



On the Role of Outside Options in Buyer-Seller Relationships

Johannes Horner, Anna Sanktjohanser

► To cite this version:

Johannes Horner, Anna Sanktjohanser. On the Role of Outside Options in Buyer-Seller Relationships. 2025. [hal-04875876](#)

HAL Id: hal-04875876

<https://hal.science/hal-04875876v1>

Preprint submitted on 9 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

On the Role of Outside Options in Buyer-Seller Relationships

Johannes Hörner and Anna Sanktjohanser*

December 22, 2024

Abstract

We consider a buyer-seller relationship, in which due to private information and a lack of flexible transfers, cooperation cannot be sustained efficiently. In each round, the buyer either purchases from the seller or takes an outside option. The fluctuating outside option may be public or private information. The seller chooses what quality to supply. 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 the buyer fails to come and fast recoveries when the buyer returns. Transparency regarding the outside option benefits both buyer and seller.

KEYWORDS: Trust, Loyalty, Imperfect Monitoring.

JEL CLASSIFICATION NUMBERS: C72, C73, C78.

*Hörner: CNRS (TSE), 1 Esplanade de l'Université, 31080 Toulouse Cedex 06, France, joh.horner@gmail.com. Sanktjohanser: Toulouse School of Economics, 1 Esplanade de l'Université, 31080 Toulouse Cedex 06, France, anna.sanktjohanser@gmail.com. This paper has previously been circulated under the title “Too Much of A Good Thing?” 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 Dario Gori, Kouros Khounsari, John Mori and Stephan Waizmann for excellent research assistance. Johannes Hörner and Anna Sanktjohanser acknowledge 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 - MARKLIM).

Long-term informal relationships do not only arise when more formal market institutions are non-existent, they often arise alongside such markets. While these markets exhibit considerable fluctuations themselves, their volatility reverberates and complicates the dynamics of these informal relationships, whose evolution must provide the requisite incentives to both parties. The goal of this paper is to understand this link, and how outside options shape these dynamics.

Examples of such coexisting arrangements abound. In the US freight industry, shippers and carriers transact either via a spot market or via (mostly informal) long-term relationships (Harris and Nguyen, 2023a,b). In the Rwanda coffee industry, upstream farmers face a choice between processing beans at home and engaging in the relationship with the downstream mill (Macchiavello and Morjaria, 2021). In OTC markets, dealers engage in bilateral trades with investors, some of whom have access to a centralized platform (Allen and Wittwer, 2021). In the Bangladeshi garment industry, exporters face a choice between “relational buyers” and spot buyers who engage in one-time transactions (Cajal-Grossi et al. 2023). The evolution of these relationships must encourage cooperation by the party occasionally tempted by the outside option, given the cost for the counterparty of maintaining the relationship. This is a challenge whether or not this outside option is observable. As we show, observability plays a central role in the relationship, although robust features obtain independently of it. Surprisingly perhaps, transparency can benefit both parties.

More formally, this paper considers a repeated game between a buyer and a seller, in which due to private information and a lack of commitment and 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 (spot market). 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. In this setting, a folk theorem holds if discounting vanishes. In contrast, we assume that players are impatient and characterize the “exactly optimal” equilibrium behavior, with a particular focus on the buyer’s favorite one.

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) with private information, the relationship undergoes recurring cycles in quality and continuation value. Periods of reduced interaction lead to gradual declines in the quality supplied by the seller and the continuation value for both parties, but these recover rapidly when the buyer resumes visiting the seller after an extended absence; (4) Despite these ongoing cycles, the relationship endures and never breaks down, even in the face of private information.

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 unobservable, the seller adjusts the quality she would supply based on the buyer’s recent visiting history, which we refer to as the buyer’s loyalty. As the buyer fails to visit over consecu-

tive periods, the quality the seller would hypothetically provide for the buyer’s first return gradually declines. This slow deterioration reflects the seller’s decreasing trust in the buyer. However, a single purchase at this reduced quality fully restores trust: if the buyer makes a second consecutive purchase, the seller immediately provides high quality again. This dynamic illustrates the relationship’s gradual deterioration in response to reduced interaction (the seller is “slow to anger”), paired with the rapid recovery following even minimal renewed engagement (the seller is “quick to forgive”).

The relationship may gradually sour over extended periods as the buyer visits less frequently, leading the seller to reduce the hypothetical quality she would provide upon the buyer’s return. This decline in quality and loyalty occurs slowly enough that a visit always eventually reoccurs, preventing the relationship from completely breaking down, even in the face of private information. Upon the buyer’s return, the seller initially offers lower quality for one period, reflecting the erosion of trust. However, this single visit fully restores trust, prompting the seller to immediately provide high quality for the buyer’s next purchase. As a result, the relationship alternates indefinitely between phases of gradual erosion in quality and loyalty, and rapid recovery after renewed interaction.

Further, we prove that both buyer and seller benefit from transparency in the buyer-preferred equilibrium. That the buyer gains from transparency is consistent with the common wisdom in repeated games (better monitoring enlarges the set of equilibrium payoffs, Kandori, 1992), though there is no reason for the seller to benefit from more transparency in the buyer-preferred equilibrium. It is, however, more surprising in light of the prevalent concerns about data privacy, since it implies that transparency might well be in the best interest of the customer.¹

These results speak to a variety of phenomena observed in long-term business relationships. 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

¹We are not the first to point out that transparency is in the best interest of the consumer. For an overview of the debate, see the survey by Acquisti et al. (2016).

such mutually beneficial trades may be optimal for consumers. The necessity for the buyer to come more often than he would like (“too much of a good thing”) echoes Harris and Nguyen’s finding that a carrier facing dynamic incentives is *ceteris paribus* more likely to accept load offers than a carrier with only static incentives. And the gradual structure of the punishment scheme that we identify explains their finding that the shipper’s punishment scheme is *soft*, as they put it. Finally, 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 the last finding we emphasized: making information common knowledge reduces volatility in the relationship, thereby ensuring regular trade relative to a situation with private information.

The marketing literature has long emphasized the importance of treating buyer-seller exchanges as ongoing relationships, and documented the cycles and spirals that they undergo. Dwyer, Schurr and Oh (1987), laid out a life-cycle theory of such relationships. Jap and Anderson (2007) provide an extensive empirical test of this theory, describing the iterative cycles (spirals) within these relationships, as well as the steady declines (that can “linger for surprisingly long periods”) followed by repairs, where partners “forgive and forget” (Jap and Anderson, p.272). Further, Autry and Golicic (2010) provide empirical evidence on alternating patterns of relationship strength on performance over time, consistent with our results on downturns and recoveries in the relationship. On these performance spirals, also see Shamsollahi, Chmielewski-Raimondo, Bell and Kachouie (2021) for an extensive review.

Our work is related to several main strands of literature in economics. The role of outside options has been extensively explored in the context of bargaining. See, among others, Muthoo (1995), Compte and Jehiel (2002) and McClellan (2021) from a theoretical point of view, and, for an empirical perspective, Caldwell and Danieli (2024). This literature examines how outside options affect the division of surplus (and, in the case of McClellan, the dynamics of what share to demand). Here instead, our focus is how the outside option affects the surplus of the relationship itself, both at the extensive margin (how often does the buyer take the outside offer?) and at the intensive margin (how does the outside offer affect the quality supplied by the seller?). The theoretical literature on this is scant. As

discussed below, the dynamics we identify are closest to Li and Matouschek (2013), although the mechanism, as well as the setting (relational contracting) differ substantially. As already mentioned, the empirical work of Harris and Nguyen (2023a,b) shares our motivation.

Related also 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 private information on the buyer’s side but moral hazard on the seller’s side.

Although we do not allow flexible transfers, some of our results are reminiscent of themes in 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.² In Board (2011), the opportunity cost is publicly observable, but the principal, who has commitment, can choose to assign a task between multiple agents. Board (2011) focuses on understanding the relative treatment of insiders vs outsiders.³ Here instead, our focus is how the dynamics of the relationship themselves can address the moral hazard arising from this outside option, in lieu of transfers. 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.⁴ Crucially, the availability of transfers in this literature

²Also 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.

³This difference in the treatment of insiders vs. outsiders is also analyzed in Kranton (1996).

⁴To 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,

means that a player can be rewarded or punished without surplus destruction. Here, on the other hand, the seller can only punish by decreasing the quality, but this is inefficient.⁵

The rest of the paper is structured as follows. The model is described in Section 2. Section 3 analyzes the complete information case (the outside option is observable). Section 4 studies the incomplete information case (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, drawn according to F , with support $\mathcal{V} = [\underline{v}, \bar{v}]$, with $\underline{v} \geq 0$.⁶ The distribution F admits a differentiable 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 (Section 2) vs. the unobservable case (Section 3).
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 $q_n \in \mathcal{Q} := [0, 1]$.

Quality is measured in the buyer’s utility and entails a marginal cost c to the seller. The seller receives a fixed, exogenous price p whenever the buyer chooses *In*, independent of the quality she picks. Hence, if the seller picks maximum quality and the buyer visits the seller,

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).

⁵In 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.

⁶The 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.

the buyer gets $1 - p$ and the seller gets $p - c$. If the buyer does not visit the seller, the buyer enjoys his outside option and the seller gets nothing. More formally, 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 $(q_n - p, p - cq_n)$.

We assume that $0 < c < p < 1$ (trading is mutually beneficial, and quality is costly), and $\bar{v} > 1 - p$ (the highest outside option exceeds the maximum benefit from trading).

Given that the price is simply a transfer, surplus is v_n when the buyer chooses *Out* and $q_n(1 - c)$ otherwise. Hence, the socially efficient outcome involves the buyer visiting whenever $v_n \leq 1 - c$, and the seller picking then $q_n = 1$. In the one-shot game, however, the buyer would only visit the seller if $v_n < q_n - p$, and the seller would supply $q_n = 0$.

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

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

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.^{7,8}

Formally, 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. (Note that the buyer's choice in past rounds is omitted, as it is implied by the seller's quality.) 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, $a = O$.⁹ A behavior strategy for the buyer is a sequence $\sigma^B := (\sigma_n^B)_n$, where σ_n^B is a probability transition from

⁷For 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.

⁸Hence, we might equivalently assume that the outside option is fixed, but the seller's cost fluctuates. Specifically, we could assume that c is a random variable, i.i.d across rounds, drawn according to G , with support $\mathcal{C} = [\underline{c}, \bar{c}]$, with $\underline{c} \geq 0$.

⁹That 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} .

$H^{n-1} \times \mathcal{V} \rightarrow \{I, O\}$. Given a strategy profile $\sigma = (\sigma^B, \sigma^S)$, expected payoffs are defined as usual. When no risk of confusion arises, we use B_0, S_0 to denote this (expected) payoff. The solution concept is subgame-perfect Nash equilibrium.¹⁰ For conciseness, we omit the qualifier “with probability one” from formal statements. A public randomization device is available, but it is omitted from the notations.¹¹

Note that the seller can secure 0 by choosing zero quality, and the buyer can secure $\mathbf{E}[\tilde{v}]$ by picking the outside option. In fact, the minmax payoff vector $(\mathbf{E}[\tilde{v}], 0)$ is an equilibrium payoff vector: if the seller expects the buyer to never visit, saving on the cost is optimal if the buyer unexpectedly shows up. Conversely, if the buyer expects the seller to shirk on quality, he might as well pick the outside option. This strategy profile is referred to as *autarky*. Under complete information, it is without loss to assume that any deviation triggers autarky. Hence, in the observable case, we focus on behavior before any deviation.

As simple as the stage game might be, describing its feasible payoff set \mathcal{F} already requires some calculations. Its boundary involves three curves. Utility transfer between buyer and seller can occur via more or less frequent visits, or via lower or higher quality. This defines the first two curves. 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, to minimize a weighted sum of payoffs: suppose the buyer uses another cut-off rule, visiting 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 both players get their

¹⁰The solution with commitment on the buyer’s side is a special case of the results derived in Section 2.2. If the seller can commit, trade becomes uncomplicated, since by assumption $p > c$, and so the seller makes positive profits.

¹¹It plays no role in the best equilibrium but facilitates its derivation. It is used for simple equilibria; see Section 4.1.




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

minmax payoff.¹² The sets \mathcal{F} and \mathcal{V} (shaded area) are illustrated in Figure 1.

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.¹³

To fix ideas, we focus mostly on the buyer-preferred equilibrium—that is, the equilibrium that maximizes the buyer's *ex ante* payoff. Solving for this equilibrium calls for solving for the entire Pareto frontier of the equilibrium payoff set. Hence, the characterization of other equilibria (for instance, the seller-preferred equilibrium) follows as a by-product.

¹²More formally, the feasible payoff is given by $\mathcal{F} := \{(B, S) \in \mathbf{R}^2 \mid B^L \leq B \leq B^U\}$, where

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

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

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

¹³Because quality and outside options are drawn from intervals, mixed strategies turn out to play no role.

Of course, interesting dynamics only arise if myopic behavior cannot be supported via Nash reversion. 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.

To rule out trivialities, we exclude myopic behavior in equilibrium (except for autarky) by assuming that the seller is impatient enough.¹⁴ Formally,

Assumption A1:¹⁵ For all $q > 0$, it holds that:

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

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 , 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

¹⁴A folk theorem holds in our setting whether outside options are observed or not: any interior feasible, individually rational payoff vector can be supported provided the buyer and seller are patient enough. The standard results under imperfect monitoring (*e.g.*, Fudenberg, Levine and Maskin, 1994) do not apply given that action and type sets are infinite, but Radner's (1985) review strategies can certainly be adapted.

¹⁵In 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. Assumption **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 $q = 1$.

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.¹⁶ Assumption **A1** is imposed throughout (including in the unobservable case) and omitted from all statements.

Of course, if the seller is *too* impatient, autarky is the only equilibrium, an outcome that is equally uncharming. Our interest lies in the middle ground, when the parties are impatient enough that the seller's incentives are not trivially satisfied but patient enough that she can be motivated.¹⁷

2 Observable Case

2.1 Main Result

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 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 2.¹⁸

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*

¹⁶This 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.

¹⁷In the special case in which the value is uniformly distributed on the unit interval,

$$1 - p < \frac{1 - \delta}{\delta} \frac{c}{p - c} < 1$$

is a sufficient condition for both **A1** and this assumption. (It ensures that $v_S^* < 1$, see (2) below.)

¹⁸Recall that only on-path behavior is specified, as any deviation triggers autarky.

(I)initial Purchase (S)ubsequent Purchases




Figure 2: Buyer favorite equilibrium.

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 *I*nitial cut-off v_I^* that involves a lower probability of visiting the seller than myopic behavior would entail and a *S*ubsequent cut-off v_S^* that leads to more visits to the seller than he would like. This necessity for the buyer to visit the seller more often than he would like to is in line with the documented pattern of behavior in Harris and Nguyen (2023a,b): carriers who face dynamic incentives are more likely to accept load offers than a carrier with only static incentives.

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 compensate her for her cost.

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 only a feature of the buyer-preferred equilibrium 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, assume for now that maximum quality is both feasible (patronizing the seller often enough delivers sufficient rents) as well as desirable, and that the buyer has no incentive to renege when his outside option calls for a visit. Further, 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 visiting 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. 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. 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.

To summarize, there are two reasons why the relationship is inefficient: the buyer comes more often than he would like, and, in fact, as is easy to show, for a wide range of parameters, more often than is efficient: the seller isn't concerned about the efficiency of the relationship,

and the rents she demands are determined by her temptation to shirk. The second reason is that, to mitigate this inefficiency, it is sometimes beneficial to cut on efficient quality provision, as a last resort. As we shall see in Section 3.2, both inefficiencies necessarily arise under incomplete information.

2.2 Other Equilibrium Payoffs

Here, our focus has been on the buyer-preferred equilibrium. This equilibrium payoff, as well as the continuation payoff once the buyer has visited the seller, both lie on the Pareto frontier of the equilibrium payoff set. The proof of the proposition follows from a general characterization of this frontier.

We now provide a brief informal description of this frontier.¹⁹ The payoffs on the frontier that give the seller less than what she receives after the first visit in the buyer-preferred equilibrium (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 equilibrium, 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.

Proposition 2 below focuses on the segment of the boundary of the equilibrium payoff set when the seller receives less than the buyer would grant her in his favorite equilibrium. This section can be skipped without loss for the reader with little time for technical details. The main ingredients of Proposition 2 are most conveniently defined via a maximization programme. Let

$$q(v) := \frac{\delta F(v)}{1 - \delta + \delta F(v)} \frac{p}{c}. \quad (4)$$

¹⁹See the additional appendix for further details, and see Figure 3 for a graphical illustration of the equilibrium payoff set.




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

The quality $q(v)$ is the solution to

$$(1 - \delta)cq = \delta F(v)(p - cq). \quad (5)$$

That is, it is the quality level that makes the seller just indifferent between defecting and not, given that defection triggers autarky, and given that the buyer visits the seller if, and only if, the outside option is below v . Note that $q(\cdot)$ is a differentiable, increasing function: the higher the cut-off, the more often the buyer visits, and so the higher the quality that the seller is willing to provide. Hence, $q(v) \leq 1$ gives an upper bound $q^{-1}(1)$ to a cut-off v ensuring the seller's indifference.

Consider the programme

$$\max_{\underline{v}} \int_{\underline{v}}^v (q(v) - p - x)F(dx) \quad (\mathcal{P})$$

over $v \leq \min\{\bar{v}, q^{-1}(1)\}$ such that

$$v > \underline{v} \Rightarrow \int_{\underline{v}}^v (q(v) - p - x)F(dx) \geq \frac{1 - \delta}{\delta}(v + p - q(v)). \quad (6)$$

Let v_S^* denote the (if need be, largest) maximizer, and set $q^* := q(v_S^*)$. $v_S^* > \underline{v}$ is *equivalent* to assuming that autarky isn't the unique equilibrium; as mentioned, we maintain as much. By Assumption **A1**, $v_S^* > q^* - p$: the buyer must patronize the seller more often than is myopically optimal.

The objective in (\mathcal{P}) is the buyer's payoff if he uses constant cut-off v , net of the outside option, given that the seller supplies constant quality $q(v)$.

The constraint (12) obtains from rearranging

$$(1 - \delta)(q(v) - p) + \delta \left(F(v)(q(v) - p) + \int_v^{\bar{v}} xF(dx) \right) \geq (1 - \delta)v + \delta \mathbf{E}[\tilde{v}],$$

stating that the buyer prefers taking the good of quality $q(v)$ at price p , when his outside option is v , rather than pocketing the outside option at the cost of triggering autarky. This is the buyer's incentive compatibility: delivering the promised utility to the seller must be sequentially rational.

Let v_I^* solve the equation

$$v + \delta \int_v^{\bar{v}} (x - v)F(dx) = q^* - p + \delta \int_{v^*}^{\bar{v}} (x - (q^* - p))F(dx). \quad (7)$$

Equation (7) states the buyer's indifference, in case his outside option is precisely equal to v , between taking the outside option, and visiting the seller, triggering thereby the loyalty loop.²⁰

Because $v_S^* > q^* - p$, and because the left-hand side is increasing in v , it follows that this equation admits a unique solution, and that $v_I^* < v_S^*$.

The values q^*, v_S^*, v_I^* are the three main components of Proposition 2. In the case of the uniform distribution, these cut-offs can be solved in closed-form, as illustrated in Figure 4. Quite generally, the programme (\mathcal{P}) can be dissected so as to elucidate in which circumstances each of the possible constraints ($q^* \leq 1$, $v \leq \bar{v}$, and (12)) bind. A systematic analysis is omitted here, but examples are provided below.

More generally, we now describe the Pareto-frontier –up to a point. Namely, let $\mathcal{B}(S)$ denote the highest buyer's payoff when the seller gets at least S . Let \check{S} denote the lowest promise S such all subgradients at $\mathcal{B}(\cdot)$ exceed $-1/c$. It can be shown that $\tilde{S} \leq \check{S}$. We restrict ourselves to the description of the boundary $(S, \mathcal{B}(S))$ up to \check{S} . The reason is that, beyond this level, quality supplied need no longer be independent of the outside option at hand: lowering the quality is a way to transfer utility from seller to buyer (precisely at

²⁰The left-hand side takes advantage of the fact that entering the loyalty loop by visiting the seller in the future is equivalent to drawing the outside option for which the buyer is indifferent between doing so and not.

rate $-1/c$) which, despite being inefficient, can nonetheless become handy when the seller's promise is high enough.

Let $\tilde{S} := \frac{1-\delta}{\delta}cq^*$ denote the seller's payoff in the loyalty loop (cf. (5), as the seller is willing to shirk). For $S \geq \tilde{S}$, consider the programme (\mathcal{P}_S) obtained by appending to (\mathcal{P}) the constraint

$$\frac{1-\delta}{\delta}cq(v) \geq S, \quad (8)$$

stating that the seller is guaranteed at least payoff S (again, given that she is willing to shirk).²¹ For $S \in [\tilde{S}, \check{S}]$, let $q(S), v(S)$ denote the (if need be, largest) solution of (\mathcal{P}_S) . For $S \in [\tilde{S}, \check{S}]$, the Pareto-efficient equilibrium delivering S to the seller involves a stationary rule: the buyer visits the seller if, and only if, $v \leq v(S)$, and the seller supplies quality $q(S)$.

If instead $S < \tilde{S}$, as is the case of the buyer-preferred equilibrium, two cut-offs are used. Once the buyer visits the seller, the same rule prevails as in the buyer-preferred equilibrium: the seller provides quality q^* , and the buyer repeats his visit whenever $v \leq v_S^*$, as given by (\mathcal{P}) . The first cut-off is determined by the solution v to

$$S = \frac{F(v)}{1 - \delta + \delta F(v)}((1 - \delta)(p - cq^*) + \delta \tilde{S}).$$

If $v < v_I^*$, that is, if S is lower than the seller's payoff \underline{S} in the buyer-preferred equilibrium (the value obtained by plugging v_I^* in the right-hand side), the buyer uses the cut-off v_I^* . Visiting the seller less often would hurt the buyer as well. If instead $S \geq \underline{S}$, the buyer visits the seller as rarely as possible, by using as cut-off the solution $v = v_I(S) \in [v_I^*, v_S^*)$ to this equation.

To summarize, the Pareto-frontier (including the buyer-preferred payoff) is achieved by the following on-path behavior.

Proposition 2. *For all $S \in [0, \check{S}]$, the Pareto-efficient equilibrium that delivers at least S to the seller involves, for $S \geq \tilde{S}$, a constant quality $q(S)$ by the seller, and one cut-off rule $v(S)$; or, for $S < \tilde{S}$, a constant quality q^* , and two cut-off rules $v_I(S), v_S^*$. In case $S < \tilde{S}$, as long as the buyer chose Out in all prior rounds, he goes to the seller in round n if, and only if, $v_n \leq v_I(S)$. Otherwise, he goes to the seller if, and only if, $v_n \leq v_S^*$. These cut-offs*

²¹By convention, the value of (\mathcal{P}_S) is set to $-\infty$ if the feasible set is empty.




Figure 4: Quantile and Quality along the Pareto frontier when $(\delta, p, c) = (\frac{13}{20}, \frac{13}{25}, \frac{1}{4})$.

are such that

$$v_I(S) < q^* - p < v_S^*.$$

Unsurprisingly, given the simplicity of the equilibrium behavior, this Pareto-frontier can be described explicitly, but such a description is omitted here. Nonetheless, it follows directly from the program (\mathcal{P}_S) that (i) the cut-off used by the buyer (after the first visit) is non-decreasing with the promised seller's payoff S , (ii) the quality supplied by the seller is non-increasing in S .

Figure 4 illustrates Proposition 2.

2.3 Discussion

Persistence We have 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.²² In practice, however, alternative trading opportunities are often persistent: the current buyer's market, say, foreshadows the buyer's market in the near future.

²²The i.i.d. assumption prevents the analysis from being conducted in continuous time. With persistence, equilibrium analysis would require the current outside option to be added to the promised utility as a second state variable. This is unlikely to be tractable.

To discuss this issue concretely, let us assume that, given an outside option worth v in round n , the outside option in the next round 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 from F . That is, the outside option follows a renewal process – a fairly specific but common way to model persistence, which includes the i.i.d. case ($\lambda = 0$) and the perfectly persistent case ($\lambda = 1$).

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.²³ 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. The higher the current outside option, the higher the buyer's expected payoff from breaking up the relationship at that point. Hence, if v^* is the prevailing cut-off, and the seller is expected to supply maximum quality, the constraint

$$(1 - \delta)v^* + \delta(1 - \lambda)\mathbf{E}[\tilde{v}] \leq (1 - \delta)(1 - p) + \delta(1 - \lambda)B,^{24}$$

which states that the payoff from pocketing the outside option and triggering autarky must be lower than the payoff from visiting the seller and continuing the relationship, becomes more problematic as λ increases.

Hence, lowering the cut-off v^* might be necessary to prevent the buyer from reneging. But lowering the cut-off means lowering the frequency of visits, and hence, all else fixed, lowering the seller's payoff. To satisfy the seller's incentive to supply the expected quality, it might become necessary to transfer utility from buyer to seller by lowering the quality supplied for

²³Of course, quality could vary over time, in which case the inequality must be adjusted accordingly.

²⁴To 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.

outside options that are so low that buyer's incentive compatibility is unproblematic. More precisely, for some parameters, the buyer-preferred equilibrium can involve a cut-off v^* and qualities $q(v) < q(v^*)$ for $v < v^*$ to reduce the seller's cost and at the same time boost her expected payoff, so as to ensure that, for all $v \leq v^*$, not only buyer incentive compatibility holds, but also

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

where $S(v)$ is the seller's continuation payoff when the buyer visits the seller given outside option v .²⁵

Flexible prices It is not the purpose of this paper to discuss why and when contractual prices do not arise in buyer-seller relationships (see Macaulay, 1963). However, we note that fixed prices are a pervasive feature of many real-world settings, often due to practical constraints or strategic considerations. For instance, Macchiavello and Morjaria (2015) document fixed-price contracts in the Kenyan rose export industry, where most firms negotiate constant prices with their main buyer throughout the year and prices are not indexed on quality nor on prices prevailing on the spot market. In our model, the dynamics in the buyer-seller relationship emerge precisely 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

²⁵Given persistence, $S(v) = \frac{1-\delta}{1-\delta\lambda}(p - cq(v)) + \frac{\delta(1-\lambda)}{1-\delta\lambda}F(v^*)(p - c\mathbf{E}[q])$.

provides an alternative explanation for such discounts, as they mitigate 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.

Applications of fixed prices abound in real-world markets. Long-term supply agreements in manufacturing, franchising arrangements, export-oriented agricultural contracts, and public procurement often rely on fixed prices to minimize transaction costs and ensure predictability.

Incomplete Information about the Distribution Loyalty programs are meant to address not only moral hazard but also incomplete information: 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 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. In a related setting, Board (2011) highlights the insider bias that arises; a similar 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, $a = O$.²⁶ 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 rules out the uninteresting equilibrium in which the buyer is supposed to always visit the seller, independent of his privately observed outside option, with a failure of visiting that triggers autarky.

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).

²⁶That 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} .

Second, because the outside option is unobservable, the cut-off structure of the buyer's strategy is now a feature of any best reply of the buyer. (This wasn't the case for all equilibria under complete information, only for some extremal ones.) 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.²⁷

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 along which the seller has not deviated.

For concreteness, we still focus on the buyer's favorite equilibrium, although this requires us to also solve for many other extremal equilibria that serve as possible continuations.

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 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.²⁸

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,

²⁷This property is definitely not shared by the buyer's favorite equilibrium under complete information, except in the initial stage.

²⁸Transitions 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.

$q^H \neq q^L$. 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: $v^H > v^L$.²⁹

It is not hard to see that when in state H , play must remain in state H whenever the buyer visits; similarly, when in state L , play must remain 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 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

²⁹Specifically, for $\mathcal{V} = [0, 1]$ and $F = \mathcal{U}[0, 1]$, one can check that these assumptions ($q^H = 1, q^L > 0$ and $S^H > S^L$) are satisfied by the buyer-preferred two-state automaton whenever

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

and δ is in the range described below; see (9) and the discussion that follows.




Figure 5: Best two-state automaton.

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 5 illustrates the best two-state automaton.

This takeaway will generalize, *mutatis mutandis*, to the best equilibrium, without any restriction to finite automata: *forgiveness is drastic, punishment is gradual*.

For those interested in quantifying this trade-off, we now illustrate this discussion by solving explicitly for this optimal two-state equilibrium with a particular parametrization. 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

$$\varepsilon := \frac{\delta}{1 - \delta} - \frac{c}{(1 - p)(p - c)}, \quad (9)$$

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.³⁰

³⁰Even 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

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 $\nu^H \in \mathbf{R}_+$.³¹ 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 from the buyer to the seller in state H , which both parties nonetheless prefer to state L . Indeed, the buyer gets³²

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

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,

involve the root of an uninspiring polynomial of degree 14—hence the use of expansions.

³¹These positive constants and those that follow are specified in Appendix B.

³²The 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.




Figure 6: Automaton with consideration stage.

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. This second take-away also generalizes: *the buyer-preferred equilibrium involves a transient consideration stage*.

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 $\nu^I, \beta^I \in \mathbf{R}_+$.³³ See Figure 6.

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,

³³Note 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 and information aggregation.

thus reducing inefficiency without lessening incentives. “Splitting” 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.³⁴ Figure 7 illustrates such an automaton.

This is the third takeaway that is generalized in the next section: *gradual punishments take the form of small degradations in the quality supplied.*

There is no reason to expect that improvements to the automaton 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. This will also allow us to address one question that cannot be conclusively settled with finite automata. We know that the relationship starts in a transient state. Does it eventually end in autarky?

³⁴More 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.

3.2 The Equilibrium Payoff Frontier

This section elucidates the seller and buyer behavior along the equilibrium payoff frontier.

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 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 8 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 quality 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, $\mathcal{S}(B)$. 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 explained below, a subset of \mathcal{D} is self-generating, and it is not the Pareto frontier.³⁵

³⁵The lower boundary of the equilibrium payoff set is also self-generating for the same reasons. The lower




Figure 8: The curve \mathcal{D} is self-generating.

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 (see Claims 3 and 5 in the appendix).

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

$$\lim_{B \downarrow \mathbf{E}[\tilde{v}]} \frac{S}{B - \mathbf{E}[\tilde{v}]} = +\infty. \quad (10)$$

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 9 and 13 below.




Figure 9: Dynamics along the curve \mathcal{D} .

To describe behavior, recall that we maintain throughout the assumptions that players are patient enough to sustain equilibria other than autarky, but not patient enough for the buyer's first best to be an equilibrium outcome. To reduce the number of possible configurations, we further assume that the buyer-preferred equilibrium involves an (initial) maximum quality.

How does behavior evolve along the curve \mathcal{D} ? Figure 9 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 10 illustrates a sample path over time: it shows how the hypothetical quality the seller would supply if the buyer went to seller 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 9 and 10, the relationship experiences gradual downturns (in terms of payoffs to both parties and the hypothetical quality provided by the seller) when the buyer breaks the seller's trust by not visiting (indicated by blue arrows in Figure 9 and blue dots in Figure 10). However, the relationship instantaneously recovers (in terms of payoffs to both parties, and also the quality provided by the seller) if loyalty is shown. Specifically, following a visit after an

extended absence, the seller initially provides lower quality, reflecting the gradual erosion of trust during the buyer’s absence. Yet, a single additional visit fully restores trust, leading the seller to immediately revert to providing high quality. This dynamic is illustrated by the jumps in payoffs and quality following a buyer’s visit, as indicated by the red arrows in Figure 9 and the red dots in Figure 10.

This gradual decline of a relationship, but also the potential for significant upturns is consistent with the behavior observed in the US truckload freight industry (Harris and Nguyen, 2023b). Shippers in this industry explicitly rank carriers, depending on past behavior, so as to incentivize carrier cooperation. The ranking pins down the order in which carriers are offered loads. If a carrier behaves opportunistically, by rejecting a load in favor of taking higher-paying loads in the spot market, the shipper may punish the carrier by demoting him in the carrier ranking. These rankings often change in the span of weeks, reminiscent of the gradual downturns of a relationship here. But equally, good behavior in a “time of need” can result in promotions in the carrier ranking.

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 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

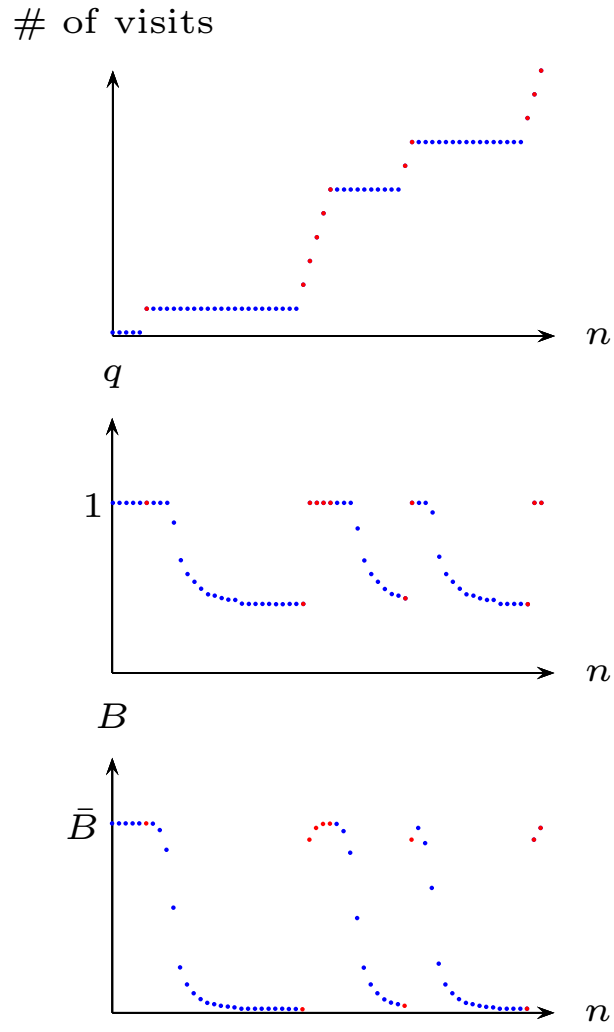


Figure 10: Sample path for $(\delta, p, c) = (\frac{13}{20}, \frac{13}{25}, \frac{1}{4})$.

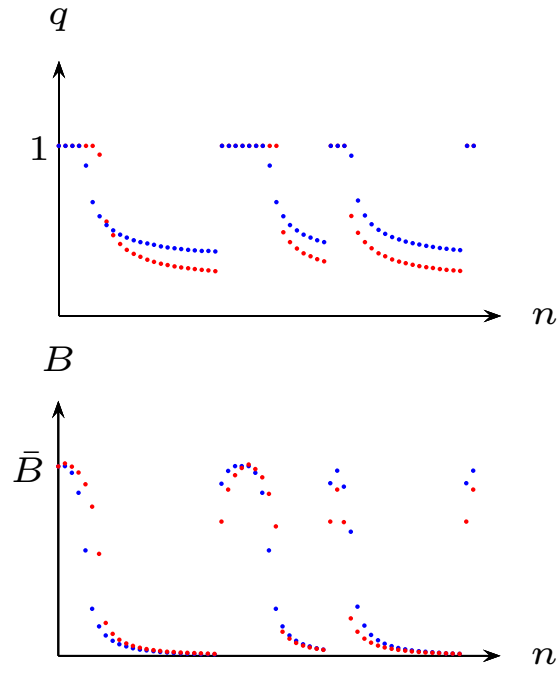


Figure 11: Sample path for $(\delta, p, c) = (\frac{13}{20}, \frac{13}{25}, \frac{1}{4})$ in blue and $(\delta, p, c) = (\frac{16}{20}, \frac{13}{25}, \frac{1}{4})$ in red.

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, Claims 11–14). 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.^{36,37}




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

We may summarize this discussion with the following Lemma, that summarizes what has been established in Appendix C. Throughout, we maintain Assumption **A1**. We also assume that, in the buyer's favorite equilibrium, $q = 1$.³⁸

Proposition 2. *Consider an equilibrium payoff vector (B, S) that maximizes the seller's payoff S , given the buyer's payoff B . Let q be the quality supplied in case the buyer visits, v the cut-off used by the buyer, and B^I (B^O) the continuation payoffs if the buyer visits (does not visit). There exist $0 < \hat{B} < \tilde{B} < \check{B} < \bar{B}$ such that:*

³⁶This 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} .

³⁷As 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 $\hat{B} < \tilde{B}$, this need not be the case. If, instead, $\hat{B} = \tilde{B}$, then any buyer visit leads to the same “resetting” of the buyer's payoff. The case $\hat{B} < \tilde{B}$ provides the “richest” dynamics; hence, we focus on it here. The analysis in Appendix C provides a systematic taxonomy of the equilibrium structure.

³⁸Section 2 already explains in detail when and why this might not be the case when outside options are observable. Here, we cut on the number of possible configurations considered in Appendix C. The other cases can be analyzed along the same lines.




Figure 13: Range of continuation payoffs following In .

1. for all $B > \check{B}$, $B^I < B^O \leq B$ ($B^O = B$ only iff $B = \bar{B}$), and so $v < q - p$, and
2. for all $0 < B < \check{B}$, $B^I < B^O \leq B$ ($B^O = B$ only iff $B = \bar{B}$), hence $v > q - p$.

For all B , $B^I \leq \tilde{B}$ (and so $B^I \leq \check{B}$). Hence, a single visit by the buyer takes the buyer's payoff below \tilde{B} forever. That is, the buyer comes more often than he wishes except possibly in the initial round. If $q \in (0, 1)$, $B^I = \hat{B}$ ($B^I < \hat{B} \Rightarrow q = 0$, $B^I > \hat{B} \Rightarrow q = 1$).

For all $B > \hat{B}$, the seller's incentive constraint binds.

As the proposition makes clear, except possibly for the first visit, the buyer comes more often than he would like. This is illustrated in Figure 14. Eventually, as he fails to come often enough and his payoff dips sufficiently below \hat{B} , so that $B^I \leq \hat{B}$, the seller cuts on the quality she supplies. These are the two potential sources of inefficiency already identified in the observable case. However, unlike in the observable case, quality undersupply necessarily occurs on path, as it is the seller's only tool to discipline the buyer. Note that for the parameters chosen in Figure 14, the buyer never comes more often than is efficient.

Appendix C further partitions the buyer's payoff into four regions, according to whether $q = 0, 1$ or $q \in (0, 1)$, and whether the seller's incentive compatibility binds or not. While it




Figure 14: Comparison of the equilibrium cut-off v to the efficient cut-off $q - cq$ and the buyer's myopic cut-off $q - p$.

seems plausible that each region is a single interval, and that these intervals are ranked, we have not been able to prove this.³⁹

Nonetheless, numerical simulations suggest a robust pattern (based on the uniform distribution on the unit interval for the valuations, but invariant to changes in the other parameters).⁴⁰ Figure 14 provides an illustration of this pattern. The lower the buyer's payoff is, the lower the seller's quality.

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 the seller would supply in case of a visit 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 the quality the seller would supply if the buyer visits 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

³⁹Appendix C contains some clues nonetheless: for instance, the interval $[\tilde{B}, \bar{B}]$ is included in one region only, and another region is included in $[0, \hat{B}]$.

⁴⁰Our numerical calculations are based on a Mathematica program that solves for the equilibrium payoff set to a degree of approximation of order 10^{-4} , established by using value iteration to get a set that is an upper bound to the equilibrium payoff set, and showing that a scaled-down version of this set is a lower bound to the equilibrium payoff set, as this smaller set expands under the Bellman-Shapley operator. The program is available upon request.

and pressing threat of a decrease in the seller's quality. For buyer's payoffs such that the seller's hypothetical 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.

3.3 The Fate of the Relationship

Section 3.1 has already touched upon the dynamics of payoffs along the equilibrium payoff frontier. Here, we further describe how behavior evolves over time, and solve for the long-run outcome.

Figures 10, 11 and 12 show how the relationship evolves over time. Specifically, Figure 10 provides a representative sample path of visits, quality levels, and continuation buyer's payoff, as the round n progresses. Figure 11 shows a representative sample path of quality levels and continuation buyer's payoff, as the round n progresses for different discount factors. The sample path for both discount factors starts with an initial value of the buyer's payoff in buyer preferred equilibrium when $\delta = 13/20$. Note that as the rounds progress, quality starts to drop earlier when the discount factor is lower.⁴¹ However, once quality drops, it drops faster for higher discount factors. Recoveries in terms of quality are always instantaneous regardless of the discount factor. Regarding continuation payoffs, downturns as well as recoveries are more gradual when the discount factor is higher. 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.⁴² It should be clear that, conditional on not dissolving, buyer (and seller) values must keep changing given that the buyer must be

⁴¹Note that this contrast would be even starker had we started the sample path with the higher discount factor with an initial value of the buyer's payoff in buyer preferred equilibrium when $\delta = 16/20$. Instead, we start the sample path when $\delta = 16/20$ with the same initial value as when $\delta = 13/20$.

⁴²This 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).

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 (10)). 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.⁴³

3.4 A Comparison with the Observable Case

Plainly, behavior under incomplete information is more complex than in the observable case. Nonetheless, focussing on the buyer preferred equilibrium, a simple comparison obtains. As before, we maintain **A1**, and focus on the case in which $q = 1$ in the buyer preferred equilibrium (as we have done in Section 3.2, 3.3).⁴⁴

Proposition 4. *In the buyer preferred equilibrium, both the buyer and the seller are better off when the outside option is observable.*

Here, we sketch the proof. Following Kandori (1992), the equilibrium payoff set expands with better monitoring.⁴⁵ The intuition for this is simple: one can replicate any equilibrium

⁴³The 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.

⁴⁴We do not know whether $q = 1$ under incomplete information implies $q = 1$ in the observable case.

⁴⁵Kandori's Proposition 1 does not literally apply, given that he focuses on finite-action, simultaneous-move, pure-strategy equilibria.

from the game with incomplete information when outside options are observable, by simply have the seller not condition on the specific outside option that prevails, but rather, only on the buyer's choice to visit or not. This immediately implies that the buyer must be better off: if the equilibrium payoff set expands, then the buyer maximum equilibrium payoff must increase. What is not obvious is that the seller also gains.

To see why this is the case, note that, in both cases, the seller incentive compatibility constraint binds (see Lemma 2); hence, her payoff once the buyer visits her is simply $(1 - \delta)c$, the rent she can secure by shirking. Therefore, the buyer is better off when the outside option is observable if, and only if, the buyer visits her sooner, that is, if the cut-off outside option \bar{v} is higher in the observable case. In both cases, this cut-off makes the buyer indifferent between visiting the seller or not; that is, letting \bar{B} denote his *ex ante* payoff in both cases,

$$\bar{B} = \int \max\{v, \bar{v}\} dF(v),$$

and so, the higher \bar{B} , the higher \bar{v} . But we have already argued that the buyer is better off with observable outside options. Hence, so is the seller.

Figure 15 compares the equilibrium payoff set in the complete information case to that in the incomplete information case. S




Figure 15: Comparison of equilibrium payoff sets when $(\delta, p, c) = (\frac{13}{20}, \frac{13}{25}, \frac{1}{4})$. The boundary of the equilibrium payoff set for the observable (unobservable) case is shown in blue (black); the boundary of 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.

However, the ranking identified in the proposition 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.

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 Li and Matouschek, 2013), they arise naturally under moral hazard and private information. 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.

References

- Abreu, D., D. Pearce, and E. Stacchetti.** 1990. “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring.” *Econometrica*, 58: 1041–1063.
- Acquisti, A., Taylor, C., and Wagman, L.** 2016. “The Economics of Privacy.” *Journal of Economic Literature*, 54: 442–492.
- Allen, J., and M. Wittwer** 2023. “Centralizing Over-The-Counter Markets?” *Journal of Political Economy*, 131: 3310–3450.
- Athey, S., and K. Bagwell.** 2001. “Optimal Collusion With Private Information.” *RAND Journal of Economics*, 32: 428–465.
- Autry, C.W., and S.L. Golicic.** 2010. “Evaluating buyer-supplier relationship-performance spirals: A longitudinal study.” *RAND Journal of Operations Management*, 28: 87–100.
- Board, S.** 2011. “Relational Contracts and the Value of Loyalty.” *American Economic Review*, 101: 3349–3367.

- Cajal-Grossi, J., Macchiavello, R., and Noguera, G.** 2023. “Buyers’ Sourcing Strategies and Suppliers’ Markups in Bangladeshi Garments.” *The Quarterly Journal of Economics*, 138(4): 2391–2450.
- Caldwell, S., and O. Danieli.** 2024. “Outside Options in the Labor Market.” *The Review of Economic Studies*, forthcoming.
- 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.
- Compte, O., and P. Jehiel.** 2002. “On the Role of Outside Options in Bargaining with Obstinate Parties.” *Econometrica*, 70: 1477–1517.
- 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.
- Harris, A., and T. M. A. Nguyen** 2023a. “Long-Term Relationships and the Spot Market: Evidence from US Trucking.” *working paper*, MIT.
- Harris, A., and T. M. A. Nguyen** 2023b. “Long-Term Relationships in the US Truckload Freight Industry.” *working paper*, MIT.
- 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.
- Jap, S.D., and E. Anderson** 2007. “Testing a Life-Cycle Theory of Cooperative Interorganizational Relationships: Movement across Stages and Performance.” *Management Science*, 53: 260–275.
- Kandori, M.** 1992. “The Use of Information in Repeated Games with Imperfect Monitoring.” *The Review of Economic Studies*, 59: 581–593.
- Kranton, R.** 1996. “The Formation of Cooperative Relationships.” *Journal of Law, Economics, & Organization*, 12: 214–233.
- Levin, J.** 2003. “Relational Incentive Contracts.” *American Economic Review*, 93: 835–857.
- Macaulay, S.** 1963. “Non-Contractual Relations in Business: A Preliminary Study.” *American Sociological Review*, 28: 55–67.

- Macchiavello, R., and A. Morjaria.** 2015. “The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports.” *The American Economic Review*, 105(9): 2911 – 2945.
- Macchiavello, R., and A. Morjaria.** 2021. “Competition and Relational Contracts in the Rwanda Coffee Chain.” *The Quarterly Journal of Economics*, 136: 1089–1143.
- 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.
- McClellan, A.** 2021. “Dynamic Outside Options and Optimal Negotiation Strategies.” *working paper*, Chicago Booth.
- Möbius, M.** 2001. “Trading Favors.” *working paper*, Harvard.
- Muthoo, A.** 1995. “On the Strategic Role of outside Options in Bilateral Bargaining.” *Operations Research*, 43: 292–297.
- 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.
- Shamsollahi, A., D.A. Chmielewski-Raimondo, S.J. Bell and R. Kachouie** 2021. “Buyer-supplier relationship dynamics: a systematic review.” *Journal of the Academy of Marketing Science*, 49: 418–436.
- 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, so that it suffices to establish the latter. Here, we prove Proposition 2 of the paper. Let $\mathcal{B}(S)$ be the buyer's highest equilibrium payoff consistent with the seller's equilibrium payoff being *at least* S ($\mathcal{B}(S) = -\infty$ if no equilibrium exists that gives the seller S or more). Let \bar{S} denote the maximum promise such that $\mathcal{B}(S) > -\infty$.

We start with a description of the buyer's favorite equilibrium, *e.g.*, yielding payoff $\bar{B} := \mathcal{B}(0)$.

Define

$$q(v) := \frac{\delta F(v)}{1 - \delta + \delta F(v)} \frac{p}{c}. \quad (11)$$

The definition of $q(v)$ (see (11)) gives the constant quality level that makes the seller just indifferent between defecting and not, given that defection triggers autarky, and given that the buyer comes if, and only if, the outside option is below v ; that is, it is equal to the quality q that solves

$$(1 - \delta)cq = \delta F(v)(p - cq).$$

We note that $q(\cdot)$ is a differentiable and increasing function. (The higher the cut-off, the more often the buyer comes, the higher the quality that the seller is willing to provide.) Hence, $q(v) \leq 1$ gives an upper bound $q^{-1}(1)$ to a cut-off v .

Consider the programme

$$\max_{\underline{v}} \int_{\underline{v}}^v (q(v) - p - x)F(dx) \quad (\mathcal{P}^S)$$

over $v \leq \min\{\bar{v}, q^{-1}(1)\}$ such that

$$\int_{\underline{v}}^v (q(v) - p - x)F(dx) \geq \frac{1 - \delta}{\delta}(v + p - q(v)). \quad (12)$$

The feasible set is non-empty and closed. Let v^* denote the (if need be, largest) maximizer,

and set $q^* := q(v^*)$. The objective of the programme is precisely the payoff from using this cut-off rule v , with quality $q(v)$ provided whenever $\tilde{v} \leq v$, net of the outside option. Equation (12) obtains from rearranging the equation

$$(1 - \delta)v + \delta \mathbf{E}[\tilde{v}] \leq (1 - \delta)(q(v) - p) + \delta \left(F(v)(q(v) - p) + \int_v^{\bar{v}} xF(dx) \right),$$

stating that the buyer prefers taking the good of quality $q(v)$ at price p , given that doing so leads to future deals with the seller at cut-off v and quality $q(v)$, when his outside option is v , rather than the outside option, given that doing so would result in autarky.

The objective is differentiable, with derivative

$$q'(v)F(v) - (v + p - q(v))f(v),$$

and while the first term is positive, the second term that is subtracted can be positive too—it is the case, in particular, when $v > q(v) - p$, so when the cut-off outside option exceeds the benefit from trade, as is to be expected in an equilibrium, given **A1**. Hence, the function that appears in the objective (and in the L.H.S. of the constraint) can have a maximizer v s.t. $q(v) < 1$. This possible non-monotonicity is also the reason why v^* need not be the highest value v that satisfies (12). Hence, v^* is either $\min\{\bar{v}, q^{-1}(1)\}$, or one of values v that solves (12), or one of the solutions to a first-order condition.

We may now restate Proposition 2.

Proposition 2. *The buyer favorite equilibrium involves a constant quality q^* by the seller, and two cut-off rules for the buyer. As long as the buyer chose Out in all prior rounds, he goes to the seller in round n if, and only if, $v_n \leq v_I^*$. Otherwise, he goes to the seller if, and only if, $v_n \leq v_S^* = v^*$. These cut-offs are such that*

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

We start the proof with a series of elementary claims. Given that we may focus on on-path histories (deviations being followed by autarky), an equilibrium specifies, for each

round n , a map from the sequence of outside options $(\tilde{v}_1, \tilde{v}_2, \dots, \tilde{v}_n)$ into a decision (I or O), and a quality q (possibly 0 after some histories). Recall that a public randomization device is available. Hence, the function $S \mapsto \mathcal{B}(S)$ is concave, for all $S \in [0, \overline{S}]$. We call an equilibrium *extremal* if it achieves payoff vector $(S, \mathcal{B}(S))$ for some S —in particular, if it achieves a Pareto-optimal equilibrium payoff.

Note that given any (on-path) history, without loss, the continuation strategy profile of an extremal equilibrium is an extremal equilibrium as well, as increasing the buyer's payoff after a given history, fixing the seller's payoff, increases the buyer's *ex ante* payoff and only relaxes the constraints that the buyer's actions must satisfy.

Henceforth, the buyer's *interim* incentive compatibility constraint, or BIC for short, given (a history and) an outside option \tilde{v} such that the buyer is supposed to visit the seller, refers to the constraint that

$$(1 - \delta)(q(\tilde{v}) - p) + \delta B(\tilde{v}) \geq (1 - \delta)\tilde{v} + \delta \mathbf{E}[\tilde{v}],$$

which states that going to the seller given continuation payoff $B(\tilde{v})$, as specified by the equilibrium, exceeds the payoff from taking the outside option, followed by the autarky payoff. Similarly, the seller incentive compatibility constraint, or SIC for short, given \tilde{v} , is the condition that

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

which states that supplying quality $q(\tilde{v})$ is better for the seller given continuation payoff $S(\tilde{v})$, rather than deviating (to quality 0), given that autarky ensues.

From here on, attention is restricted to extremal equilibria with payoff $B > \mathbf{E}[\tilde{v}]$.

Claim 1. *Wlog, the buyer's decision in round n is a cut-off value $v_n := v_n(\tilde{v}_1, \dots, \tilde{v}_{n-1})$ such that he chooses I if, and only if, $\tilde{v} \leq v_n$.*

Proof. If not, make a measure-preserving permutation of the values for which the buyer chooses I (with the corresponding permutation of the values of q) that leads to a distribution of values triggering visits that is first-order stochastically dominated by the original one. This does not change the seller's payoff, but improves the buyer's payoff (and relaxes BIC given

\tilde{v} , given that he is supposed to come for lower values of \tilde{v}). \square

Claim 2. *If (S, B) is an equilibrium payoff, with $B > \mathbf{E}[\tilde{v}]$, then $(S + \varepsilon, B - \varepsilon/c)$ is also an equilibrium payoff, for $\varepsilon > 0$ small enough.*

Proof. Fix an equilibrium σ with $B > \mathbf{E}[\tilde{v}]$. Let H' be the set of finite on-path histories $h = (v_1, \dots, v_n)$ (all n) such that (i) $q(h) > 0$, (ii) BIC does not bind along the sequence (that is, for all truncations (v_1, \dots, v_m) , $m \leq n$). Because $B > \mathbf{E}[\tilde{v}]$, $\mathbf{P}^\sigma(H') > 0$.⁴⁶ Pick a subset of H' with positive probability under σ , and decrease the specified quality by $\varepsilon/((1 - \delta)c) > 0$ at any such history. SIC is relaxed, and by assumption, for $\varepsilon > 0$ small enough, BIC continues to hold along the path to such a history. This change increases the seller payoff by ε , and decreases the buyer's payoff by $-\varepsilon/c$. \square

We note that this claim implies that the map $S \mapsto \mathcal{B}(S)$ has slope at least $-1/c$. (Recall that $\mathcal{B}(\cdot)$ is concave). Hence, any extremal equilibrium with $B > \mathbf{E}[\tilde{v}]$ maximizes some weighted sum $\lambda B + (1 - \lambda)S$ over equilibrium payoff vectors (S, B) , for some $\lambda > c/(1 + c)$.

We let $\check{S} := \inf\{S \mid -1/c \in \partial\mathcal{B}(S)\}$, that is, \check{S} is the smallest promise S such that $-1/c$ is a subgradient of $\mathcal{B}(S)$. By Claim 2, $\{(S, \mathcal{B}(S)) : S \in [\check{S}, \bar{S}]\}$ is a straight line segment with slope $-1/c$. We relax our analysis by assuming that payoff vectors (S, B) on this extended line segment, with $S \geq \bar{S}$, are available as continuation payoffs (verifying *ex post* that such vectors are not used).

Given some extremal equilibrium with payoff (S, B) , let v_1 denote the cut-off used in the initial round (see Claim 1). To emphasize that (S, B) is achieved in the initial round, we also write (S_1, B_1) . Recall that, given outside option \tilde{v} , $B(\tilde{v})$ (resp., $S(\tilde{v})$) denote the continuation equilibrium payoff of the buyer (resp., seller).

Claim 3. *Fix \tilde{v} such that either (i) $\tilde{v} > v_1$ or (ii) $\tilde{v} \leq v_1$ and neither BIC nor SIC binds at \tilde{v} . Then, wlog, $B(\tilde{v}) = B_1$, $S(\tilde{v}) = S_1$.*

Proof. Whatever weighted sum $\lambda B_1 + (1 - \lambda)S_1$ the (extremal) equilibrium payoff maximizes, it is optimal to set $B(\tilde{v}), S(\tilde{v})$ to maximize the same weighted sum, unless some

⁴⁶If $\mathbf{P}^\sigma(H') = 0$, then along any history, until BIC binds for the first time (if ever), the buyer must receive zero quality whenever he comes. Given that $\underline{v} > 0$, this implies that $B \leq \mathbf{E}[\tilde{v}]$.

constraint binds. Similarly, increasing $q(\tilde{v})$ slightly, and increasing $S(\tilde{v})$ slightly also, to keep S_1 constant, increases the weighted sum, given Claim 2 (namely, the slope of $\mathcal{B}(\cdot)$ is at least $-1/c$). \square

Claim 4. *Fix $\tilde{v} \leq v_1$ such that BIC (resp., SIC) does not bind. Then, wlog, $S(\tilde{v}) \geq S_1$ (resp., $S(\tilde{v}) \leq S_1$).*

Proof. This is a simple variation of the previous claim. Otherwise, increase (resp., decrease) slightly $S(\tilde{v})$ (adjusting $B(\tilde{v})$ accordingly, given that these are extremal equilibria). \square

Claim 5. *Fix $\tilde{v}, \tilde{v}' \leq v_1$ such that BIC does not bind at \tilde{v}, \tilde{v}' . Then, $(q(\tilde{v}), S(\tilde{v}), B(\tilde{v})) = (q(\tilde{v}'), S(\tilde{v}'), B(\tilde{v}'))$, wlog. In particular, either SIC binds at both \tilde{v}, \tilde{v}' , or at neither.*

Proof. Otherwise, replace $(q(\tilde{v}), S(\tilde{v}), B(\tilde{v})), (q(\tilde{v}'), S(\tilde{v}'), B(\tilde{v}'))$ by their (density-weighted) average, so that S_1 remains constant. Given concavity of $\mathcal{B}(\cdot)$, B_1 cannot decrease. \square

Claim 6. *Fix $\tilde{v} \leq v_1$ such that BIC binds at \tilde{v} . Then BIC binds at all $\tilde{v}' \in [\tilde{v}, v_1]$, wlog.*

Proof. Otherwise, rearrange $(q(\tilde{v}), S(\tilde{v}), B(\tilde{v})), (q(\tilde{v}'), S(\tilde{v}'), B(\tilde{v}'))$. \square

Claim 7. *Fix $\tilde{v} \leq v_1$ such that SIC does not bind at \tilde{v} . Then $q(\tilde{v}) = 1$, wlog.*

Proof. Increasing $q(\tilde{v})$ slightly, and increasing $S(\tilde{v})$ slightly also, to keep S_1 constant, relaxes BIC given Claim 2 (namely, the slope of $\mathcal{B}(\cdot)$ is at least $-1/c$), and increases the weighted sum $\lambda B_1 + (1 - \lambda)S_1$ that the equilibrium maximizes. \square

Claim 8. *Given an equilibrium with payoff $S < \check{S}$, the set $\{\tilde{v} : \tilde{v} \leq v_1, \text{BIC and SIC bind}\}$ is nowhere dense in $[\underline{v}, v_1]$.*

Proof. Given that BIC and SIC are continuous in $q(\tilde{v}), S(\tilde{v}), B(\tilde{v})$, this is equivalent to stating that BIC and SIC cannot identically hold over an interval $I \subset [\underline{v}, v_1]$. Suppose otherwise. Then, over this interval,

$$0 = \delta S(\tilde{v}) - (1 - \delta)cq(\tilde{v}),$$

and BIC holds, namely,

$$0 = (1 - \delta)(q(\tilde{v}) - p - \tilde{v}) + \delta(B(\tilde{v}) - \mathbf{E}[\tilde{v}]).$$

If $q(\tilde{v}) = 0$ or 1 over a subinterval, then it is constant over that interval; hence, $S(\tilde{v})$ is then constant as well (by SIC), but since \tilde{v} varies, $B(\tilde{v})$ cannot be constant, by BIC; a contradiction, since an improvement is then possible, by keeping $S(\tilde{v}), q(\tilde{v})$ the same over this subinterval, but by replacing $B(\tilde{v})$ by its maximum.

Hence $q(\tilde{v}) \in (0, 1)$ over this interval (except possibly at some isolated points), then by considering a local variation in $q(\tilde{v}), S(\tilde{v})$, one can increase $B(\tilde{v})$ in the last display, while keeping the first equation displayed satisfied, unless $B(\tilde{v}) + S(\tilde{v})/c$ is constant over that interval, that is, $S(\tilde{v}) > \check{S}$. If $S(\tilde{v})$ for some \tilde{v} (outside of this interval), then the range of $S(\tilde{v})$ over $[\underline{v}, \bar{v}]$ does not lie on a line segment, and an improvement is possible by considering a convex combination of $q(\cdot), S(\cdot)$, using strict concavity. \square

Claim 9. *There cannot be $\tilde{v}, \tilde{v}' \leq v_1$ such that only SIC binds at \tilde{v} , and only BIC binds at \tilde{v}' .*

Proof. Suppose otherwise. Given Claims 5, 6 and 8, there exists $\tilde{v}'' \in (\underline{v}, v_1)$ such that SIC (only) binds for all $\tilde{v} \in [\underline{v}, \tilde{v}'']$ and BIC (only) binds for all $\tilde{v} \in (\tilde{v}'', v_1)$, with $q(\cdot) = 1$ in the latter case, given Claim 7. Hence, $(1 - \delta)c < \delta S(\tilde{v}) \leq \delta S_1$, for $\tilde{v} \in (\tilde{v}'', v_1)$, using Claim 4. By Claim 4 still, $S_1 \leq S(\tilde{v})$ for $\tilde{v} \in [\underline{v}, \tilde{v}'']$, and so $(1 - \delta)c < \delta S(\tilde{v})$ for such \tilde{v} , contradicting that SIC binds. \square

Claim 10. *Suppose SIC, and only SIC binds at $\tilde{v} \leq v_1$. Then either $q(\tilde{v}) = 1$ (and so, by SIC, $S(\tilde{v}) = (1 - \delta)c/\delta$) or $S(\tilde{v}) = \check{S}$ (and so, by SIC, $q(\tilde{v}) = \delta\check{S}/(1 - \delta)c$), wlog. Specifically, $q(\tilde{v}) = 1$ if, and only if, $(1 - \delta)c \leq \delta\check{S}$.*

Proof. Suppose $q(\tilde{v}) < 1$ and $S(\tilde{v}) < \check{S}$. Increase $q(\tilde{v})$ slightly, and increase $S(\tilde{v})$ so that SIC still binds. That is, the changes $\Delta q(\tilde{v}) > 0, \Delta S(\tilde{v})$ are such that

$$(1 - \delta)c\Delta q(\tilde{v}) = \delta\Delta S(\tilde{v}).$$

It follows that, conditional on \tilde{v} , the payoff of the buyer changes by

$$(1 - \delta)\Delta q(\tilde{v}) + (1 - \delta)\Delta B(\tilde{v}),$$

where $(S(\tilde{v}) + \Delta S(\tilde{v}), B(\tilde{v}) + \Delta B(\tilde{v}))$ is an extremal equilibrium payoff. Given the discussion that follows Claim 2 (namely, the map $S \mapsto \mathcal{B}(S)$ has slope greater than $-1/c$ for $S < \check{S}$), it follows that, conditional on \tilde{v} , this change benefits the buyer; it also benefits the seller, since SIC holds, yet the payoff from defecting has increased. Hence, such a change improves payoffs. Plainly, this implies that either $S(\tilde{v}) = \check{S}$, or $q(\tilde{v}) = 1$, whichever constraint binds first. Hence also, $q(\tilde{v}) = 1$ if, and only if, $(1 - \delta)c \leq \delta\check{S}$. \square

By Claims 6–9, either (i) SIC, and only SIC, binds at all $\tilde{v} \in [\underline{v}, v_1]$; (ii) SIC, and only SIC, binds at all $\tilde{v} \in [\underline{v}, v_1)$, and both SIC and BIC bind at v_1 ; (iii) neither SIC nor BIC bind, for any $\tilde{v} \in [\underline{v}, v_1]$; (iv) there exists $\tilde{v}'' \in (\underline{v}, v_1]$ such that neither SIC nor BIC bind for $\tilde{v} < \tilde{v}''$, and BIC only binds for $\tilde{v} \in [\tilde{v}'', v_1]$; (v) BIC, and only BIC, binds at all $\tilde{v} \in [\underline{v}, v_1]$. Note that in all cases $v_1 > \underline{v}$, as otherwise the payoff cannot be extremal (being a weighted average of the autarky payoff and the continuation payoff).

Case (v) can be ignored, as it implies that the buyer's payoff is $\mathbf{E}[\tilde{v}]$: indeed, either he is indifferent between trading with the seller, and initiating autarky, or his outside option is above the cutoff v_1 , in which case he takes the outside option, and finds himself with the same payoff B_1 , and so the same choices, from the next round onward; hence, he is always willing to take the outside option.

We now turn to ordering these possibilities as a function of S , the equilibrium payoff they yield to the seller. We note that in cases (iii)–(iv), there exists $\tilde{v} < v_1$ such that $q(\tilde{v}) = 1$ and $S(\tilde{v}) \leq S_1$ (Claims 3 and 4). Hence, it must be that $(1 - \delta)c \leq \delta S(\tilde{v}) \leq \delta S_1$. Given that, in cases (i)–(ii), there exists \tilde{v} such that SIC binds, and so $(1 - \delta)cq(\tilde{v}) = \delta S(\tilde{v}) \geq \delta S_1$, it follows that if case (i) or (ii) occurs for some S , and case (iii), (iv) occurs for some S' , then $S \leq S'$.

Next, we argue that if S is the seller's payoff in an instance of case (iii) (with cutoff v_1), and S' is the seller's payoff in an instance of case (iv) (with cutoff v'_1 , buyer's payoff B'_1),

then $S < S'$.

To see this, considering case (iv), we first note that the seller's payoff when $\tilde{v} \in [\tilde{v}'', v'_1]$ is at least as high as his payoff when $\tilde{v} > v'_1$. To see this, observe that for $\tilde{v} \in [\tilde{v}'', v'_1]$, the buyer would be strictly better off if the equilibrium specified that he take the outside option with continuation payoff B'_1 (as is the case for $\tilde{v} > v'_1$). If the seller were also strictly better off with the alternative specification, there would exist an equilibrium that strictly Pareto-dominates the one considered, a contradiction.

Hence, the seller's payoff satisfies

$$S' \geq (1 - \delta)F(\tilde{v}'')(p - c) + \delta(1 - F(\tilde{v}''))S',$$

while

$$S = (1 - \delta)F(v_1)(p - c) + \delta(1 - F(v_1))S.$$

Hence, if $\tilde{v}'' > v_1$, $S' > S$. So suppose instead $\tilde{v}'' \leq v_1$. Because BIC is slack at v_1 in case (iii),

$$(1 - \delta)(1 - p) + \delta B_1 > (1 - \delta)v_1 + \delta \mathbf{E}[\tilde{v}],$$

and because it is not at \tilde{v}'' in case (iv),

$$(1 - \delta)(1 - p) + \delta B'_1 \leq (1 - \delta)\tilde{v}'' + \delta \mathbf{E}[\tilde{v}].$$

Hence, $B'_1 < B_1$, and so $S' > S$.

Next, we argue that if S is a payoff obtained in an instance of case (i), and S' is a payoff obtained in an instance of case (ii), then $S < S'$.

To see this, note first that, in case (ii), we can wlog pick $q(v_1) = q(\tilde{v})$ (and so $S(v_1) = S(\tilde{v})$), for $\tilde{v} < v_1$ (independent of the choice of \tilde{v} , by Claim 10) –that is, we can pick the same quality (and continuation payoffs) for all types up to *and including* v_1 . Suppose not, that is, $q(v_1) \neq q := q(\tilde{v})$ for $\tilde{v} < v_1$. Since BIC is slack for $\tilde{v} < v_1$, it must be satisfied at $\tilde{v} = v_1$ with the choice $q(v_1) = q$ (and $S(v_1) = S(\tilde{v})$, $\tilde{v} < v_1$). Hence, conditional on v_1 , the buyer must be weakly better off with this specification $q(v_1) = q$. The seller is also better off if $q > q(v_1)$ (given that SIC binds with both specifications, a higher quality yields a higher

seller's payoff, since a higher quality means a higher cost saving in case of defection). Hence, we must have $q < q(v_1)$. By Claim 10, this means that $S(\tilde{v}) = \check{S}$, all $\tilde{v} < v_1$. Since SIC binds at v_1 , and $q < q(v_1)$, it then holds that $S(v_1) > \check{S}$. Thus, we can reduce $S(v_1)$ slightly (and hence $q(v_1)$, while preserving SIC), and keep both seller and buyer indifferent; indeed, we can do this up to the point when $S(v_1) = \check{S}$, and so $q(v_1) = q$.

Returning to the claim that $S < S'$, note that the quality $q(\tilde{v})$ (constant in both cases up to the cut-off) is the same in both cases (and so are continuation payoffs), given Claim 10. Hence, since BIC binds at the cut-off v_1 in case (ii), but not in case (i), it follows that the cut-off is higher in case (ii). Note that, in both cases, the seller's payoff is a weighted average of the payoff she gets if the buyer comes, and the payoff if he does not visit, which is δS (or $\delta S'$, depending on the case), with weights $F(v_1)$, $1 - F(v_1)$. Since $S > \delta S$ (and similarly, $S' > \delta S'$), her payoff is higher when the buyer comes than when he does not; hence, her payoff is higher, the higher the cut-off. Hence, $S' > S$.

We note that SIC and BIC binding at v_1 in case (ii), together with Claim 10 that pins down $q(\tilde{v})$, $S(\tilde{v})$, all $\tilde{v} \leq v_1$, pins down v_1 ; hence, case (ii) occurs (at most) for one value of S . If this case occurs, case (i) must prevail for all $S' < S$.

We now divide the analysis in two cases, depending upon which of the two cases Claim 10 distinguishes prevails.

1. $\delta\check{S} < (1 - \delta)c$. Then, considering cases (i) or (ii) $q := q(\tilde{v})$, all $\tilde{v} \leq v_1$, is such that $q = \delta\check{S}/((1 - \delta)c) < 1$. In particular, $\delta\check{S} < (1 - \delta)c$. Let S^c the highest value of S for which case (i) or (ii) applies. Consider a sequence $S^n \downarrow S^c$, assuming that case (iii) or (iv) applies. Then, $\delta S^n = \delta S^n(\tilde{v}) \geq (1 - \delta)c$ for any \tilde{v} in the interval for which neither SIC nor BIC applies. Taking limits, $\delta S^c \geq (1 - \delta)c$. Hence, $S^c > \check{S}$, and cases (iii) or (iv) are irrelevant, since they would occur in the range of promises S for which the slope of \mathcal{B} is $-1/c$.

Hence, we can summarize the case when $\delta\check{S} < (1 - \delta)c$ as follows. When $S < \check{S}$, case (i) applies, the continuation payoff of the seller remains at S if the buyer does not visit, and moves to \check{S} if he comes to the seller. When $S = \check{S}$, regardless of whether (i) or (ii) applies, and regardless of whether the buyer comes to the seller or not, the seller's

continuation payoff remains at \check{S} . Hence, it is implemented by a two-state automaton with a transient state. The following steps allow us to check whether $\delta\check{S} < (1 - \delta)c$ applies: if case (i) applies at \check{S} , then we can solve for the buyer's payoff B as a function of $S < \check{S}$ (\check{S} a parameter to be determined); that is, for the function $\mathcal{B}(S)$, and hence for its derivative; hence, for the value of S at which this derivative is equal to $-1/c$; if this value satisfies BIC at the cut-off v_1 , then this case holds; if not, then either case (ii) applies at \check{S} (and still $\delta\check{S} < (1 - \delta)c$), or $\delta\check{S} \geq (1 - \delta)c$. Recall that if case (ii) applies, the value \check{S} can be solved for; it is then immediate to check whether this value is such that $\delta\check{S} < (1 - \delta)c$. If not, $\delta\check{S} \geq (1 - \delta)c$ must hold.

2. $\delta\check{S} \geq (1 - \delta)c$. Then, considering cases (i) or (ii), $q := q(\tilde{v}) = 1$, for all $\tilde{v} \leq v_1$. In particular, $\delta S \leq \delta S(\tilde{v}) = (1 - \delta)c$ (the first inequality follows from Claim 4). Let S^c the highest value of S for which case (i) or (ii) applies. Taking limits for any sequence $S^n \uparrow S^c$, we get $\delta S^c \leq (1 - \delta)c$. Consider now a sequence $S^n \downarrow S^c$, assuming that case (iii) or (iv) applies. Then, $\delta S^n = \delta S^n(\tilde{v}) \geq (1 - \delta)c$ for any \tilde{v} in the interval for which neither SIC nor BIC applies. Taking limits, $\delta S^c \geq (1 - \delta)c$. Hence, $S^c = (1 - \delta)c/\delta$, and (generically) this value cannot be the unique one for which case (ii) applies (see discussion above), so we can ignore case (ii). Hence, case (iii) applies when $S > S^c$, for S sufficiently close to S^c .

Hence, we can summarize the case when $\delta\check{S} \geq (1 - \delta)c$ as follows: starting from $S < (1 - \delta)c/\delta$, case (i) prevails; the continuation payoff of the seller remains at S if the buyer does not visit, and moves to $(1 - \delta)c/\delta$ if he comes to the seller; at that point, whether he comes to the seller or not, the seller's continuation payoff remains at $(1 - \delta)c/\delta$ (the transition value from case (i) to case (iii)). Hence, it is implemented by a two-state automaton with a transient state. Calculations in that case are immediate as well: solve for $\mathcal{B}(\cdot)$ for $S < (1 - \delta)c/\delta$ assuming the structure is as described; check that $\mathcal{B}'((1 - \delta)c/\delta) > -1/c$.

These computations are straightforward and omitted. We note that in all cases SIC binds when the buyer comes (provided the initial promise S is below \check{S} or $(1 - \delta)c/\delta$, depending on the case), and that, plainly BIC must hold. Hence, the buyer's favorite

equilibrium must involve his favorite quality, or equivalently cut-off, given that SIC gives a one-to-one relationship between both, subject to BIC, and subject to quality being less than one, and the cut-off being less than \bar{v} . It is immediate that this gives a value S below \tilde{S} or $(1 - \delta)c/\delta$, depending on which case occurs.

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, in the statements (though, for clarity, we do not always do so in the explanations), we normalize 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. Working with this weighted sum of payoffs has the advantage that the quality cancels from the objective (given that the quality q has benefit q and cost cq). We let \overline{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.

Throughout, we take for granted that $E_\delta \neq \{\mathbf{0}\}$, the autarky payoff, though we repeat this when useful.

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 $[\underline{v}, \bar{v}]$, and that the distribution F is assumed to be twice continuously differentiable, with density f bounded away from 0.

Claim 11. *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 visits the seller with positive probability ("wpp"). Note that $n < \infty$, for otherwise $\phi(0) = 0$. Second, the average 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 . Since $\mathbf{E}[\tilde{v}]$ is both his average payoff and his minmax payoff, his reward in round n cannot exceed $\mathbf{E}[\tilde{v}]$. Without loss, take $n = 1$.

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

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

Claim 12. *It holds that, for all $B \in (0, \bar{B}]$, $\phi(B) > 0$.*

Proof. This is immediate, as if the buyer gets more than $\mathbf{E}[\tilde{v}]$ (that is, when his normalized payoff is 0), he must visit 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 on (\underline{v}, \bar{v}) , with left- and right-derivatives everywhere. In fact, it is continuous on $[\underline{v}, \bar{v}]$ by standard arguments (omitted). 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 13. *It holds that*

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

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

Proof. The fact that $B^O = \bar{B}$ is without loss when $B = \bar{B}$ 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 benefits from it if he keeps using the same cut-off; *a fortiori* for the cut-off that is optimal given the change). In fact, this gain is strict if $\bar{B} > 0$, as we assume throughout. Hence, B^O must equal \bar{B} .

To show that the slope of ϕ is infinite at \bar{B} , consider setting $B^O = (1 - \alpha)\bar{B}$, for some (small) $\alpha > 0$ (keeping B^I and q constant). Indifference of the buyer at the cut-off v_α is equivalent to⁴⁷

$$(1 - \delta)(q - p - v_\alpha) = \delta((1 - \alpha)\bar{B} - B^I).$$

This implies that $(1 - \delta)dv_\alpha = \delta\bar{B}d\alpha$ at $\alpha = 0$. Using the envelope theorem, the buyer's payoff changes by

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

while the seller's payoff change by

$$dS = F(v_\alpha)((1 - \delta)(p - cq) + \delta S^I)dv_\alpha - \delta\bar{F}(v_\alpha)\bar{B}S'(\bar{B})d\alpha,⁴⁸$$

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

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

or

$$S'(\bar{B}) := \lim_{\alpha \downarrow 0} \frac{dS}{dB} = -\lim_{\alpha \downarrow 0} \frac{F(v_\alpha)}{(1 - \delta)\bar{F}(v_\alpha)}((1 - \delta)(p - cq) + \delta S^I) + S'(\bar{B}),$$

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

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

⁴⁷For α small, the buyer must visit 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 visits wp1, and so if he were to visit 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 $v_\alpha \in \mathcal{V}$.

⁴⁸Here, 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 14. *If $E_\delta \neq \{\mathbf{0}\}$, then*

$$\delta \hat{\phi} > (1 - \delta)p. \quad (13)$$

Proof. Since $E_\delta \neq \{\mathbf{0}\}$, it must be that some quality $q > 0$ occurs in some equilibrium, with continuation payoffs B^I, B^O . 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. □

Claim 15. *It holds that*

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

Further, $\lim_{B \downarrow 0} B^O$ exists and is equal to 0.⁴⁹ Similarly, the buyer's cut-off v satisfies $\lim_{B \downarrow 0} v = \underline{v}$.

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.

⁴⁹More 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.

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 visits the seller if, and only if, $v \geq \underline{v} + \varepsilon$. If the buyer is indifferent for some $v > \underline{v}$ between visiting or 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 $\underline{v} + \varepsilon \rightarrow \underline{v}$ and $B^O \rightarrow 0$, using claim 4. Hence, the seller's incentive compatibility is satisfied in case (a). 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 of payoffs, and taking $\varepsilon \rightarrow 0$ yields the desired result.

The claim that $\lim_{B \downarrow 0} B^O = 0$ is immediate, as the buyer could otherwise secure a payoff bounded above $\mathbf{E}[\tilde{v}]$ by simply not visiting in the first round. Similarly, if the cut-off $v_n \not\rightarrow \underline{v}$ along some subsequence of (B_n) tending to 0, the seller would secure a payoff bounded above $\mathbf{E}[\tilde{v}]$, given that his payoff from visiting the seller for, say, all $\tilde{v} \in ((v_n + \underline{v})/2, v_n)$ is bounded away from his payoff from visiting when $\tilde{v} = \underline{v}$. \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 $\{(B, B + S/c) \mid (B, S) \in E_\delta\}$.

Claim 16. *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 be in the strict hypograph of ϕ , a contradiction. \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 visit the seller, given some vector $(B, \phi(B))$ and corresponding equilibrium. By the previous claim, $v > \underline{v}$. In fact:

Claim 17. *Fix an equilibrium payoff vector $(B, \phi(B))$. 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 visit the seller, then we may increase B^O without loss. Hence, we may increase it either to the point at which indifference obtains (for this type $v = \bar{v}$), 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

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

Claim 18. *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 $q = 1$ (recall that this was assumed throughout for \bar{B})—and thus can be ignored. Recall from Claim 3 that also $B^O = \bar{B}$. This implies that the best equilibrium payoff for the buyer is obtained by setting $q = 1$ always and the buyer behaving myopically, contradicting Assumption A1. \square

C.2 Equilibrium Behavior for Extremal Payoffs

Given Claim 19, 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.

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

$$B = (1 - \delta) \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) + \delta B^O, \quad (14)$$

as well as

$$\frac{S(B)}{c} = (1 - \delta)F(v) \left(\frac{p}{c} - q \right) + \delta F(v) \frac{S(B^I)}{c} + \delta \bar{F}(v) \frac{S(B^O)}{c}, \quad (15)$$

and, rearranging the indifference condition that defines v ,

$$0 = (1 - \delta)(q - p) + \delta B^I - ((1 - \delta)v + \delta B^O). \quad (16)$$

Multiply (16) by $F(v)$, add the resulting equations to (14) and (15), and we get

$$\phi(B) = (1 - \delta) \int_{\underline{v}}^v (\kappa - \tilde{v}) dF(\tilde{v}) + \delta F(v)\phi(B^I) + \delta \bar{F}(v)\phi(B^O), \quad (17)$$

where $\kappa := p \frac{1-c}{c}$, $B = (1 - \delta) \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) + \delta B^O$, $q \in [0, 1]$, 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 in the proof of Claim 5. The constraint $q \in [0, 1]$ can be rewritten as, using (16)

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

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

$$\min\{\delta\phi(B^I), 1 - \delta + \delta B^I\} \geq (1 - \delta) \left(p + v - \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right) + B \geq \delta B^I. \quad (18)$$

We can rewrite (17) as

$$\phi(B) = (1 - \delta) \int_{\underline{v}}^v (\kappa - \tilde{v}) dF(\tilde{v}) + \delta F(v)\phi(B^I) + \delta \bar{F}(v)\phi \left(\frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right). \quad (19)$$

Motivated by this discussion, 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

$$\min\{\delta\phi(B^I), 1 - \delta + \delta B^I\} \geq (1 - \delta) \left(p + v - \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v}) \right) + B \geq \delta B^I, \quad (20)$$

s.t. $B \geq (1 - \delta) \int_{\underline{v}}^v (v - \tilde{v}) dF(\tilde{v})$ (a nonempty set, since we can pick $v = \underline{v}$). We note that, given that $\phi(B^I)$ is increasing in B^I over $[0, \hat{B}]$, and given that the left-hand inequality of (20) is slack for $B^I \leq \hat{B}$, we always have $B^I \geq \hat{B}$, unless the right-hand inequality of (20) binds for $v = \underline{v}$. Similarly, because $\phi(B^I)$ is decreasing in B^I over $[\hat{B}, \bar{B}]$, and because decreasing B^I relaxes the left-hand inequality for $B^I > \tilde{B}$, it holds that $B^I \leq \tilde{B}$. Accounting for the right-hand inequality of (20), we have established:

Claim 19. *It holds that*

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

Depending on the ordering between the bounds that appear in Claim 19, up to four regions can occur.⁵⁰

Claim 20. *For all $B \in [0, \bar{B}]$, $B^O \leq B$, and the inequality is strict if, and only if, $B \notin \{0, \bar{B}\}$.*

⁵⁰Plainly, if the left-hand side of (20) binds, the R.H.S. does not, and vice-versa.

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 21, 22, 23 and 24).

Definition 1. *Let:*

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

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

i.e., $B^I = \hat{B}$, $q \in (0, 1)$;

\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., $B^I \in (\hat{B}, \tilde{B})$, $q = 1$;

\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., $B^I = \tilde{B}$, $q = 1$ (and the seller IC binds).

\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 $\mathcal{B}^L, \mathcal{B}^M, \mathcal{B}^H, \mathcal{B}^{LL}$ must be unions of intervals (with \mathcal{B}^H being closed).

Claim 21. *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 cut-off 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 16.

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{B}^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

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

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

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

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,

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

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,

$$\phi'(B) = \bar{F}(v)\phi' \left(\frac{B}{\delta} - \frac{1 - \delta}{\delta} \int_{\underline{v}}^v F(\tilde{v}) d\tilde{v} \right). \quad (21)$$

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

Claim 22. *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 and strictly concave 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

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

and since the R.H.S. increases in v , and so decreases in B , $B - B^O$ decreases as well, and so, in particular, $B^O \leq B$ if this is true for the highest B in \mathcal{B}^H . 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 22) and is therefore omitted. \square

Claim 23. *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 and strictly concave 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)

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

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.⁵¹ 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.

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,

$$\phi'(B) = F(v)\phi'(B^I) + \bar{F}(v)\phi'(B^O), \quad (22)$$

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 $B^O < B < B^I$. \square

Claim 24. *The set \mathcal{B}^{LL} is a subset of $[0, \hat{B})$. Further, $B^O < B < B^I \leq \hat{B}$ for $B > 0$. The function ϕ is differentiable and strictly concave on any interval in \mathcal{B}^{LL} .*

Proof. Differentiability is shown as before. Given that $q = 0$, $B^I \leq \hat{B}$ (otherwise, increase q and decrease B^I , keeping v constant: this would increase the seller's payoff and therefore

⁵¹The case in which $v = \bar{v}$ because $\frac{\kappa - a(1-p)}{1-a} > \bar{v}$ is analogous.

$\phi(B)$). 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 $B^I > B$ for $B > 0$, and given that $B^I \leq \hat{B}$, 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 16, $v > \underline{v}$, 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 23. Hence we must have $B^O < B$ for $B > 0$. \square

By elimination, given Claims 11, 13 and 14, for all $B > \tilde{B}$, $B \in \mathcal{B}^H$, and so $q = 1$. But obviously, \mathcal{B}^H cannot be the only interval, else $q = 1$ after all histories, and the buyer would behave myopically. Note that $B \geq \tilde{B}$ implies that the incentive compatibility condition is violated for $B^I = B$ if $q = 1$. To sum up:

Claim 25. *The set \mathcal{B}^H includes $[\tilde{B}, \bar{B}]$ (but does not include 0). Further, on $[\tilde{B}, \bar{B}]$, it holds that $B^I = \tilde{B} \leq B \leq B^O$ (with each equality holding only at one of the endpoints).*

The existence of a unique $B > \tilde{B}$ (the \check{B} from Lemma 2) such that $B^O = B^I$ (and so equal to \tilde{B}) follows from this claim.

C.3 Asymptotic Behavior

Claim 26. *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 22, $\min\{B_n^I, B_n^O\} < B_n$, $\min\{B_n^I, B_n^O\} \leq \tilde{B}$ 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 21 and 23, 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 $\mathcal{B}^L \cup \mathcal{B}^M$,⁵² 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

⁵²And $B - B^O$ 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 $B^O = B$.