I$, hence $v < 1 - p < \kappa$, a contradiction. Hence, $B^L \subset [0, \hat{B}]$, and $B^O < B$. \square

Claim 12. Consider values of B in \mathcal{B}^H . It holds that $B^O \leq B$ is increasing in B , with $B < \bar{B} \Rightarrow B^O < B$; in fact, $B - B^O$ is decreasing in B , with $B - B^O = 0$ as $B = \bar{B}$ (which is in \mathcal{B}^H). Also, v is decreasing over \mathcal{B}^H . The function ϕ is differentiable over any interval in \mathcal{B}^H .

Proof. Since $1 - \delta + \delta\tilde{B} = (1 - \delta) \left(p + v - \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right) + B$, then if $B_1, B_2 \in \mathcal{B}^H$, and $B_1 > B_2$, then $v_1 < v_2$, showing that v is decreasing. A simple calculation shows that

$$B - B^O = -\frac{1 - \delta}{\delta} \left((1 - \delta)(1 - p - v) + \delta\tilde{V} - \delta \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right),$$

and since the R.H.S. increases in v , and so decreases in B , $B - B^O$ decreases as well. If $B = \bar{B} \in \mathcal{B}^H$, it must be that $B^O = B$, for otherwise increasing B^O without changing v would lead to an increase

in B , a contradiction by definition of \bar{B} . The proof that ϕ is differentiable mimics the earlier proof of differentiability (see Claim 12) and is therefore omitted. \square

Claim 13. *The set \mathcal{B}^M is a subset of $[0, \tilde{B}]$. Further, $B^O < B < B^I$, with $B^I - B^O$ (and so v) decreasing in B . The function ϕ is differentiable on any interval in \mathcal{B}^M .*

Proof. Differentiability is shown as before.

We first show that ϕ cannot be affine over any interval that contains B , as well as either the corresponding B^O or B^I . Note that, by the envelope theorem,

$$\phi'(B) = F(v)\phi'(B^I) + \bar{F}(v)\phi'(B^O),$$

and so if ϕ is affine on an interval that contains either B^O or B^I , it must also contain the other possible continuation payoff.

Now, suppose for the sake of contradiction that $\phi(B) = aB + b$ on such an interval that we take to be maximal.

Recall that, by definition of this region,

$$B^I = \frac{B}{\delta} - \frac{1-\delta}{\delta} \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} + \frac{1-\delta}{\delta}(p+v-1).$$

Hence, $\phi(B)$ is given by (the maximum over v of)

$$\begin{aligned} \phi(B) = aB + b &= (1-\delta) \int_{\underline{v}}^v (\kappa - \tilde{v}) dF(\tilde{v}) + \delta F(v)\phi(B^I) + \delta \bar{F}(v)\phi(B^O) \\ &= (1-\delta) \int_{\underline{v}}^v (\kappa - \tilde{v}) dF(\tilde{v}) + \delta b + a \left(B - (1-\delta) \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right) \\ &+ (1-\delta)a\bar{F}(v)(p+v-1), \end{aligned}$$

and so

$$b = \int_{\underline{v}}^v (\kappa - a(1-p) - (1-a)\tilde{v}) dF(\tilde{v}),$$

a differentiable function of v maximized when

$$v = \frac{\kappa - a(1-p)}{1-a},$$

which is constant.⁵³ Hence, $B^I - B^O = \frac{1-\delta}{\delta}(p+v-1)$ is also constant, and from the definition of B^I and B^O , $B^I - B$ and $B^O - B$ are strictly increasing in B for constant v . By maximality of the interval, either (in fact, both) difference must be positive at the lower extremity of the interval, and negative at the upper extremity, a contradiction.

⁵³The case in which $v = \bar{v}$ because $\frac{\kappa - a(1-p)}{1-a} > \bar{v}$ is analogous.

Next, we show that $B^O \leq B \leq B^I$, with at least one strict inequality. By maximality of v , we must have

$$0 = (1 - \delta)(\kappa - v) + (1 - \delta)F(v)\bar{F}(v) (\phi'(B^I) - \phi'(B^O)) + \delta f(v) (\phi(B^I) - \phi(B^O)).$$

If $B^I < B^O$ then $v < 1 - p < \kappa$, as well as $\phi'(B^I) - \phi'(B^O) \geq 0$ and $\phi(B^I) - \phi(B^O) \geq 0$ (since $B^I \geq \hat{B}$). Hence, all three terms on the right are positive, a contradiction. Hence, $B^I \geq B^O$, and given the envelope theorem, namely,

$$\phi'(B) = F(v)\phi'(B^I) + \bar{F}(v)\phi'(B^O), \quad (45)$$

and concavity of ϕ (with the property that ϕ is not affine on any subinterval), $B \in [B^O, B^I]$. Finally, those three values cannot be equal, as v would remain constant over time, a contradiction. But if two are distinct, then again, given that ϕ is not affine on any subinterval containing B and either B^I or B^O , the three subgradients must be distinct, implying that $B^O < B < B^I$.

We note that $B \geq \tilde{B} \Rightarrow B^I \leq B$ immediately implies that B cannot belong to \mathcal{B}^M for $B \geq \tilde{B}$; recall also from Claim 11 that $B \geq \hat{B} \Rightarrow B \notin \mathcal{B}^L$; hence, for all $B \geq \tilde{B}$, $B \in \mathcal{B}^H$. \square

Claim 14. *The set \mathcal{B}^{LL} is a subset of $[0, \hat{B}]$. Further, $B^O < B < B^I$ for $B > 0$. The function ϕ is differentiable on any interval in \mathcal{B}^{LL} .*

Proof. Differentiability is shown as before. We note that, from the definition of \mathcal{B}^{LL} , it holds that

$$B^I = \frac{B}{\delta} + \frac{1 - \delta}{\delta} \left(p + v - \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right) > \frac{B}{\delta},$$

so that $B^I > B$, and given that $B^I < \hat{B}$, it follows that $\mathcal{B}^{LL} \subseteq [0, \hat{B}]$. By the envelope theorem, over any interval in \mathcal{B}^{LL} , it holds that

$$\phi'(B) = F(v)\phi'(B^I) + \delta\bar{F}(v)\phi'(B^O),$$

with, by Claim 6, $v > \underline{v}$, and so $F(v) > 0$. Ruling out that ϕ is affine over an interval containing B and B^I given B (and strict concavity) follows exactly the same steps as in Claim 13. Hence we must have $B^O < B$. \square

C.3 Asymptotic Behavior

Claim 15. *For any initial condition $B_0 \in (0, \bar{B}]$, the Markov chain $\{B_n : n \in \mathbf{N}\}$ defined by*

$$B_{n+1} = \begin{cases} B_n^I & \text{with prob. } F(v_n), \\ B_n^O & \text{with prob. } \bar{F}(v_n), \end{cases}$$

where v_n, B_n^I, B_n^O solve $\mathcal{P}(B_n)$, is such that

$$\limsup_{n \rightarrow \infty} B_n = \tilde{B}, \quad \liminf_{n \rightarrow \infty} B_n = 0.$$

Proof. As is clear from Claim 12, $\min\{B_n^I, B_n^O\} < B_n$, $\min\{B_n^I, B_n^O\} \leq \tilde{B}$ if $B_n \in \mathcal{B}^H$ (and recall that both B_n^I, B_n^O have positive probability). Hence, \mathcal{B}^H is exited almost surely whenever visited. Further, for $B_n \in \mathcal{B}^L \cup \mathcal{B}^M$, consider the (stopped) Markov chain $\{\check{B}_{n+m} : m \in \mathbf{N}\}$ defined by $\check{B}_n = B_n$, $\check{B}_{n+m+1} = B_{n+m+1}$ if $B_{n+m} \in \mathcal{B}^L \cup \mathcal{B}^M$, $= B_{n+m}$ if $B_{n+m} \in \mathcal{B}^H$. We note that, by the envelope theorem, there exists a selection of $\partial\phi(\check{B}_{n+m})$ that is a martingale. While it is not bounded (since $\phi(B)/B \rightarrow +\infty$ as $B \rightarrow 0$), it is bounded below (because $\mathcal{B}^L \cup \mathcal{B}^M \subset [0, \tilde{B}) \subsetneq [0, \bar{B})$). Hence it converges to a limit $\phi'(B_\infty)$ with finite expectation. This limit cannot be achieved for $B \rightarrow 0$ (again, since $\phi(B)/B \rightarrow +\infty$ as $B \rightarrow 0$). Because as shown in Claims 11 and 13, ϕ is not affine on any interval that would contain both some value B_{n+m} and its continuation payoff B_{n+m+1} , it must be that B_{n+m} converges, and this limit is in \mathcal{B}^H , *i.e.*, the region $\mathcal{B}^L \cup \mathcal{B}^M$, whenever visited, must be eventually exited again also.

Because $B^O < B$ for every B in $\mathcal{B}^L \cup \mathcal{B}^M$,⁵⁴ and the event that the continuation payoff is B^O has probability bounded away from 0, and because $B^I = \tilde{B}$ for every $B \in \mathcal{B}^H$ (and B^I is the continuation payoff with probability bounded away from 0 for B bounded away from 0), the conclusion follows. \square

⁵⁴And $B - B^O$ is bounded away from 0 if B is bounded away from 0; if not, by the maximum theorem, there would exist $B > 0$ such that $B^O = B$.