Dynamic Screening: Why Weight isn't Volume

Johannes Hörner and Anna Sanktjohanser*

May 9, 2024

Abstract

We consider a dynamic buyer-seller interaction. Instead of the buyer's valuation, it is the *frequency* with which he needs to trade that is the buyer's private information. The difference matters. With commitment, in particular, full surplus extraction is possible, via, for instance, limited-time offers. Without commitment, ratcheting is mitigated, as not buying is not necessarily a sign of strength. Because time is informative, the seller learns and may adjust her behavior over time. She starts with a pooling offer, before occasionally experimenting with separating offers. The seller's payoff is not monotone in her belief about the buyer's type.

KEYWORDS: Loyalty, Imperfect monitoring, Repeat customers.

JEL CLASSIFICATION NUMBERS: C72, C73, C78.

1 Introduction

Buyers come in all shapes and sizes, just like the products they purchase. Most of the economic literature on adverse selection has focused on the case in which consumers differ in their taste for quality or quantity –attributes of the product that a seller supplies. The starting point of this paper is that buyers also differ in how often they need to trade, an inherently dynamic trait that cannot be reduced to the analysis of Mussa and Rosen (1978). While largely overlooked by the economics literature, the importance of heterogeneity in repeat-buying has been emphasized in marketing (Ehrenberg, 2000). The literature segments buyers into *heavy* vs. *light* buyers. A heavy (category) airline customer, for instance, is someone who often travels across

*Hörner: CNRS (Toulouse School of Economics), 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. We thank Jim Dana for many helpful conversations, as well as Marina Halac, Michihiro Kandori, Ali Shourideh, Jean Tirole, Alex Wolitzky and especially Juuso Toikka for enlightening discussions. Johannes Hörner and Anna Sanktjohanser acknowledge funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program).

the country, whereas a light (brand) customer might be an individual who occasionally travels on a single route using the same airline.

The purpose of this paper is to understand how this unobserved heterogeneity shapes the seller's pricing strategy, the buyer's welfare and the dynamics of the relationship. The difference between weight and volume (quantity, as in Mussa and Rosen) matters. Results markedly differ, whether the seller commits or not.

To highlight the role of buyer purchase frequency, we abstract from the usual heterogeneity in valuations: consumers place the same value on all attributes of the good. However, because heavy buyers anticipate more frequent needs, they are more willing to experiment with competing products, to find those that are most suitable. If the seller wishes to deter them from doing so, and so from losing them eventually, she must supply a requisite premium. Hence, although consumers share the same value for the good, their heterogeneity in purchase frequency endogenously generates heterogeneity in their reservation utility. Unlike in Mussa and Rosen, it is unclear what kind of customer is more desirable: heavy buyers purchase more often, but entail a lower markup.

Suppose first that the seller commits. Surprisingly, the seller can extract all the buyer's surplus, yet guarantee that the buyer trades with the seller whenever he needs to trade. That is, unlike in the Mussa-Rosen paradigm, here, the buyer derives no rent from his private information. Many pricing strategies achieve this. We focus on a simple one that is prevalent in practice: limited-time offers. Buying triggers a countdown timer. Attractive terms are offered to a buyer who takes action before the window closes. Those terms are attractive to a heavy buyer, but the window's brevity is dissuasive for a light buyer.[1]

Under non-commitment, the seller learns from the buyer's purchase behavior. On

[1]Leaving no rents to the buyer is optimal for the seller if her profit is linear in the utility that she supplies, as is the case if price is one of the instruments that she controls. If she cannot adjust the price, it might not be profit-maximizing to leave no rents to the buyer, as this requires the utility to vary over time. If the seller's profit is strictly concave in the utility she supplies, this might be unappealing.

the one hand, no trade can be statistical evidence that the buyer has a reservation utility that is too high given the seller's deal and, hence, that he is a heavy buyer. On the other hand, a heavy buyer might have sampled alternatives: conditioning on the event that he is not a lost customer is conditioning on the event that he had few opportunities to do so, suggesting that he is a light buyer. Hence, unlike with unobserved valuations, foregoing trade is not a show of strength and does not automatically trigger more advantageous offers. This curtails ratcheting and facilitates screening.

We solve for the Markov perfect equilibrium of the game, using the seller's belief as a state variable. This belief may drift up or down in the absence of trade, depending on the strategy she follows, and which buyers she targets. In particular, it can drift down whenever the seller caters to both types, but drift up whenever the seller targets light buyers only (the "no news is good news" scenario). At the juncture of these two events, the seller intersperses pooling offers, which if accepted leads to further pooling, at least for a spell, with separating offers, which only a light buyer might accept. Attracting the light buyer is not cheap, however, since such a buyer anticipates that more attractive deals will occasionally arise if she turns down a separating offer.

For low beliefs, targeting the light buyer is attractive, but involves taking a chance of losing the buyer if the buyer is heavy, a cost that increases with optimism (higher belief). For high beliefs, catering to both buyers' types is optimal, but repeat buying does not often occur if the buyer is light, a revenue loss that increases with pessimism. This leads to a non-monotone (quasiconvex) seller's payoff as a function of the belief she attaches to a buyer being heavy.

These peculiarities in the seller's payoff and strategy differ markedly from the predictions of the model with heterogeneous values. Hart and Tirole (1988) show that, absent commitment, the standard ratchet effect forces the seller to make a pooling offer, at least when the buyer is sufficiently patient, the case on which they focus.[2]

[2] Otherwise, she makes a separating offer that only the light type accepts, leading to perfect price discrimination from the second round onward.

Here, these Coasian forces are diminished because the seller learns about the buyer's type over time. This is either because the buyer was supposed to visit if the opportunity arose but failed to do so, or because the buyer was likely to find a suitable alternative yet remains in the market. A detailed comparison is carried out in Section 4.6, accounting for the differences in settings across papers.

A snapshot description of the formal model is as follows. The game is cast in continuous time. The buyer needs to trade at random times that are private, drawn according to a Poisson process whose intensity can be either high or low, depending on the buyer's type. When the buyer needs to trade, the buyer can either turn to the seller or to an outside option. In the latter case, this outside option is suitable with some exogenous, i.i.d. probability, in which case the buyer exits the relationship (and ends the game), or not. The seller observes all the buyer's trades with her, and nothing else. Given her information, she makes an offer at every point in time, directly formulated in terms of the utility it provides to the buyer, which the buyer observes if he needs to trade at that moment.

With commitment, we may invoke the revelation principle to simplify the seller's problem, and we simply exhibit a scheme that extracts all the buyer's surplus, while ensuring that the buyer always wants to trade with the seller (Theorem 1). In the absence of commitment, attention is restricted to Markov perfect equilibria. That is, the deal the seller offers is a function of her belief about the buyer only. The Markov perfect equilibrium (MPE) need not be unique. While our description above pertains to the arguably more interesting MPE of the game (see Theorems 2 and 4), there might exist another equilibrium, in which the seller makes a pooling offer for every belief she may hold (see Theorems 3 and 5). Depending on the parameters, one or the other equilibrium (or both) might exist. Focusing on the limiting equilibrium of the finite-horizon game as the horizon length increases restores uniqueness.

Examples of limited-time offers abound. Such deals are common in many markets, from retail to the service industry.[3] Without commitment, our results address the

[3]For instance, Burger King or Taco Bell regularly offer buy-one-get-one-free deals to repeat

personalization of pricing and services in business-to-business (B2B) relationships. For instance, firms that make regular purchases over a certain period receive discounted prices on future orders, priority customer service or extended payment terms.[4]

There are a number of papers looking at the economics of repeat customers with one-sided private information. In these papers, the private information pertains to the buyer's value for the good and flexible transfers are typically available. In contrast, in this paper, consumers do not differ in their value for the good but in the frequency of their consumption need. Regardless of their frequency of trade, buyers in our model prefer higher over lower quality, and they receive the same utility from any given provided quality. Moreover, we do not allow for flexible transfers. Crémer (1984) analyzes a two-period model in which consumers learn their own value for the good through consumption, and transfers are available. Kennan (2001) analyses repeated contract negotiations between a buyer and a seller, where the buyer has persistent private information about his value for the good. Hart and Tirole (1988) analyze the optimal contract between a seller and a buyer whose (persistent) valuation for a per period service is private. Battaglini (2005) looks at the optimal long-term contract when consumers' preferences evolve following a Markov process.[5]

The type-dependent reservation utility of the buyer clearly plays a key role in the

customers through their mobile app. Similarly, the gym chain Orangetheory Fitness regularly offers deals where purchasing a certain number of classes entitles the consumer to a complimentary future session. Supercuts, a chain offering haircuts, regularly offers a discount on a future haircut if the customer books within a specific timeframe. While we are not aware of economic papers on firms screening by purchase frequency specifically, there are a number of papers on time-limited offers being used as a screening device: among others, see Baik and Larson (2023), Chevalier and Kashyap (2019), Gerstner and Hess (1991), Gerstner et al. (1994), and Nevo and Wolfram (2002).

[4]While companies offering such personalization are often not transparent about these schemes, or may provide such customized pricing and services through individual negotiations or account management, there are a number of companies known to offer such personalization. For instance, Amazon Business (the B2B arm of Amazon), Office Depot Business Solutions (the B2B arm of Office Depot), Grainger (a large industrial supply distributor), and Cisco (a technology and networking equipment provider) all offer customized pricing and services based on purchase patterns. While largely overlooked by the economics literature, several marketing papers examine the link between quality and purchase frequency. See Athanasopoulou (2009) for a survey.

[5]Further references that are more tangentially related are provided in Fudenberg and Villas-Boas (2006) in their survey article on behavior-based price discrimination. Hosios and Peters (1993) is also tangentially related.

analysis. The difficulties and peculiarities that type-dependent reservation utilities raise for adverse selection models are thoroughly examined in Jullien (2000).

2 Model

Time $t \geq 0$ is continuous and the horizon is infinite. There are two players: one *buyer* (he) and one *seller* (she).

This is a game of incomplete information. The buyer is one of two types, which is private information. Either the buyer is a *heavy* (h-type, for short), or a *light* buyer (l-type). Heavy buyers have more frequent needs to trade. More specifically, the buyer needs to trade at times that are exponentially distributed, with parameter λ^h, λ^l, depending on the type, with $\lambda^h > \lambda^l > 0$.

Arrival times are private information as well. A buyer who needs to trade at time τ has the choice between the seller's offering, and an outside option. We follow Schmalensee (1982) by modeling the outside option as the offering of an alternative seller, whose good's quality or fit is uncertain.[6] If he picks the outside option, with probability $\alpha \in (0, 1)$, the buyer exits the game. This is interpreted as the event in which the alternative seller who is sampled turns out to be a good match, and so the buyer becomes a lost customer for the seller. Such an event is i.i.d. across those times –think of a large number of possible alternatives, of which the outside option is the reduced-form, each of which is equally likely to be suitable given our buyer's preferences. If the match is bad, the buyer remains attentive to the seller's future offerings.

Hence, we have in mind a perishable product that is an experience good, so the only way that the consumer can resolve uncertainty about its fit is to purchase a brand and try it, and for simplicity, in the product class considered, a product "works" or "does not work" for the consumer, e.g., stainless steel razor blades. That uncertainty no longer applies to the seller in our game, whose suitability for the buyer is already established.

We have assumed that a buyer who eschews trading with the seller necessarily samples

[6] See also, in particular, Villas-Boas (2004).

the outside option. This is for convenience: it is readily checked that, in equilibrium, the buyer prefers sampling the outside option rather than sitting out. Conversely, we have assumed that a buyer cannot pretend to have a need when he does not. This is more restrictive, but innocuous provided that the utility from consuming a unit when it is not needed is low enough. To be sure, taking a flight at the end of the year to maintain or improve one's tier is not unheard of, but it is the exception rather than the rule. Finally, we assume that a buyer who finds a suitable match elsewhere never returns –a simplification to be sure, but one that captures the conventional wisdom that "acquiring a new customer is anywhere from five to 25 times more expensive than retaining an existing one" (Gallo, 2014).

At any time t, the seller offers a distribution over utilities $\mathcal{U}_t$. Because the buyers' types do not differ in their valuation for the good, we do not need to take a stance regarding the actual levers that the seller uses to control this utility, which might include price, quality, a quantity/price schedule, etc. It could also depend on other random variables, such as the actual volume the buyer wants to buy on that particular occasion. All that matters is that the buyer's type does not affect his preference over these instruments, so that the derived expected utility is a sufficient statistic. A buyer who needs to trade gets to observe both the distribution over utilities and the realized draw from this distribution, before choosing between taking this offering, or the outside option. In case the seller offers a nondegenerate distribution over utilities, she does not see the realization that the buyer has drawn, in case the buyer had a need to trade (which she does not observe, as mentioned), *unless* the buyer takes this offering.

All rewards are discounted at the common interest rate $r > 0$. Whenever the buyer takes the offering $\mathcal{U}_t$, this is his reward. If the buyer samples the outside option, he obtains some fixed utility $\mathcal{U}^o > 0$ in the case of a good match, that is, with probability α, and 0 otherwise.[7] If the match is good, the buyer exits the game, and

[7] We take the utility of a good match to be type-independent, in line with our focus on the trading frequency as the only feature distinguishing heavy from light buyers. Picking this utility constant is convenient, but inessential. We might as well take the net present value of a good match as a primitive, at the cost of introducing additional parameters.

his continuation payoff is the discounted sum of these utilities $\mathcal{U}^o$, which he obtains whenever the need for a trade occurs. That is, his expected continuation payoff is then[8]

$$\mathbf{E}\left[\sum_{\tau_n} e^{-r\tau_n}\mathcal{U}^o\right] = \frac{\lambda^k}{r}\mathcal{U}^o =: Z^{o,k},$$

with $k = h, l$, where $(\tau_n)_{n\in\mathbf{N}}$ are the arrival times of the buyer. Note that this continuation payoff is higher for the h-type. That is, the reservation utility is type-dependent, as a consequence of the difference in arrival rates.

The seller's reward when her offering $\mathcal{U}_t$ is accepted by the buyer is denoted $\Pi(\mathcal{U}_t)$. We assume that Π is continuous, strictly decreasing in $\mathcal{U}$ and strictly positive under complete information (see below). In Section 4.4, concerned with equilibrium uniqueness, and for the figures, we assume that $\Pi(\mathcal{U}) = R - c\cdot\mathcal{U}$, for some constants $R, c \in \mathbf{R}_+$. The buyer and seller's rewards at every instant when no trade takes place for either party are normalized to 0. The players' objective is to maximize their payoff, that is, the expected discounted sum of rewards.

We conclude this section with the benchmark of complete information. The seller has no incentive to offer more than the bare minimum which keeps the buyer from sampling the outside option. Therefore, unless she chooses to let him go, the utility $\mathcal{U}^k$ the seller offers to the buyer of type k solves

$$\mathcal{U}^k + \frac{\lambda^k}{r}\mathcal{U}^k = \alpha\big(\mathcal{U}^o + Z^{o,k}\big) + (1-\alpha)\frac{\lambda^k}{r}\mathcal{U}^k.$$

Indeed, by the one-shot principle, the buyer is indifferent between staying with the seller forever, or trying the outside option once, and reverting to the seller forever in case of a bad match. Solving,

$$\mathcal{U}^k = \frac{\lambda^k + r}{\alpha\lambda^k + r}\alpha\mathcal{U}^o. \tag{1}$$

The utility level $\mathcal{U}^k$ that the seller must offer to the buyer is higher if the buyer is known to be heavy (as $(\lambda^h + r)/(\alpha\lambda^h + r) > (\lambda^l + r)/(\alpha\lambda^l + r)$). While it might well be that, in many applications, the outside option $\mathcal{U}^o$ is also type- dependent, with

[8] We omit the standard definitions of outcomes –infinite histories– and realized payoffs.

heavy buyers having better outside opportunities, we do not need to assume as much: their higher trade frequency already makes experimentation with other sellers more attractive to them.

Let $\Pi^k := \Pi(\mathcal{U}^k)$ denote the seller's reward from serving a buyer of type k. Because $\mathcal{U}^h > \mathcal{U}^l$, $\Pi^h < \Pi^l$. We assume that $\Pi^h > 0$, so the seller has no incentive to exclude either type of buyer under complete information. Because $\Pi^h < \Pi^l$, but a buyer of h-type comes more often than a buyer of l-type, the seller might find either buyer type more desirable:

$$\frac{\lambda^h}{r}\Pi^h \lesseqgtr \frac{\lambda^l}{r}\Pi^l.$$

Hence, a higher belief is not necessarily "good news," regarding the seller's payoffs, although it is convenient to describe higher beliefs this way, as we do.

3 Commitment

Here, we assume the seller can commit. Then, without loss, we may assume that the optimal mechanism has the buyer reveal his type, and as a function of this report, specify a utility level after each sequence of arrival times. Formally, a mechanism is a pair of maps $\mathcal{U}_h, \mathcal{U}_l$, where, for $\theta = h, l$,

$$\mathcal{U}_\theta \colon \cup_{n\in\mathbf{N}} \mathbf{R}_+^n \to \mathbf{R}$$

$$(\tau_1, \tau_2, \ldots, \tau_n) \mapsto \mathcal{U},$$

where τ_j is the interarrival time between the $j-1$-th visit and the j-th visit. That is, $\mathcal{U}_\theta(\tau_1, \tau_2, \ldots, \tau_n)$ is the utility provided at time $t_n = \sum_{j=1}^n \tau_j$ to a visiting buyer who reported type θ, and already visited at times $t_j = \sum_{i=1}^j \tau_i$, $j < n$. Without loss, the mechanism satisfies incentive compatibility; that is, type θ prefers reporting truthfully.

It is tempting to think that, under commitment, the problem is "static." However, the seller does not simply want to frontload the payment (say, if this is one of the instruments determining utilities) to the first encounter. To see why this is suboptimal, note that she then either rations one of the types (one of the prices involving a cap on the "net present number of future times" the buyer comes for the service), in which case she leaves some surplus on the table (there are histories after which either type arrives very often in a short span of time), or she does not, but then she has to charge the same price to both types, again leaving some surplus to at least one of the buyer's types. Backloading is not an option either, both because the buyer's needs can't wait –the utility must be delivered when he needs it– and because the buyer is not committed –if payment was an instrument available to the seller, the buyer would just leave the seller for the outside option if she had a debt that increased sufficiently over time.

Yet, with commitment, the seller can reduce the buyer's utility to his payoff under complete information, *and* have the buyer always visit the seller. In the case of linear cost, this implies that the seller receives the same payoff as under complete information, i.e., she leaves no rents to the buyer.[9]

The construction is informally described as follows. The buyer first truthfully reports his type: if he reports light, he obtains $\mathcal{U}^l$ whenever he visits the seller, which, as a light type, he does whenever the opportunity arises. If he reports heavy, he obtains the following deal: at the first visit, he obtains utility $\mathcal{U}^I$, to be determined; if he then visits again *within* a delay of $T \in \mathbf{R}_+$, he obtains utility $\mathcal{U}^F$. As soon as he does, or if he fails to visit within this delay, he faces the same deal as was offered to start with after his initial report: the utility $\mathcal{U}^I$ is supplied at the next visit, followed by $\mathcal{U}^F$ if another visit takes place within T of that next visit, etc.[10]

We choose the triple $(T, \mathcal{U}^I, \mathcal{U}^F)$ so that the heavy type is indifferent between visiting the seller or not when visiting the seller entails utility $\mathcal{U}^I$; he strictly prefers to visit the seller over the outside option if a visit entails utility $\mathcal{U}^F$. For the light type, the

[9]If the cost is strictly convex, this conclusion does not follow –indeed, it is is critical in the construction below that the utility level provided to the buyer is not constant over time –hence, the seller's average cost is not equal to the cost of supplying the average utility.

[10]Note that, with this scheme, the buyer obtains $\mathcal{U}^F$ at best half the time. Perhaps it would be more realistic to assume that *any* visit opens a window of length T over which $\mathcal{U}^F$ is supplied if a visit ensues, whether or not $\mathcal{U}^F$ has already been supplied at the last visit. We chose our formulation for convenience, but many alternative mechanisms exist, with the desired properties.

same holds: intuitively, the delay T is chosen to be sufficiently short such that the light buyer discounts the opportunity to take advantage of the better deal $\mathcal{U}^F$.

Hence, time is used as a screening device. This is reminiscent of contract theory with hard evidence: here, the light buyer cannot replicate the probabilistic frequency of visits of the heavy type, and the seller can use this to eliminate rents that would otherwise arise. However, note that the same result also holds if the buyer can trade even when he has no need, provided that he derives little enjoyment from such a trade.

Solving for such a triple is a matter of simple algebra. Fix $T > 0$. Let us denote by Z_t^λ the *ex ante* payoff of a buyer with arrival rate $\lambda = \lambda^l, \lambda^h$, who visits the seller whenever he has the chance to do so, and has another $t \leq T$ time-to-go over which the next visit would yield utility $\mathcal{U}^F$. Therefore, Z_T^λ is the payoff when the buyer just visited the seller and derived utility $\mathcal{U}^I$, and Z_0^λ is his payoff when this time-to-go runs out, so that the next utility is (back to) $\mathcal{U}^I$. Keeping in mind that time-to-go decreases as time passes by, Z_t^λ satisfies

$$(r+\lambda)Z_t^\lambda + \dot{Z}_t^\lambda = \lambda(\mathcal{U}^F + Z_0^\lambda), \tag{2}$$

with boundary condition

$$Z_0^\lambda = \frac{\lambda}{\lambda + r}(\mathcal{U}^I + Z_T^\lambda). \tag{3}$$

The solution of this differential equation is

$$Z_0^\lambda = \frac{\lambda}{r}\left(\mathcal{U}^I + \frac{\lambda(e^{(\lambda+r)T} - 1)}{e^{(\lambda+r)T}(2\lambda + r) - \lambda}(\mathcal{U}^F - \mathcal{U}^I)\right). \tag{4}$$

Given T, we then pick $\mathcal{U}^I$,$\mathcal{U}^F$ to satisfy $Z_0^{\lambda^l} = \frac{\lambda^l}{r}\mathcal{U}^l$, $Z_0^{\lambda^h} = \frac{\lambda^h}{r}\mathcal{U}^h$, that is, both types are held to their outside option; in particular, lying does not benefit either type at the reporting stage. It is readily shown that $\mathcal{U}^F \geq \mathcal{U}^I$ (see appendix).[11] Hence, either type strictly prefers to visit the seller if he gets the opportunity to do so during the window of length T that opens after a visit. Hence, conditional on reporting heavy, either type of buyer visits the seller whenever the opportunity arises. We summarize

[11]The difference $\mathcal{U}^F - \mathcal{U}^I$ is actually decreasing in T. A seller with convex cost prefers longer windows T.

this discussion with the following theorem.

Theorem 1 *There exists a mechanism such that the buyer comes whenever he has a need to trade, and is held down to his reservation utility, independent of his type.*

As mentioned, this mechanism is not unique. The one we have focused on has the benefit of being simple to describe, common in practice, and rich enough that we can choose its parameters so that the players' payoffs are positive at all times.

If the seller's profit function $\Pi(\cdot)$ is strictly concave, extracting all surplus is not necessarily desirable for the seller. While solving for the optimal mechanism is beyond the scope of this paper, a few observations simplify the seller's problem. For inference purposes, conditional on the buyer's n-th arrival occurring at time t, the exact earlier arrival times are uninformative. Hence, the seller might as well offer utility functions $\{\mathcal{U}^n(\cdot)\}_{n\in\mathbf{N}}$ to the h-type, which specify a utility level he gets if his n-th arrival occurs at time t. (Meanwhile, it is optimal for the seller to offer $\mathcal{U}^l$ always to the professed l-type.) Indeed, if the buyer could commit (so that only his *ex ante* individual rationality constraint must be satisfied), one can characterize these functions; not surprisingly, they are decreasing functions of time, which implies that the buyer would want to renege and take the outside option, at least after some histories. Such *ex post* inefficient exclusion serves screening, and is *ex ante* desirable from the seller's point of view, to an extent that depends not only on her cost, but also on her revenue. That is, the optimal mechanism also specifies times $\{T(n)\}_{n\in\mathbf{N}} \in \overline{\mathbf{R}}_+$ such that the buyer must have arrived at least n times by time $T(n)$ for the seller to offer a utility level that is acceptable to either type. In turn, this significantly complicates the calculation of the probability that the buyer arrives for the n-time at time t, since he will have had to fulfill the time-limits during the meantime. Hence, we do not push the analysis further here.

We conclude this section with four remarks.

Commitment: Commitment is not actually required to achieve the outcome described in Theorem 1. As is often the case in this literature, non-commitment only bites to the extent that some refinement (usually, Markov perfection) is imposed. In the spirit of the durable-goods literature, here, reputational equilibria can be devised that implement the commitment solution, relying on the threat of reversion to some Markov equilibrium. However, as discussed in Section 4.4, such equilibria are not robust to a finite horizon, no matter how long.

Storable Units: Theorem 1 relies critically on units being nonstorable. If the buyer can store units, buying and consuming them as he wishes, there is no longer a distinction between weight and volume. The analysis of Mussa and Rosen applies, and rents must be granted to the buyer given his private information.

More Types: Here, we have focused on two types only. One might wonder whether the analysis extends to more types. In particular, consider an intermediate type, with arrival intensity $\lambda^m \in (\lambda^l, \lambda^h)$. Because the expected number of needs of such a buyer's type is a convex combination of the expected number of the other two types, he could "mimic" such a convex combination by mixing his initial report. This suggests that he might be able to secure rents in this fashion. Yet, this is not the case. The contract can depend on richer statistics than the expected number of needs. In particular, by conditioning on, say, the variance in the number of units demanded, the seller could separate a truthful m-report from the mixture over l- and h-reports. Formally, a type defines a distribution over sample paths of realized times at which needs arise, and the space of sample paths is sufficiently rich not only to distinguish three types, but, indeed, an arbitrary continuum of arrival intensities. Of course, the richer the set of types, the richer the set of contracts.

Formally, in Online Appendix C, we extend Theorem 1 to the case of a bounded interval of types $\lambda \in [\underline{\lambda}, \overline{\lambda}]$, $0 < \underline{\lambda} < \overline{\lambda}$. A mechanism exists such that the buyer comes whenever he has a need to trade, and is held down to his reservation utility, independent of his type. The contract is a simple variation of the one we have used with two types only. It involves an initial utility $\mathcal{U}^I(\hat{\lambda})$, given a report $\hat{\lambda} \in [\underline{\lambda}, \overline{\lambda}]$, and a second higher utility $\mathcal{U}^F(\hat{\lambda})$ if such a demand arises during a window that closes either when this demand materializes, or at a random time, occurring at a constant rate $\kappa \cdot \hat{\lambda}$, for some $\kappa > 0$ (as with two types, this deal repeats over time). Theorem 1 might be reminiscent of the full surplus extraction result of Crémer and McLean (1988). Indeed, both rely on a separation argument, but it is not obvious that the analogy extends further.

Deterministic vs. Stochastic Needs: Of course, if the set of types is the set of sample paths itself, separation is no longer possible without conceding rents. This is what happens if the buyer has foreknowledge of the timing of his needs, and the support of this set of possible types is rich. On the other hand, if the set of types is small enough (e.g., one type wishes to receive the newspaper every day, the other type only cares about the weekend edition), separating types without granting rents remains possible.

4 No Commitment

4.1 Equilibrium Concept

To define the non-cooperative game, additional notation is needed. A seller's history h_S^t specifies the distribution of utilities she has offered at all times $s < t$, the times $(\tau_n)_n$ at which the buyer traded with the seller and the realized utility the buyer drew at those times. A buyer's history h_B^t includes, in addition to h_S^t, the times at which the outside option was sampled, unsuccessfully, as well as the realized seller's utility whenever he had a trade need (including possibly at time t). We ignore histories that occur after the buyer exits, if he does.

We assume that exit is observable. This is without loss of generality. Even if this is rarely the case in reality, it does not matter: the seller will not have to deliver on the promised utility ever again if exit has occurred, so it is irrelevant what she does after the buyer exits. Hence, the seller might as well condition on the event that exit has not occurred, which is equivalent to assuming that exit is observable.

At any time t (throughout, conditional on no exit), given the seller's history h_S^t, the seller has a belief $\mu_t = \mathbf{P}(\lambda = \lambda^h \mid h_S^t)$ that the buyer is of the h-type. Throughout, it is more convenient to work with the likelihood ratio $\ell_t := \ln \frac{\mu_t}{1-\mu_t}$, despite the cost of having to deal with the domain $\overline{\mathbf{R}}$, the extended real line: we interpret $\ell_t = \pm\infty$ as a degenerate belief –the seller believes with probability one that the buyer is the h- or l-type. With some abuse we also refer to ℓ_t as the belief. The prior belief is denoted μ^0, or ℓ^0.[12]

We focus on equilibria in (stationary) Markov strategies, with the seller's belief as state variable. We omit the definition of strategies that are not Markov, but note that, as is standard, equilibria in Markov strategies remain equilibria when a larger class of strategies is allowed –in particular, players are allowed to deviate to non-Markov strategies. From here on, we drop the qualifier Markov.

The seller's strategy is a map $\sigma_S : \overline{\mathbf{R}} \mapsto \Delta(\mathbf{R}_+)$, mapping her belief into a distribution over (positive) utilities.[13] Without loss, we restrict attention to distributions with finite support. Measurability assumptions are made below (see footnote 14).

A buyer's strategy is a map $\sigma_B^k : \overline{\mathbf{R}} \times \mathbf{R}_+ \mapsto \Delta(\{In, Out\})$, specifying for each time at which an arrival might occur, a (possibly random) choice between the outside option and the seller's offering, as a function of his type $k = h, l$, the seller's belief ℓ, and the realized seller's offering $\mathcal{U}$. Specifically, we let $\sigma_B^k(\ell, \mathcal{U})$ denote the probability that the buyer picks the outside option.

Absent any trade, the belief evolves according to the differential equation[14]

$$\dot{\ell}_t = \sum_{\mathcal{U}} \mathbf{P}(\mathcal{U})\big((1 - \sigma_B^l(\ell_t, \mathcal{U})(1-\alpha))\lambda^l - (1 - \sigma_B^h(\ell_t, \mathcal{U})(1-\alpha))\lambda^h\big), \qquad (5)$$

[12]Equivalently, we may assume that ℓ^0 is the seller's belief when a buyer first approaches her, an assumption that is perhaps more palatable for applications.

[13]While the notation does not account for this, we allow duplicate utilities, so to speak, to avoid issues of existence: that is, the seller is allowed to offer two identical utility levels, with labels that induce different behaviors by the buyer: one, say, that both types accept, and one that only one type is supposed to accept.

[14] Requiring (5) to be Lipschitz continuous in ℓ is too strong an assumption, as it turns out: stationary points of the belief dynamics are (sometimes) reached in finite time, so that the assumption fails in interesting cases.[15] We assume that the RHS of (5) is one-sided Lipschitz continuous in ℓ, so that, given ℓ_t, (5) has at most a unique solution on $[t, t+\varepsilon]$, for some $\varepsilon > 0$, and all ℓ_t and t. More precisely, we assume that the domain of the belief, $\overline{\mathbf{R}}$ admits a finite partition in intervals, possibly including points, such that the RHS of (5) is one-sided Lipschitz continuous on each nondegenerate interval. Primitive assumptions on σ^B, σ^S that ensure that this condition is satisfied are omitted.

given any posted lottery over utilities, $\{\mathbf{P}(\mathcal{U}) : \mathcal{U} \in \mathbf{R}_+\}$. To better understand this equation, consider the special case in which the seller assigns probability one to utility $\mathcal{U}$ only, which the l-type is supposed to accept (in case he gets the opportunity to trade) and the h-type is supposed to reject. No-trade is an event of probability $1 - \alpha\lambda^h \,\mathrm{d}t$ in case the buyer is an h-type, and of probability $1 - \lambda^l \,\mathrm{d}t$ if he is an l-type, so that the likelihood ratio satisfies

$$\frac{\mathrm{d}\ell}{\ell} = \frac{1 - \alpha\lambda^h \,\mathrm{d}t}{1 - \lambda^l \,\mathrm{d}t} \simeq (\lambda^l - \alpha\lambda^h) \,\mathrm{d}t.$$

Hence, the belief drifts up or down depending on the sign of $\lambda^l - \alpha\lambda^h$. Indeed, no trade is good news: on the one hand, the l-type is less likely to get opportunities to trade, and so the absence of trade is bad news (in terms of the probability of the h-type). On the other hand, exit has not occurred, and exit could have occurred if the buyer was an h-type, so no-exit is good news in this regard. This dichotomy plays an important role in the sequel, as $\lambda^l - \alpha\lambda^h < 0$ implies that, in equilibrium, beliefs always drift down in the absence of trade. We call this scenario the *no news is bad news* case.

In case the buyer trades with the seller, the seller's belief jumps, in a direction and to an extent that reflects the buyer's strategy. More precisely, given the realized $\mathcal{U}$, and belief ℓ_t, the updated belief is

$$\ell_{t+} = \ell_t + \ln \frac{\lambda^h(1 - \sigma_B^h(\ell_t, \mathcal{U}))}{\lambda^l(1 - \sigma_B^l(\ell_t, \mathcal{U}))}, \tag{6}$$

with the convention that $\ell_{t+} = \pm\infty$ in case $\sigma_B^k(\ell_t, \mathcal{U}) = 0$ for exactly one of $k = h, l$. If both types are supposed to reject the offer, which is the only off-path event that the buyer can trigger, the posterior belief could be set arbitrarily. For definiteness, with one exception, the equilibria we construct under no commitment assume that such an offer is interpreted as evidence that the buyer is of the l-type.[16] Note that this belief is common knowledge. This is why we assume that the buyer observes the distribution of utilities that the seller offers, in addition to the realized utility, so that the buyer's information includes everything the seller knows.

[16] Refinements are discussed below. Note that updating (5)–(6) already encapsulates "no signaling what you don't know," as the seller's action is not informative *per se*.

An equilibrium refers to a Markov perfect equilibrium. That is, fixing the equilibrium strategies, and so the updating rules (5)–(6), we may view the game as a stochastic game with state variable ℓ, and require the players' strategies to be mutual best-replies; that is, given any time t, any history up to that time, and the resulting belief ℓ_t, each player's continuation strategy maximizes his or her expected payoff, conditional on the equilibrium strategies and the transition rules (5)–(6).

It is worth pointing out that, perhaps surprisingly, the equilibrium outcome under complete information (described at the end of Section 2) is not the unique (Markov) equilibrium outcome of the game with incomplete information, when the seller believes the buyer to be of type k with probability one. Some refinement must be imposed to rule out other, arguably less plausible equilibria. We return to this issue in subsection 4.4. Until then, we describe equilibria of the game with incomplete information in which the equilibrium outcome of the continuation game with degenerate belief coincides with the outcome of the game with complete information.

4.2 No News is Bad News

This section assumes that $\lambda^l < \alpha\lambda^h$: that is, when only the l-type is supposed to trade, absence of trade is bad news, even if the concomitant absence of exit is good news: the belief ℓ drifts down at rate $\alpha\lambda^h - \lambda^l$ in this event. *A fortiori*, if both types are supposed to trade, absence of trade is bad news too: the seller's belief then drifts down at the higher rate $\lambda^h - \lambda^l$.

A natural equilibrium candidate involves the seller using a cutoff strategy: she offers $\mathcal{U}^h$, and only $\mathcal{U}^h$, whenever $\ell \geq \ell^*$, for some ℓ^* to be determined, and she offers $\mathcal{U}^l$ at lower beliefs. The h-type buyer accepts any offer that is at least equal to $\mathcal{U}^h$; hence, in equilibrium, he accepts offers as long as the belief remains above ℓ^*. The l-type buyer accepts any offer above $\mathcal{U}^l$ for $\ell < \ell^*$, and uses a cutoff $\check{\mathcal{U}}(\ell) \leq \mathcal{U}^h$ that is nondecreasing in ℓ for $\ell > \ell^*$, with $\check{\mathcal{U}}(\ell^*) = \mathcal{U}^l$. Its exact value does not matter here (see (15)): it makes such a buyer indifferent between accepting the offer, thereby disclosing his type, and thus obtaining $\mathcal{U}^l$ forever in the continuation game, or rejecting the offer, but enjoying the opportunity to trade at utility $\mathcal{U}^h$ in the time window until the belief drifts down past ℓ^*, at which point he will content himself with $\mathcal{U}^l$.[17] Fortunately, along the equilibrium path, the l-type buyer will be spared this dilemma, as for $\ell > \ell^*$ he can expect to get $\mathcal{U}^h$. This calculation implicitly assumes that the seller interprets any acceptance of an offer strictly below $\mathcal{U}^h$ as a clear indication that the buyer is of the l-type. While our seller could draw more creative conclusions in such an event, accepting lower offers is strictly dominated for the h-type, unless he aspires to get offers strictly above $\mathcal{U}^h$ by doing so, which is rather wishful thinking.[18]

What pins the cutoff belief ℓ^*? Let $V(\ell;\ell^*)$ denote the seller's payoff, given belief ℓ, under the cutoff ℓ^*. Typically, considering that the buyer adjusts his behavior to the cutoff, one cannot expect the seller's cutoff to be time-consistent.[19] That is, depending on ℓ, the seller might prefer different values for ℓ^*. Provided her favorite cutoff is interior, this favorite ℓ^* must solve the first-order condition $V_2(\ell;\ell^*)$, where $V_2(\cdot;\cdot)$ denotes the derivative with respect to the second argument of V. In particular, if she is willing to switch the utility level she offers at $\ell = \ell^*$, this belief must solve

$$V_2(\ell;\ell^*)|_{\ell=\ell^*} = 0. \tag{7}$$

Fixing ℓ^*, the payoff function V can be solved explicitly, but as a stepping stone it is easier to work with the rescaled payoff $f(\ell;\ell^*) = (1+e^\ell)V(\ell;\ell^*)$ (so, dividing the payoff by the probability $1-\mu$ that the buyer is of the l-type): on the range $\ell \geq \ell^*$, it is the unique solution of the delay differential equation, with $j := \ln(\lambda^h/\lambda^l)$,

$$rf(\ell) = (\lambda^l + e^\ell\lambda^h)\Pi^h + \lambda^l(f(\ell+j) - f(\ell)) - (\lambda^h - \lambda^l)f'(\ell), \tag{8}$$

(where we drop the argument ℓ^* from f), with boundary condition $\lim_{\ell\to+\infty} f(\ell)/(1+$

[17] That is, unless this value would exceed $\mathcal{U}^h$, in which case we set $\check{\mathcal{U}}(\ell) = \mathcal{U}^h$.

[18] We note that our updating rules (5)–(6) (which are stronger than the usual requirements perfect Bayesian equilibrium imposes) *imply* what we posit here as the seller's inference, and the buyer's expectations –except in the event in which an offer below $\check{\mathcal{U}}(\ell)$ is unexpectedly accepted.

[19] To be more precise, the seller's payoff depends on the prevailing belief, ℓ, the cutoff ℓ^*, and the cutoff the buyer expects the seller to use, call it ℓ^*_e. In equilibrium, $\ell^*_e = \ell^*$, but the seller takes ℓ^*_e as given, so the first-order condition is with respect to ℓ^*, keeping ℓ^*_e, and then solving for the value of ℓ such that $\ell = \ell^* = \ell^*_e$.

$e^\ell) = \lambda^h \Pi^h/r$. This equation has a familiar interpretation: the annuity $(rf(\ell))$ is equal to sum of the flow payoff $((\lambda^l + e^\ell \lambda^h)\Pi^h$, where $(\lambda^l + e^\ell \lambda^h)$ is the probability of an arrival and hence an acceptance, times $1 + e^\ell)$, of the capital gain in case an arrival materializes $(\lambda^l(f(\ell+j) - f(\ell)))$, and of the change in value due to the drift in belief $(-(\lambda^h - \lambda^l)f'(\ell))$.[20]

Equation (8) admits as unique solution[21]

$$f(\ell) \quad = \quad \frac{\lambda^l}{r}(1+e^{\ell+j})\Pi^h + \left(f(\ell^*) - \frac{\lambda^l}{r}(1+e^{\ell^*+j})\Pi^h\right)e^{-\kappa(\ell-\ell^*)}, \tag{9}$$

where κ is the unique positive root of

$$r + \lambda^l = \lambda^l\left(\frac{\lambda^h}{\lambda^l}\right)^{-\kappa} + (\lambda^h - \lambda^l)\kappa. \tag{10}$$

As for $\ell < \ell^*$, given the rescaling, it holds that $f(\ell) = \lambda^l \Pi^l/r$. By continuity, this is also the value of the constant $f(\ell^*)$ that appears in (9).

As it turns out, the necessary condition (7) pins down a unique candidate for ℓ^*. Namely,

$$e^{\ell^*} = \frac{\kappa}{1+\kappa}\frac{\lambda^l}{\lambda^h}\frac{\Pi^l - \Pi^h}{\Pi^h}. \tag{11}$$

Both κ and ℓ^* are increasing in r, in λ^l and decreasing in λ^h.

Formally, the seller's strategy is

$$\sigma_S(\ell) = \begin{cases} \mathcal{U}^h & \text{if } \ell \geq \ell^*, \\ \mathcal{U}^l & \text{if } \ell < \ell^*, \end{cases} \tag{12}$$

[20] The only perhaps surprising feature of this equation is the coefficient λ^l in front of the capital gain, as opposed to the rescaled probability of a jump. Recall that $f(\cdot)$ is rescaled, so that, in particular $f(\ell + j)$ is the payoff $V(\ell + j; \ell^*)$, *times* $1 + e^{\ell+j}$. Yet, by Bayes rule, the ratio $(1 + e^{\ell+j})/(1 + e^\ell)$, which is the probability that the buyer is of the l-type after an arrival, relative to that prior probability, is simply λ^l divided by the probability of a jump.

[21] Continuous time is key to obtaining closed forms, and hence to verifying that candidate strategies form an equilibrium. Even in continuous time, delay-differential equations such as (8) rarely admit explicit solutions, and studying them requires considerable care (see Keller and Rady, 2001).

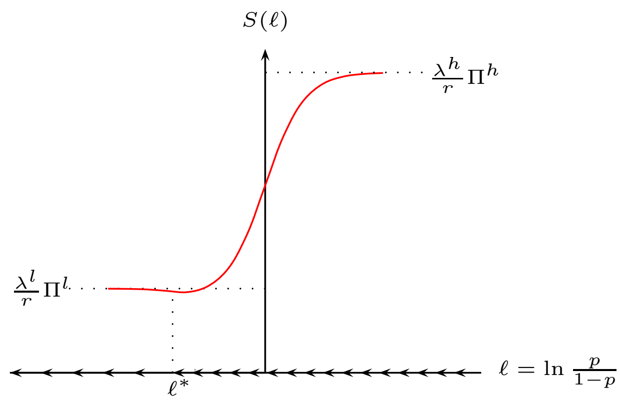

Figure 1: Seller's payoff in the bad news case.[23]

where ℓ^* is given by (11). The buyer's strategy is

$$\sigma_B^h(\ell,\mathcal{U}) = \begin{cases} In & \text{if } \mathcal{U} \geq \mathcal{U}^h, \\ Out & \text{otherwise,} \end{cases} \tag{13}$$

$$\sigma_B^l(\ell,\mathcal{U}) = \begin{cases} In & \text{if } \ell \geq \ell^* \text{ and } \mathcal{U} \geq \check{\mathcal{U}}(\ell), \\ In & \text{if } \ell < \ell^* \text{ and } \mathcal{U} \geq \mathcal{U}^l, \\ Out & \text{otherwise,} \end{cases} \tag{14}$$

where $\check{\mathcal{U}}(\ell)$ is given by, for $\ell \geq \ell^*$,

$$\check{\mathcal{U}}(\ell) = \mathcal{U}^l + (\mathcal{U}^h - \mathcal{U}^l) \cdot \min\left\{1, \frac{\lambda^l}{r}(1-\alpha)\big(1-e^{-\kappa(\ell-\ell^*)}\big)\right\}. \tag{15}$$

It is readily verified that the value function $V(\cdot)$ (that is, the seller's payoff function given tis nothe optimal ℓ^*) is strictly quasiconvex and decreasing at ℓ^*.[22] This is not too surprising: for $\ell < \ell^*$, the seller bets on the buyer being of the l-type, giving up on the buyer in the event he is of the h-type. Such a strategy is particularly costly when the buyer is likely to be of the h-type. By continuity, the same calculus applies for ℓ close to, but above ℓ^*. As for high enough ℓ, the seller's payoff increases with ℓ simply because she anticipates more frequent trade. Figure 1 illustrates.

We summarize this discussion with the following theorem.

[22] It is differentiable everywhere, including at ℓ^*.

Theorem 2 *The strategies given by* (12), (13) *and* (19), *and the belief updating rules given by* (5)–(6) *constitute an equilibrium, provided* $\lambda^l < \alpha\lambda^h$.

Sufficiency is established in appendix.

We note that asymptotic learning is complete, but it does come at a cost to the seller, in the sense that she might lose the heavy customer with positive probability. If $\ell \leq \ell^*$, she loses him with probability 1. If $\ell > \ell^*$, she loses him with probability

$$\mathbf{P}^\sigma[\ell^* \text{ is eventually hit} \mid k = h] = e^{-(\ell-\ell^*)}. \tag{16}$$

Indeed, it is readily verified that this expression solves the functional equation

$$\lambda^h y(\ell) = \lambda^h y(\ell + j) - (\lambda^h - \lambda^l) y'(\ell),$$

subject to $y(\ell^*) = 1$ and $\lim_{\ell \to +\infty} y(\ell) = 0$, which is the defining property of this probability. Perhaps surprisingly, the arrival intensities λ^k do not play a role in these formulas. But, of course, they affect the probabilities via ℓ^*. We note that, as one might have surmised, the probability that the heavy customer is lost is decreasing in the seller's patience (since ℓ^* is increasing in r). The probability that an outside option is suitable plays no role whatsoever, except in determining $\mathcal{U}^h$,$\mathcal{U}^l$.

There is another simpler equilibrium, which exists for some parameters, and can be viewed as the special case in which $\ell^* \to -\infty$: the seller offers $\mathcal{U}^h$ for any belief $\ell > -\infty$. The h-type buyer accepts any offer over $\mathcal{U}^h$. What would entice the l-type to reveal his type? The seller would have to offer at least $\hat{\mathcal{U}}$, which solves

$$\hat{\mathcal{U}} + \frac{\lambda^l}{r}\mathcal{U}^l = \alpha\big(\mathcal{U}^o + Z^{o,l}\big) + (1-\alpha)\frac{\lambda^l}{r}\mathcal{U}^h,$$

or, rearranging, and using the definition of $\mathcal{U}^l$,

$$\hat{\mathcal{U}} = \mathcal{U}^l + (1-\alpha)\frac{\lambda^l}{r}(\mathcal{U}^h - \mathcal{U}^l). \tag{17}$$

That is, $\hat{\mathcal{U}}$ is the value that makes the l-type buyer indifferent between trading now, and revealing his type, or forfeiting the trade, and getting $\mathcal{U}^h$ forever after, in case of a bad match. Note that if $r < (1-\alpha)\lambda^l$, the seller would have to offer more than $\mathcal{U}^h$ to induce the l-type buyer to reveal his type. But such an offer would also be accepted by the h-type. Hence, $r < (1-\alpha)\lambda^l$ is a sufficient condition for such an equilibrium to exist, as the seller is then unable to screen the buyer's types. It is also necessary, as otherwise, for low enough ℓ, the seller finds it profitable to offer $\hat{\mathcal{U}}$ as a way to screen the buyer's types, and save on the utility she must supply.

To summarize, the seller's strategy is[24]

$$\sigma_S(\ell) = \mathcal{U}^h \text{ for all } \ell > -\infty \;\; \sigma_S(-\infty) = \mathcal{U}^l, \tag{18}$$

and the buyer's strategy is, for $k = h, l$, and all $\ell > -\infty$,

$$\begin{aligned} \sigma_B^k(\ell,\mathcal{U}) &= \begin{cases} In & \text{if } \mathcal{U} \geq \mathcal{U}^h, \\ Out & \text{otherwise,} \end{cases} \\ \sigma_B^k(-\infty,\mathcal{U}) &= In \text{ if, and only if, } \mathcal{U} \geq \mathcal{U}^k. \end{aligned} \tag{19}$$

We then have

Theorem 3 *The strategies given by* (18) *and* (19)*, and the belief updating rules given by* (5)–(6) *constitute an equilibrium, provided* $\lambda^l < \alpha\lambda^h$ *and* $r < (1-\alpha)\lambda^l$.

Needless to say, the seller prefers the cutoff equilibrium described first, and the buyer prefers the simpler equilibrium without screening (although the h-type is indifferent).

4.3 No News is Good News

We now turn to the more complex case in which $\lambda^l > \alpha\lambda^h$, so that the seller's belief drifts up in case only the l-type is supposed to accept. Hence, if the seller uses a cutoff strategy as in the bad news case, offering $\mathcal{U}^h$ above it, and targeting the l-type only below it, belief dynamics in the absence of trade will converge to the cutoff ℓ^* from either side of it.

If so, ℓ^* must be a rest point of these dynamics. At the cutoff, the seller must offer a lottery over the utilities that she supplies on either side, to ensure that the belief does not budge, until one of these utilities is accepted: depending on the accepted realized utility, the belief ℓ then either jumps up by j (as both types type accept $\mathcal{U}^h$, but the

[23]The parameters in this example are $(r, \alpha, \lambda^l, \lambda^h, R, c) = (1/6, 2/3, 3/5, 1, 3, 1)$.

[24]Of course, $\sigma_S(-\infty) = \mathcal{U}^l$, $\sigma_B^k(-\infty,\mathcal{U})$ if, and only if, $\mathcal{U} \geq \mathcal{U}^k$.

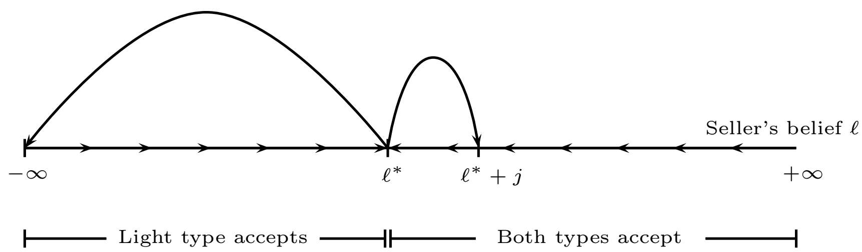

Figure 2: Belief dynamics at the cutoff ℓ^*.

h-type is more likely to get such an opportunity), or down to $-\infty$ (as only the l-type accepts the lower offer). This is why we define the seller's actions as lotteries over utilities, as opposed to a single utility level. Note that this would not occur under the same dynamics in discrete time: the belief would jump back and forth above and below ℓ^*, with the seller incessantly switching between the lower and higher offer, with frequencies that converge to the weights attached to the lottery's utility as the period length vanishes.[25] Yet, despite being an artifact of continuous time, this lottery admits a simple interpretation: the seller uses different channels to sell her product (say, direct to consumer vs. wholesale), but she only finds out which channel the buyer has sampled in case this buyer makes a purchase. Indeed, there is a large literature in marketing that emphasizes how pricing and consumer targeting should be adjusted to the channel. See Figure 2 for an illustration of the belief dynamics at ℓ^*.

These dynamics have two consequences that markedly distinguish such an equilibrium from the cutoff equilibrium in the bad news case. First, when $\ell < \ell^*$, the seller does not

[25]Formally, consider the discrete-time game in which the period length is $\Delta > 0$ (see Section 4.4 for details), and generic parameters such that the belief jump when both types are supposed to accept (but fail to do so) and the belief jump when only the light type is supposed to accept (but does not) are not multiples of each other. Fixing the strategies given in Theorem 4 (which we do not claim are equilibrium strategies in discrete time), then, absent trade, the discrete-time dynamics of the beliefs satisfy

$$\lim_{\varepsilon\downarrow 0}\lim_{\Delta\to 0}\lim_{n\to\infty}\frac{\sum_n \#\mathbf{1}_{\{\ell_n\in(\ell^*,\ell^*+\varepsilon)\}}}{\sum_n \#\mathbf{1}_{\{\ell_n\in(\ell^*-\varepsilon,\ell^*)\}}}=\frac{\lambda^l-\alpha\lambda^h}{(1-\alpha)\lambda^h},$$

cf. (20). The lottery is the familiar chattering control from optimal control, which allows to sidestep the issue of an infinite number of switches on any finite-time interval that arises otherwise.

give up on the h-type altogether: this type is dormant, sampling the outside option until the belief is back to ℓ^*. Second, when $\ell < \ell^*$, enticing the l-type to disclose his identity requires a utility strictly above $\mathcal{U}^l$: by stubbornly sampling the outside option until the belief is back to ℓ^*, the l-type can anticipate getting eventually $\mathcal{U}^h$. The utility the seller supplies must be sufficiently high to deter him from doing so.

Formally, the seller's strategy is

$$\sigma_S(\ell) = \begin{cases} \mathcal{U}^h & \text{if } \ell \geq \ell^*, \\ \mathcal{U}^h & \text{wp. } \frac{\lambda^l - \alpha\lambda^h}{(1-\alpha)\lambda^h}, \hat{\mathcal{U}}(\ell^*) \text{ wp. } \frac{\lambda^h - \lambda^l}{(1-\alpha)\lambda^h}, \text{if } \ell = \ell^*, \\ \hat{\mathcal{U}}(\ell) & \text{if } \ell < \ell^*, \end{cases} \tag{20}$$

where $\hat{\mathcal{U}}(\ell)$ is given by, for $\ell \leq \ell^*$,

$$\hat{\mathcal{U}}(\ell) = \mathcal{U}^l + (\mathcal{U}^h - \mathcal{U}^l)\frac{\lambda^l}{r}\frac{(1-\alpha)(\lambda^l - \alpha\lambda^h)\kappa}{r + \alpha\lambda^l + (\lambda^l - \alpha\lambda^h)\kappa} e^{-\frac{r+\alpha\lambda^l}{\lambda^l - \alpha\lambda^h}(\ell^* - \ell)}, \tag{21}$$

and ℓ^* is given by

$$e^{\ell^*} = \frac{\lambda^l\kappa}{\lambda^h(\alpha\lambda^h + r)(1+\kappa)} \frac{r(\Pi(\hat{\mathcal{U}}(\ell^*)) - \Pi^h) + \lambda^l(\Pi^l - \Pi^h)}{\Pi^h}, \tag{22}$$

This cutoff is pinned down, here as in the bad news case, by the first-order condition (7). Perhaps surprisingly, $\hat{\mathcal{U}}(\ell^*)$ is decreasing in α, fixing $\mathcal{U}^h$,$\mathcal{U}^l$.[26] This is because a higher arrival rate of suitable outside opportunities decreases the probability at which the high utility $\mathcal{U}^h$ must be supplied at $\ell = \ell^*$ (the drift "up" is lower, so the probability on the low offer must be higher). Hence, the buyer gains less from waiting for the high utility, as it is less likely to arise soon. In turn, this means that the seller has to compensate him less via the low offer. As one would expect, $\hat{\mathcal{U}}(\ell^*)$ is decreasing in r: the more patient the buyer, the more she must be offered as a compensation to reveal his type.

This also means that the seller's payoff from the low offer $\hat{\mathcal{U}}(\ell^*)$ is increasing in α; as a result, the net effect of α on ℓ^* is ambiguous: the first term of (22) is decreasing in α, reflecting the heightened risk of losing the heavy buyer when suitable outside opportunities arise often. Similarly, a higher discount rate can either decrease or

[26]Taking into account how $\mathcal{U}^h$,$\mathcal{U}^l$ adjusts, the overall effect can be of either sign.

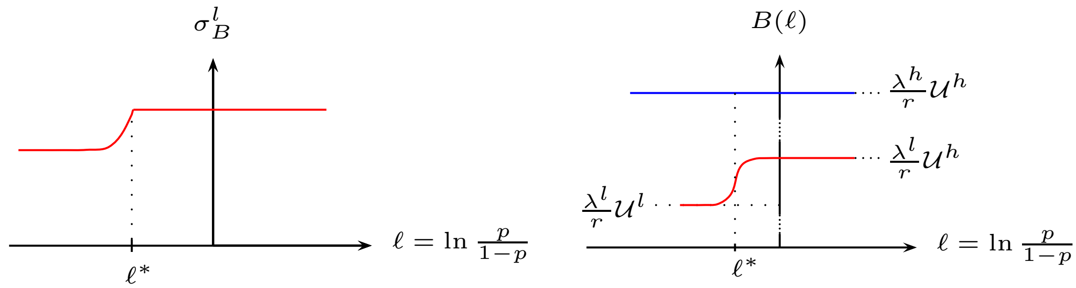

Figure 3: (Light) buyer's strategy (left) and buyer's payoff (right) in the good news case.[27]

increase ℓ^*, depending on parameters.

The buyer's strategy is

$$\sigma_B^h(\ell, \mathcal{U}) = \begin{cases} In & \text{if } \mathcal{U} \geq \mathcal{U}^h, \\ Out & \text{otherwise,} \end{cases} \tag{23}$$

$$\sigma_B^l(\ell, \mathcal{U}) = \begin{cases} In & \text{if } \ell > \ell^* \text{ and } \mathcal{U} \geq \check{\mathcal{U}}(\ell), \\ In & \text{if } \ell \leq \ell^* \text{ and } \mathcal{U} \geq \hat{\mathcal{U}}(\ell), \\ Out & \text{otherwise,} \end{cases} \tag{24}$$

where $\check{\mathcal{U}}(\ell)$ is given by, for $\ell > \ell^*$,

$$\check{\mathcal{U}}(\ell) = \mathcal{U}^l + (\mathcal{U}^h - \mathcal{U}^l) \cdot \min\left\{1, \frac{\lambda^l}{r}(1-\alpha)\left(1 - \frac{(r + \alpha\lambda^l)e^{-\kappa(\ell - \ell^*)}}{r + \alpha\lambda^l + (\lambda^l - \alpha\lambda^h)\kappa}\right)\right\}. \tag{25}$$

The light buyer's strategy and the buyer's payoff are illustrated in Figure 3.

Note that the weights attached to the two utility levels of the lottery offered at ℓ^* are pinned down by the requirement that ℓ^* be a rest point of the belief dynamics. It might well be that the higher utility level $\mathcal{U}^h$ appears with a high weight. This makes it tempting for the l-type seller to pass on such realizations, and the utility level $\hat{\mathcal{U}}(\ell^*)$ reflects this: the higher is this weight, the higher is $\hat{\mathcal{U}}(\ell^*)$. For some parameters, notably, when the l-type is sufficiently patient, this could lead to $\hat{\mathcal{U}}(\ell^*)$ above $\mathcal{U}^h$ (recall that $\hat{\mathcal{U}}(\ell^*)$ must compensate the l-type from revealing his type, which is particularly costly if he is patient). This cannot be, for otherwise the h-type would also accept this offer. This leads to the necessary, and, as we show in appendix, sufficient condition for this equilibrium to exist, namely, $\hat{\mathcal{U}}(\ell) < \mathcal{U}^h$ for all $\ell < \ell^*$, or, in terms of primitives, given (25),

$$r(\alpha\lambda^l + r) \geq (\lambda^l - \alpha\lambda^h)\big((1-\alpha)\lambda^l - r\big)\kappa. \tag{26}$$

This defines a lower bound on r (fixing other parameters), as well as on α.

Perhaps more unexpectedly, there is another necessary condition: for the seller's strategy given by (20) to be optimal, it must be the case that she never prefers to make an offer that both types reject. While $\Pi^h > 0$ ensures that, under complete information, exclusion is never optimal, this is not entirely obvious here, in particular, when $\ell < \ell^*$ (as it turns out, she never wants to make an unacceptable offer for $\ell > \ell^*$, see the proof of Theorem 4). Indeed, by making an unacceptable offer, the seller lets the belief drift down, which means that the offer that she needs to make in order to screen the l-type, $\hat{\mathcal{U}}(\ell)$ goes down as well. Of course, this is a risky course of action, as the buyer might find an attractive outside option in the meantime; nonetheless, it cannot be ruled out without making an assumption on parameters. Specifically, let

$$H(\ell) := \Pi(\hat{\mathcal{U}}(\ell)) + \frac{\lambda^l}{r}\Pi^l$$

denote the seller's continuation payoff when the buyer visits her (recall that $\hat{\mathcal{U}}(\cdot)$ is given by (25)).

Assumption A. It holds that $\ell < \ell^* \Rightarrow (r + \alpha\lambda^l)H(\ell) + \alpha(\lambda^h - \lambda^l)H'(\ell) > 0$.

Assumption **A** reduces to an assumption on the slope of $\Pi(\cdot)$, and is sufficient to ensure that the seller does not want to make an unacceptable offer. Weaker sufficient conditions exist (as well as less compact, but stronger assumptions stated directly in terms of primitives), but *some* condition is needed: there are parameters for which exclusion is optimal for some beliefs, despite $\Pi^h > 0$. This is another difference with the standard model in which private information pertains to valuations: here, the

[27] The parameters in this example are $(r, \alpha, \lambda^l, \lambda^h, R, c) = (1/6, 1/2, 3/4, 1, 3, 1)$. The light buyer's payoff is in red, the heavy buyer's payoff is in blue.

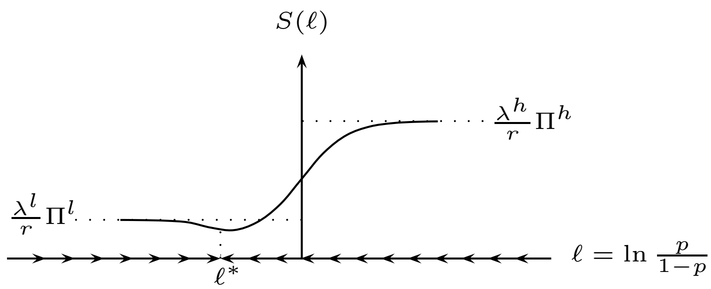

Figure 4: Seller's payoff in the good news case.[28]

seller might benefit from (temporarily) excluding *all* types of buyers. Yet, we will refrain from exploring what happens when the seller wants to do so and maintain Assumption **A** throughout.

Theorem 4 *The strategies given by* (20), (23) *and* (24)*, and the belief updating rules given by* (5)–(6) *constitute an equilibrium, provided* $\lambda^l > \alpha\lambda^h$ *and* (26) *holds, as well as Assumption **A**.*

For completeness, the rescaled seller's value function is given by, for $\ell \geq \ell^*$,

$$f(\ell) = (\lambda^l + \lambda^h e^\ell)\frac{\Pi^h}{r} + \frac{\lambda^l}{r}\frac{r(\Pi(\hat{\mathcal{U}}(\ell^*)) - \Pi^h) - \lambda^l(\Pi^h - \Pi^l)}{(1+\kappa)(\lambda^l + r + (\lambda^l - \alpha\lambda^h)\kappa)}e^{-\kappa(\ell-\ell^*)}, \tag{27}$$

and for $\ell < \ell^*$,

$$\begin{aligned} f(\ell) &= e^{-\frac{r+\lambda^l}{\lambda^l-\alpha\lambda^h}(\ell^*-\ell)}f(\ell^*) \\ &+ \frac{\lambda^l}{\lambda^l - \alpha\lambda^h}\int_\ell^{\ell^*} e^{-\frac{\lambda^l+r}{\lambda^l-\alpha\lambda^h}(\tilde{\ell}-\ell)}\left[\Pi(\hat{\mathcal{U}}(\tilde{\ell})) + \frac{\lambda^l}{r}\Pi^l\right]\mathrm{d}\tilde{\ell}, \end{aligned} \tag{28}$$

where $f(\ell^*)$ is given by (27).

The value function itself ($V(\ell) = f(\ell)/(1+e^\ell)$) is differentiable and strictly quasiconvex, but the minimum can occur on either side of ℓ^*, depending on the parameters. Figure 4 illustrates.

Determining the probability that the heavy customer is eventually lost requires more work than in the bad news case. Let $p(\ell^*)$ denote the probability that the heavy customer remains indefinitely with the seller, starting from belief ℓ^*, and $q(\ell^* + j)$ the probability that the seller's belief eventually hits ℓ^* starting from $\ell^* + j$. Note that $p(\ell^*)$ solves

$$p(\ell^*) = \beta(1 - q(\ell^* + j) + q(\ell^* + j)p(\ell^*)) + (1 - \beta)(1 - \alpha)p(\ell^*), \tag{29}$$

where $\beta = \frac{\lambda^l - \alpha\lambda^h}{(1-\alpha)\lambda^h}$ is the probability that the buyer's utility from the seller is $\mathcal{U}^h$ when he has a need to trade. Note that, from (16), $q(\ell^* + j) = e^{-j} = \frac{\lambda^l}{\lambda^h}$, and so $p(\ell^*) = 1 - \frac{\alpha\lambda^h}{\lambda^l}$. Next, we note that, for $\ell < \ell^*$

$$\mathbf{P}^\sigma[\ell^* \text{ is eventually hit} \mid k = h] = e^{-\frac{\alpha\lambda^h}{\lambda^l - \alpha\lambda^h}(\ell^* - \ell)},$$

so that the probability that the heavy customer is eventually lost starting from $\ell < \ell^*$ is given by

$$1 - e^{-\frac{\alpha\lambda^h}{\lambda^l - \alpha\lambda^h}(\ell^* - \ell)}p(\ell^*) = 1 - \left(1 - \frac{\alpha\lambda^h}{\lambda^l}\right)e^{-\frac{\alpha\lambda^h}{\lambda^l - \alpha\lambda^h}(\ell^* - \ell)}.$$

If instead the initial belief is $\ell \geq \ell^*$, then, recalling (16), this probability is equal to

$$(1 - p(\ell^*))e^{-(\ell - \ell^*)} = \frac{\alpha\lambda^h}{\lambda^l}e^{-(\ell - \ell^*)}.$$

The comparative statics of these probabilities are ambiguous, because the impact on ℓ^* of changes in α, r are ambiguous (see discussion below (22)).

Finally, the equilibrium described in Theorem 3 also exists in the good news case, under the same restriction on parameters. The seller's payoff in these two equilibria is compared in Figure 5. Formally:

Theorem 5 *The strategies given by* (18) *and* (19)*, and the belief updating rules given by* (5)–(6) *constitute an equilibrium, provided* $\lambda^l > \alpha\lambda^h$ *and* $r < (1 - \alpha)\lambda^l$*.*

Note that, if $r \geq (1 - \alpha)\lambda^l$, so that the equilibrium of Theorem 5 does not exist, then (26) is automatically satisfied (the RHS of the inequality is negative), so that the equilibrium of Theorem 4 does exist. Whenever both exist, the seller is better off (and the buyer worse off) in the equilibrium described by Theorem 4.

[28] The parameters in this example are $(r, \alpha, \lambda^l, \lambda^h, R, c) = (1/6, 1/2, 3/4, 1, 3, 1)$.

[29] The parameters in this example are $(r, \alpha, \lambda^l, \lambda^h, R, c) = (1/6, 1/2, 3/4, 1, 3, 1)$. The pure pooling equilibrium is illustrated in red, the equilibrium where the seller uses a cutoff strategy is illustrated in blue.

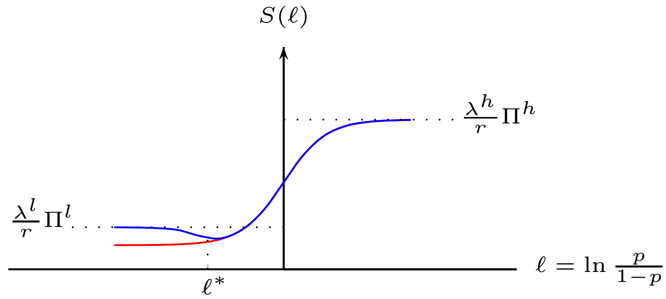

Figure 5: Seller's payoff comparison in the good news case.[29]

4.4 Uniqueness

Markov equilibria are not unique –indeed, as mentioned, the equilibria of Theorem 5 and Theorem 4 exist in parameter regions that overlap. Worse, there are many more Markov equilibria.

To understand the cause of multiplicity, consider the following variation on the equilibrium of Theorem 3. Replace $\mathcal{U}^h$ in (18) and (19) by $(1+\gamma)\mathcal{U}^h$, for some $\gamma > 0$ such that $\Pi((1+\gamma)\mathcal{U}^h) > 0$. That is, the buyer offers strictly more than $\mathcal{U}^h$ at every belief. Suppose that any lower offer, if accepted, leads the seller to the conclusion that the buyer is of the l-type. Then buyers of either type have no incentive to accept such lower offers, provided that r is low enough. To be sure, an offer in the interval $(\mathcal{U}^h, (1+\gamma)\mathcal{U}^h)$ is tempting to the h-type buyer as well, but he would rather forfeit such an opportunity, if it means giving up on all the future opportunities to trade at $(1+\gamma)\mathcal{U}^h$.

A similar, perhaps more perverse example can be constructed even when the seller believes that the buyer is of the l-type. The seller offers $(1+\gamma)\mathcal{U}^h$ ($\gamma > 0$, $\Pi((1+\gamma)\mathcal{U}^h) > 0$), and interprets the acceptance of any lower offer as evidence that the buyer is of the h-type, in which case $\mathcal{U}^h$ is offered forever. Here again, if r is low enough, no buyer's type has an incentive to deviate.

We view these two examples as unconvincing. The seller's deviations to lower offers are deterred by beliefs that "punish" both buyers' types. We could lower γ, and thus the offer the seller makes, and maintain the same on-path inferences without disrupting the equilibrium. The offer is maintained at an arbitrarily high level in a way that we perceive as artificial. We now argue that the equilibria described in Theorems 2–5 are those that survive a refinement, which further selects a unique one among them when two co-exist for a given set of parameters.

Following Hart and Tirole (1988), we focus on the equilibria that arise as limits of equilibria of the finite-horizon game. On the one hand, there are many reasons to be suspicious of such a selection; on the other hand, this refinement is sufficiently stringent that the Markov assumption need no longer be imposed. Markov behavior is a result, not an assumption.

Formally, consider the game in which, up to some time $T = J\Delta$, where $J \in \mathbf{N}$, $\Delta > 0$, the seller gets to choose a utility (or a lottery over utilities) at all times $t = j\Delta$, $j = 1, \ldots, J$. In each interval $((j-1)\Delta, j\Delta]$, $j \in \mathbf{N}$, the buyer with type k gets a need to trade with probability $1 - e^{-\lambda^k \Delta}$. Needs to trade are independent across time intervals. In the event that such a need arises for some $j \leq J$, the buyer gets to choose at time $t = j\Delta$ between the seller's supplied utility $\mathcal{U}_j^{J,\Delta}$ and the outside option, which is modeled as before, namely, a lottery between $\mathcal{U}^o$ and 0, with probability α on $\mathcal{U}^o$. As in the baseline model, the buyer irreversibly leaves the seller if a suitable outside option is found. Note that we do not restrict the seller to trade only in those rounds $j \leq J$, but after time T only the outside option remains available. Let $\sigma_S^{J,\Delta}(\cdot)$ denote the seller's strategy, with J rounds to go, and interval length Δ, as a function of the initial belief; similarly, define $\sigma_B^{l,J,\Delta}(\cdot), \sigma_B^{h,J,\Delta}(\cdot)$, and $\sigma^{J,\Delta}(\cdot) = (\sigma_S^{J,\Delta}(\cdot), \sigma_B^{l,J,\Delta}(\cdot), \sigma_B^{h,J,\Delta}(\cdot))$.

Theorem 6 *Assume $\lambda^l - \alpha\lambda^h < 0$. The game admits an essentially unique equilibrium for all $J \in \mathbf{N}$, $\Delta > 0$.*[30] *Moreover, the equilibrium strategy profile $\sigma^{J,\Delta}(\cdot)$ converges to the Markov equilibrium of Proposition 2, that is, for all $\ell \in \overline{\mathbf{R}}$,*

$$\lim_{\Delta \to 0} \lim_{J \to \infty} \sigma^{J,\Delta}(\ell) = \sigma(\ell),$$

where $\sigma = (\sigma_S(\cdot), \sigma_B^l(\cdot), \sigma_B^h(\cdot))$ is given by (12), (13) *and* (19).

[30] "Essentially" refers to the seller's indifference at the cutoff belief.

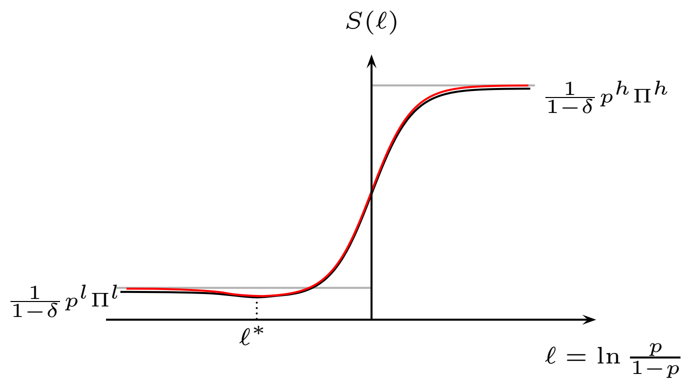

Figure 6: Seller's payoff with 50 periods to go (infinite horizon in red) in the good news case.[31]

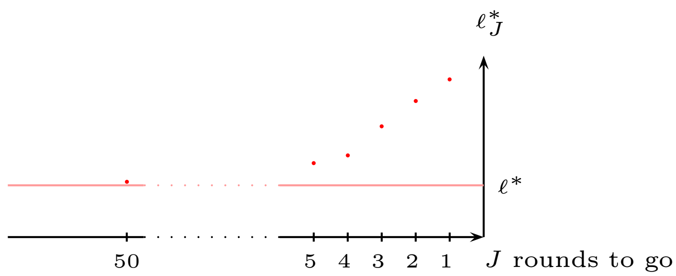

Figure 7: Seller's cutoff ℓ^*_J with J rounds to go.[32]

We prove Theorem 6 in Online Appendix D.

In the good news case, because of the recurrence of the random walk that describes the seller's belief in discrete time, there is little hope to derive a formal result along the lines of Theorem 6.[33] Nonetheless, as we illustrate in Figures 6 and 7, numerical simulations suggest that the equilibrium of the discrete-time game converges (as $J \to \infty$ and $\Delta \to 0$) to the equilibrium strategies as described in Theorem 4 whenever it exists, and to those of Theorem 5 whenever it does not.[34] Figures 6 and 7 illustrate

[31] Parameters are $(r, \alpha, \lambda^l, \lambda^h, R, c, \Delta) = (5.2680, 0.4232, 38.6252, 75.7875, 3, 1, 0.02)$.
[32] Parameters are $(r, \alpha, \lambda^l, \lambda^h, R, c, \Delta) = (5.2680, 0.4232, 38.6252, 75.7875, 3, 1, 0.02)$.
[33] Recall that this random walk has unequal step size; thus, its distribution cannot be solved in closed form.
[34] This assumes that parameters satisfy Assumption **A**, so that no exclusion ever occurs in the continuous-time game. We have not explored what happens when Assumption **A** is violated, whether in continuous time, or in the finite-horizon, discrete-time game.

convergence of the seller's payoff as well as the seller's cutoff.

4.5 Discussion

As mentioned, the particular microfoundation for the reservation utilities $\mathcal{U}^h$,$\mathcal{U}^l$ highlights that no heterogeneity among buyers has to be assumed, aside from the frequency with which trading needs arise. It reflects the fact that informational asymmetries are one of the main sources of switching costs (see the review article of Klemperer, 1995).

But one could also replace the uncertainty about the suitability of the alternative product with a switching cost $c > 0$: $\mathcal{U}^k$ then solves

$$\mathcal{U}^k + \frac{\lambda^k}{r}\mathcal{U}^k = \mathcal{U}^o + \frac{\lambda^k}{r}\mathcal{U}^o - c,$$

which also yields $\mathcal{U}^h > \mathcal{U}^l$.

These reservation utilities could also be the net present value delivered by non-stationary, history-dependent strategies from competing sellers –in particular, these could be equilibrium objects in a larger game between sellers. It is unclear how much additional insight would be gained from fleshing out such a model: regardless of the reservation utilities that would arise, each seller's problem would be identical to the one considered here, and so would be the structure of the equilibrium strategies.

4.6 A Comparison with Hart and Tirole (1988)

In the Section 4 of their paper, Hart and Tirole consider the problem of a seller without commitment who faces a buyer whose binary type pertains to the value of the object, $v^h > v^l$. There is no uncertainty about the arrival rate, which, for the sake of comparison, we take to be $\lambda > 0$.[35] Suppose that the *statu quo* involves the seller posting a price $p = v^l$ –so, a "pooling" offer that leads to no updating. Under which circumstances would the high type accept a (one-shot) seller's deviation to a slightly higher offer, say, $v^l + \varepsilon$? The buyer's tradeoff would be

$$v^h - (v^l - \varepsilon) + W^h \gtrless \alpha\left(\mathcal{U}^{o,h} + \frac{\lambda}{r}\mathcal{U}^{o,h}\right) + (1-\alpha)\frac{\lambda}{r}(v^h - v^l), \tag{30}$$

[35] In their discrete-time model, they assume that the buyer comes in every round, and there is no outside option, equivalently, $\alpha = 0$.

where W^h is the buyer's payoff from always sampling the outside option –if the buyer accepts the offer $v^l + \varepsilon$, the seller's later offers will keep him down to this payoff. The left-hand side of (30) is the payoff from accepting the offer, whereas the right-hand side is the payoff from once sampling the outside option instead. Explicitly,[36]

$$W^h = \frac{\lambda}{r}\frac{\lambda + r}{\alpha\lambda + r}\alpha\mathcal{U}^{o,h}. \tag{31}$$

Here and in (30), we take the outside option $\mathcal{U}^{o,h}$ to depend on h, as there is no reason to assume that it doesn't depend on the buyer's valuation. Indeed, let us take $\mathcal{U}^{o,h} = v^h - v^l$, to capture the idea that the suitable alternatives are sellers posting the same price. Taking $\varepsilon \downarrow 0$ in (30), such an offer by the seller is unattractive if

$$r \leq (1 - \alpha)\lambda. \tag{32}$$

Strikingly, this is the same condition as the one in Theorem (3) and (5), *mutatis mutandis*. The more patient the buyer, or the more likely the outside option turns out to be suitable, the less likely it is that the seller can screen the types. Hart and Tirole (1988) is the special case $\alpha = 0$.[37] Indeed, they assume that (32) holds strictly, and prove that the seller posting v^l in every round is the unique equilibrium outcome that is the limit of the equilibrium outcome in the finite horizon game, as the horizon grows large.[38]

If (32) does not hold, and *fixing* $U^{o,h}$ to $v^h - v^l$, then, provided that the seller places a high enough prior probability on the buyer being of the high type, there exists an equilibrium in which the seller screens the types with an initial offer above v^l: indeed, she then sets a price that makes (30) hold with equality. No further learning occurs. However, it might make more sense to assume that the outside option consists of sellers posting the same price p to new customers, followed by either v^l if that offer is rejected, or whatever price yields continuation payoff W^h to the high-type buyer –so that other sellers follow exactly the same strategy as the seller under consideration.

[36] This is simply the net present value of (1), with $\lambda^h = \lambda$.

[37] In discrete time, $r \leq \lambda$ is equivalent to $\delta \geq 1/2$, which is the assumption of their Proposition 5.

[38] Taking the horizon length to infinity is a reasonable refinement, but so are many others, and it is an approach that is not particularly tractable in continuous time. See Section 4.4.

In that case, if (32) is violated, the "Diamond paradox" kicks in, and the only solution involves $p = v^h$.[39]

Therefore, while the condition that delineates when pooling is an equilibrium or not is the same whether the private information pertains to the buyer's arrival rate or the buyer's valuation, the dynamics are quite different, since no arrival is not informative *per se* in the latter case: all the "action" takes place in the first round.

Under commitment (see Proposition 1 of Hart and Tirole), the solution to the seller's problem is simple: depending on the prior, she sets the price equal to either v^l, so that the buyer purchases, independent of his type, or to v^h, excluding the low type from trade. It simply is not possible to have both types trade, yet extract their surplus nonetheless.

5 Concluding Comments

There are several limitations to our analysis. First, competition is modeled as an exogenous outside option. In particular, there is an obvious tension here between the seller's pricing strategy, which varies over time, and the way this outside option is assumed to be constant, once a suitable match is found. In applications in which there are only a few sellers, it might be fruitful to explicitly model the game these sellers play. Of course, this raises the standard difficulties with repeated games: whether other sellers' prices are observed or not, "folk-theorem" constructions can be devised.[40] The Markov restriction would not help: in the unobserved case, higher-order beliefs become payoff-relevant; in the observed case, it loses its bite, since histories can be encoded in the price vector, which is payoff-relevant. These difficulties do not make such an analysis less relevant. As mentioned, however, our entire analysis goes through while treating $\mathcal{U}^h, \mathcal{U}^l$ as exogenous parameters. Therefore, whatever (possibly history-dependent) strategies other sellers might use does not affect the conclusions of our

[39] Eqn. (30) becomes $v^h - p + W^h \gtrless \alpha\big(v^h - p + W^h\big) + (1-\alpha)\frac{\lambda}{r}(v^h - v^l)$, with $W^h = \frac{\lambda}{r}\alpha(v^h - p)$. Simplification yields (32).

[40] See Bergemann and Välimäki (2002) for an exception (when attention is restricted to pure strategies) to this rule of thumb.

analysis, provided that the environment is stationary, and that a customer who leaves for a (not necessarily uncertain) outside option never returns. These two assumptions seem plausible in environments with many competing sellers.

Second, we do not allow the buyer to hide past purchases. We note, however, that the optimal mechanism under commitment is robust to such an option, as the buyer can only gain from revealing past trades. Under non-commitment, matters are more subtle: on the one hand, the seller can no longer fleece a buyer who is revealed to be light, decreasing her incentives to screen; on the other hand, screening becomes easier, because a light buyer is no longer bothered by the information he might reveal. Non-pooling equilibria can be devised, under some parameter restrictions.

The optimal mechanism under commitment is also robust to other modeling variations. Replacing the uncertainty about the match quality with a switching cost, for instance, as a friction the buyer faces, does not affect its structure. Under non-commitment, a search cost would yield similar results to the "no news is bad news case," since beliefs would drift down independently of the seller's strategy.[41]

An interesting direction of future research involves the dynamics in case fundamentals change. Commitments have little value when crises hit. Evidently, airlines have no qualms about changing their loyalty programs if a pandemic abruptly affects business. If, say, the seller's costs go up, and so her profit function shifts down, how should the seller adjust her policy? Shall she sacrifice the heavy customer, decreasing quality supplied to all of them, or focus on the light ones?

Appendices

A Proof of Theorem 1

Fix $T > 0$. We pick $\mathcal{U}^I$,$\mathcal{U}^F$ such that $Z_0^{\lambda^l} = \frac{\lambda^l}{r}\mathcal{U}^l$, $Z_0^{\lambda^h} = \frac{\lambda^h}{r}\mathcal{U}^h$, that is, both types are held to their outside option; in particular, lying does not benefit either type at the

[41] A buyer who is still attentive is a buyer who had no other offer, which is statistical evidence that he is a light buyer.

reporting stage. Note that if $\mathcal{U}^F \geq \mathcal{U}^I$, either type strictly prefers to visit the seller if he gets the opportunity to do so during the window of length T that opens after a visit. Hence, conditional on reporting heavy, either type of buyer visits the seller whenever the opportunity arises. To confirm $\mathcal{U}^F \geq \mathcal{U}^I$, simple algebra gives

$$\mathcal{U}^F - \mathcal{U}^I = \frac{r(1-\alpha)\alpha(\lambda^h-\lambda^l)e^{-rT}((2\lambda^h+r)e^{(r+\lambda^h)T}-\lambda^h)((2\lambda^l+r)e^{(r+\lambda^l)T}-\lambda^l)\mathcal{U}^o}{(\alpha\lambda^l+r)(\alpha\lambda^h+r)\big(r(\lambda^h-\lambda^l)e^{(r+\lambda^h+\lambda^l)T}-\lambda^h(\lambda^l+r)e^{\lambda^l T}+\lambda^l(\lambda^h+r)e^{\lambda^h T}\big)},$$

which is of the sign of the denominator, that is, of

$$g(r) := r(\lambda^h-\lambda^l)e^{(r+\lambda^h+\lambda^l)T}-\lambda^h(\lambda^l+r)e^{\lambda^l T}+\lambda^l(\lambda^h+r)e^{\lambda^h T},$$

which is positive. To see this, note that (i)

$$g(0) = \lambda^h\lambda^l\Big(e^{\lambda^h T}-e^{\lambda^l T}\Big) \geq 0,$$

(ii) g is convex since

$$g''(r) = T(2+rT)(\lambda^h-\lambda^l)e^{(\lambda^h+\lambda^l+r)T}$$

and (iii)

$$g'(0) = \frac{e^{(\lambda^h+\lambda^l)T}}{T}\Big(\lambda^h T\Big(1-e^{-\lambda^h T}\Big)-\lambda^l T\Big(1-e^{-\lambda^l T}\Big)\Big) \geq 0,$$

(the function $x \mapsto x(1-e^{-x})$ being increasing).

Algebra also gives

$$\lim_{T\downarrow 0}\mathcal{U}^I = \frac{\alpha\big(\alpha\lambda^h\lambda^l+r\alpha(\lambda^h+\lambda^l)+r^2\big)}{(\alpha\lambda^h+r)(\alpha\lambda^l+r)}\mathcal{U}^o \geq 0,$$

and so $\mathcal{U}^I$,$\mathcal{U}^F$ are both positive for small enough T.[42]

B Proofs of Theorems 2–5

All four theorems require verification that the strategy profiles referred to in the statements, alongside the updating rules given by (5) and (6), constitute equilibria of the game. We note that, fixing the strategy profiles, and the updating rules, we can view the game as a stochastic game with the seller's belief as a state variable, as

[42]Whether or not $\mathcal{U}^I$ is positive for T arbitrarily large depends on parameters. The difference $\mathcal{U}^F - \mathcal{U}^I$ is actually decreasing in T: a seller whose cost is convex prefers a larger T among this class of mechanisms.

explained at the end of Section 2, and so we can apply the one-shot deviation principle to verify optimality. Recall that we assume that, given any deviation, play proceeds as it would on path, given the updated belief (the strategies have been specified for *any* belief $\ell \in \overline{\mathbf{R}}$).

Further, we note that the buyer's best-reply must be a cutoff rule, given the specified strategy profile: an acceptance of a higher offer leads to a weakly higher posterior belief, and the buyer's payoff is non-decreasing in the belief (as we verify below, see *e.g.*, (33)). Hence, if a buyer of type $k = h, l$ accepts offer $\mathcal{U}$ with positive probability, given belief ℓ, it is also optimal for him to accept $\mathcal{U}' > \mathcal{U}$. Hence, it suffices to verify indifference at the specified cutoffs ((15), (17), (21) and (25)).

Further, we note that the buyer's h-type cutoff must be $\mathcal{U}^h$: indeed, he can at any time secure the equivalent continuation payoff from sampling the outside option, and the seller never supplies any strictly higher utility given the specified profile, so there is no potential "upside" from turning down an off-path offer above $\mathcal{U}^h$. Hence, the buyer's h-type problem is trivial, and we ignore it from now on.

Hence, accepting any offer $\mathcal{U} \geq \mathcal{U}^h$ leads to an increase in the seller's belief (whereas the belief evolves continuously in case of a rejection). Therefore, because in the specified equilibria the l-type's payoff is non-decreasing in the belief, the l-type must also accept any such offer; hence, his cutoff is no larger than $\mathcal{U}^h$. Recall that, as stated in Section 2, we assume throughout that the acceptance of an offer both types should reject leads to an updated belief $\ell = -\infty$. Hence, in all four equilibria, accepting an offer $\mathcal{U} < \mathcal{U}^h$ leads to belief $\ell = -\infty$.

In each case, we first derive the l-type's payoff as a function of the belief, and verify that he is indifferent at the specified cutoffs. For clarity, in the derivations, we assume that this payoff is differentiable in ℓ, but this is directly verified, since these derivations yield differentiable closed-form solutions. We then turn to the seller, and show that, in each relevant interval of belief, the one-shot deviation principle holds. Note that only three offers are candidate optima: $\mathcal{U}^h$ (which both types accept, according to the specified strategies), the l-type's cutoff (denoted $\hat{\mathcal{U}}(\ell)$ or $\check{\mathcal{U}}(\ell)$), which is the lowest offer the l-type accept, or any lower offer, that both types reject. We first ignore this last possibility, but verify *ex post* that it is unprofitable.

Throughout, let us write $Z^l(\ell)$ for the l-type's *ex ante* equilibrium payoff, given ℓ.

B.1 Proof of Theorem 2

Let δ_l denote the probability with which the seller makes a screening offer (assuming such an offer exists, *i.e.*, it is less than $\mathcal{U}^h$); let δ_b denote the probability with which the seller makes the pooling offer ($\mathcal{U}^h$), and $\delta_\varnothing$ denote the probability with which she makes an unacceptable offer. On the seller's side, taking the buyer's behavior as given by the equilibrium strategies given by (13) and (19), we have to verify that the value function $f(\cdot)$ given by (9) for $\ell \geq \ell^*$ and $f(\ell) = \lambda^l \Pi^l / r$ for $\ell < \ell^*$ satisfies the following HJB equation:

$$(r+\lambda^l) f(\ell) = \max_{\delta_\varnothing, \delta_b, \delta_l} \delta_b \big((\lambda^h e^\ell + \lambda^l) \Pi^h + \lambda^l f(\ell + j) \big) + \lambda^l \delta_l \left(\hat{\Pi}(\ell) + \frac{\lambda^l}{r} \Pi^l \right)$$
$$+ \quad \big(\delta_l (\lambda^l - \alpha \lambda^h) - \delta_\varnothing \alpha (\lambda^h - \lambda^l) - \delta_b (\lambda^h - \lambda^l) \big) f'(\ell) + (1-\alpha) \lambda^l \delta_\varnothing f(\ell).$$

The inequalities (35), (36) and (37) directly verify this equation.

Buyer's problem: Given $\ell < \ell^*$, $Z^l(\ell) = \mathcal{U}^l$, and it immediately follows that the l-type is indifferent between the outside option and the seller's offer if, and only if, $\mathcal{U} = \mathcal{U}^l$.

Given $\ell \geq \ell^*$, l-type's payoff satisfies

$$(r+\lambda^l) Z^l(\ell) = \lambda^l \big(\mathcal{U}^h + Z^l(\ell + j) \big) - (\lambda^h - \lambda^l)(Z^l)'(\ell),$$

where we recall that $j = \ln \frac{\lambda^h}{\lambda^l}$. Given the boundary conditions $\lim_{\ell \to \infty} Z^l(\ell) = \frac{\lambda^l}{r} \mathcal{U}^h$ and $Z(\ell^*) = \frac{\lambda^l}{r} \mathcal{U}^l$, the unique solution to this delay-differential equation is

$$Z^l(\ell) = \frac{\lambda^l}{r} \mathcal{U}^h - \frac{\lambda^l}{r} \big(\mathcal{U}^h - \mathcal{U}^l \big) e^{-\kappa(\ell - \ell^*)}, \tag{33}$$

where κ is given by (10). This is increasing in ℓ, which, as discussed above, implies that he accepts any offer $\mathcal{U} \geq \mathcal{U}^h$. The l-type's choice is to accept or reject an offer

$\mathcal{U} < \mathcal{U}^h$ depending on

$$\mathcal{U} + \frac{\lambda^l}{r}\mathcal{U}^l \gtrless \alpha\big(\mathcal{U}^o + Z^{o,l}\big) + (1-\alpha)Z^l(\ell), \tag{34}$$

an equation that holds with equality if, and only if, $\mathcal{U} = \check{\mathcal{U}}(\ell)$, where $\check{\mathcal{U}}(\ell)$ is given by (15).

Seller's problem: Given $\ell < \ell^*$, the seller's (re-scaled) payoff function f is given by $\frac{\lambda^l}{r}\Pi^l$. Given $\ell \geq \ell^*$, this function solves the delay-differential equation given in the text, see (8), whose solution, given the boundary conditions, is given by (9).

For $\ell < \ell^*$ (resp., $\ell > \ell^*$), we must check that deviating to $\mathcal{U}^h$ (resp., $\check{\mathcal{U}}(\ell)$) is suboptimal. For $\ell < \ell^*$, this means verifying that

$$rf(\ell) + (\lambda^h - \lambda^l)f'(\ell) \geq (\lambda^l + \lambda^h e^\ell)\Pi^h + \lambda^l(f(\ell+j) - f(\ell)). \tag{35}$$

Since $f(\ell)$ is constant, while $f(\ell + j)$ is non-decreasing in ℓ, it suffices to establish this inequality for $\ell = \ell^*$.

For $\ell > \ell^*$, we must show that

$$\lambda^l\left(\check{\Pi}(\ell) + \frac{\lambda^l}{r}\Pi^l - f(\ell)\right) - (\alpha\lambda^h - \lambda^l)f'(\ell) \leq rf(\ell), \tag{36}$$

where $\check{\Pi}(\ell) := \Pi(\check{\mathcal{U}}(\ell))$. Since $\check{\mathcal{U}}(\ell)$ is non-decreasing in ℓ, $\check{\Pi}(\ell)$ is non-increasing in ℓ, whereas it is readily verified that $(r + \lambda^l)f(\ell) - (\lambda^l - \alpha\lambda^h)f'(\ell)$ is increasing in ℓ (see the good news case, where the calculation is carried out), given (9). Hence, here as well, it suffices to check this equation for $\ell = \ell^*$.

Yet, $\ell = \ell^*$ is precisely the value of ℓ such that both (35) and (36) hold with equality, concluding the proof.

It remains to show that, as mentioned above, deviating to an unacceptable offer (that is, one that neither type accepts) is unattractive. This requires showing that, for all ℓ,

$$g(\ell) := (r + \alpha\lambda^l)f(\ell) + \alpha\big(\lambda^h - \lambda^l\big)f'(\ell) \geq 0. \tag{37}$$

This is trivial for $\ell < \ell^*$ (since $f(\ell) = 0$), and is immediate upon computation (using (9)) for $\ell > \ell^*$.

B.2 Proof of Theorem 4

The proof follows exactly the same steps as the proof of Theorem 2, but involves more tedious calculations. The same remarks as at the start of Section B.1 apply. Recall that Theorem 4 assumes that (26) holds.

Buyer's problem: Given $\ell > \ell^*$, l-type's payoff satisfies

$$(r+\lambda^l)Z^l(\ell) = \lambda^l\big(\mathcal{U}^h + Z^l(\ell+j)\big) - \big(\lambda^h - \lambda^l\big)\big(Z^l\big)'(\ell), \tag{38}$$

where we recall that $j = \ln\frac{\lambda^h}{\lambda^l}$. At ℓ^*, because the l-type accepts either offer

$$Z^l(\ell^*) = \frac{\lambda^l}{\lambda^l + r}\left(\delta\big(\mathcal{U}^h + Z^l(\ell^*+j)\big) + (1-\delta)\left(\hat{\mathcal{U}}(\ell^*) + \frac{\lambda^l}{r}\mathcal{U}^l\right)\right), \tag{39}$$

where the buyer's utility $\hat{\mathcal{U}}(\ell^*)$ solves

$$\hat{\mathcal{U}}(\ell^*) + \frac{\lambda^l}{r}\mathcal{U}^l = \alpha\big(\mathcal{U}^o + Z^{o,l}\big) + (1-\alpha)Z^l(\ell^*), \tag{40}$$

and $\delta = \frac{\lambda^h - \lambda^l}{(1-\alpha)\lambda^h}$.

Plugging (40) into (39), we get $Z^l(\ell^*)$ as a function of $Z^l(\ell^*+j)$, which is pinned down by the solution to the delay-differential equation. Given the boundary condition $\lim_{\ell\to\infty} Z^l(\ell) = \frac{\lambda^l}{r}\mathcal{U}^h$, the unique solution to this delay-differential equation is

$$Z^l(\ell) - \frac{\lambda^l}{r}\mathcal{U}^h = \left(Z^l(\ell^*) - \frac{\lambda^l}{r}\mathcal{U}^h\right)e^{-\kappa(\ell-\ell^*)}, \tag{41}$$

where κ is the unique root of

$$r + \lambda^l = \lambda^l\left(\frac{\lambda^h}{\lambda^l}\right)^{-\kappa} + \big(\lambda^h - \lambda^l\big)\kappa.$$

This implies in particular that $Z^l(\ell^*+j)$ is given by:

$$Z^l(\ell^*+j) - \frac{\lambda^l}{r}\mathcal{U}^h = \left(Z^l(\ell^*) - \frac{\lambda^l}{r}\mathcal{U}^h\right)\left(\frac{\lambda^h}{\lambda^l}\right)^{-\kappa}. \tag{42}$$

Plugging (42) into (39) gives $Z^l(\ell^*)$ in terms of primitives:

$$Z^l(\ell^*) = \frac{\lambda^l\big((\alpha\lambda^l + r)\mathcal{U}^l + \big(\lambda^l - \alpha\lambda^h\big)\mathcal{U}^h\kappa\big)}{r(r + \lambda^l\kappa + \alpha(\lambda^l - \lambda^h\kappa))}. \tag{43}$$

Given ℓ, the l-type's choice is to accept or reject an offer $\mathcal{U} < \mathcal{U}^h$ depending on

$$\mathcal{U} + \frac{\lambda^l}{r}\mathcal{U}^l \gtrless \alpha\big(\mathcal{U}^o + Z^{o,l}\big) + (1-\alpha)Z^l(\ell), \tag{44}$$

an equation that holds with equality if, and only if, $\mathcal{U} = \check{\mathcal{U}}(\ell)$, where $\check{\mathcal{U}}(\ell)$ is given by (21).

Given $\ell < \ell^*$, the l-type is indifferent between visiting the seller offering $\mathcal{U}(\ell)$ and the outside option if, and only if,

$$\begin{aligned}\mathcal{U}(\ell) + \frac{\lambda^l}{r}\mathcal{U}^l =& \alpha\big(\mathcal{U}^o + Z^{o,l}\big)\\ &+ (1-\alpha)\bigg(\int_0^{\tau_\ell} \alpha\lambda^l e^{-(r+\alpha\lambda^l)t}\big(\mathcal{U}^o + Z^{o,l}\big)dt + e^{-(r+\alpha\lambda^l)\tau_\ell} Z^l(\ell^*)\bigg),\end{aligned}$$

where $\tau_\ell := \frac{\ell^*-\ell}{\lambda^l-\alpha\lambda^h}$ is the time it takes for the belief to reach ℓ^* when starting at ℓ. This can be simplified to

$$\mathcal{U}(\ell) = \mathcal{U}^l + (1-\alpha)e^{-\frac{r+\alpha\lambda^l}{\lambda^l-\alpha\lambda^h}(\ell^*-\ell)}\bigg(Z^l(\ell^*) - \frac{\lambda^l}{r}\mathcal{U}^l\bigg). \tag{45}$$

Using (44),we can then write it as

$$\mathcal{U}(\ell) = \mathcal{U}^l + (1-\alpha)e^{\frac{(\alpha\lambda^l+r)(\ell-\ell^*)}{\lambda^l-\alpha\lambda^h}}\frac{\lambda^l(\lambda^l-\alpha\lambda^h)\kappa}{r(r+\lambda^l\kappa+\alpha(\lambda^l-\lambda^h\kappa))}\big(\mathcal{U}^h - \mathcal{U}^l\big). \tag{46}$$

Seller's problem: Given $\ell \geq \ell^*$, the seller's re-scaled payoff function f solves the delay-differential equation given in the text; that is, equation (8), whose solution is given by (27) and (28).

For $\ell < \ell^*$, the seller's payoff function f is given by (24).

For $\ell < \ell^*$, we must check that deviating to $\mathcal{U}^h$ is suboptimal, that is, that it holds that

$$rf(\ell) + (\lambda^h - \lambda^l)f'(\ell) \geq (\lambda^l + \lambda^h e^\ell)\Pi^h + \lambda^l(f(\ell+j) - f(\ell)). \tag{47}$$

Using the defining differential equation for f', this is equivalent to

$$\begin{aligned}&\bigg(r + \frac{(\lambda^h-\lambda^l)(r+\lambda^l)}{\lambda^l-\alpha\lambda^h}\bigg)f(\ell) + \frac{(\lambda^h-\lambda^l)\lambda^l}{\lambda^l-\alpha\lambda^h}H(\ell)\\ &\geq (\lambda^l + \lambda^h e^\ell)\Pi^h + \lambda^l(f(\ell+j) - f(\ell)),\end{aligned} \tag{48}$$

where $H(\ell) := \Pi(\hat{\mathcal{U}}(\ell)) + \frac{\lambda^l}{r}\Pi^l$ is the seller's continuation payoff when the buyer visits her (a non-increasing function of ℓ, since $\hat{\mathcal{U}}(\cdot)$ is non-decreasing and $\Pi(\cdot)$ is decreasing).

To prove this inequality, we shall introduce an auxiliary payoff function for the seller, namely, the one in which the screening offer $\mathcal{U}(\ell)$ for $\ell < \ell^*$ is equal to $\hat{\mathcal{U}}(\ell^*)$, rather than $\hat{\mathcal{U}}(\ell)$ (and is accepted by the l-type) for all $\ell < \ell^*$. This is the only change: all dynamics, etc., remain unchanged, except for this collected (one-time) reward. Denote by $\tilde{f}$ the corresponding normalized payoff. Evidently, $f(\ell) \geq \tilde{f}(\ell)$, since $\hat{\mathcal{U}}(\ell^*) \geq \hat{\mathcal{U}}(\ell)$. Further, a coupling argument shows that

$$f(\ell) - \tilde{f}(\ell) \geq f(\ell + j) - \tilde{f}(\ell + j).$$

(The gain from the screening offer being $\hat{\mathcal{U}}(\ell)$ rather than $\hat{\mathcal{U}}(\ell^*)$ is greater, the lower the belief, since $\hat{\mathcal{U}}(\cdot)$ is increasing, and the hitting time of ℓ^* is larger.)[43] Hence,

$$f(\ell + j) - f(\ell) \leq \tilde{f}(\ell + j) - \tilde{f}(\ell).$$

Hence, to show (48), it suffices to prove the corresponding inequality with f replaced by $\tilde{f}$. We have, for $\ell < \ell^*$,

$$\tilde{f}(\ell) = \frac{\lambda^l}{r + \lambda^l} H(\ell^*) + \frac{(r + \lambda^l) f(\ell^*) - \lambda^l H(\ell^*)}{r + \lambda^l} e^{\frac{r+\lambda^l}{\lambda^l - \alpha\lambda^h}(\ell - \ell^*)}.$$

Hence, computing the difference, we obtain, bounding $H(\ell)$ below by $H(\ell^*)$,

$$\begin{aligned} & \left(r + \frac{(\lambda^h - \lambda^l)(r + \lambda^l)}{\lambda^l - \alpha\lambda^h}\right)\tilde{f}(\ell) + \frac{(\lambda^h - \lambda^l)\lambda^l}{\lambda^l - \alpha\lambda^h} H(\ell^*) \\ & -(\lambda^l + \lambda^h e^\ell)\Pi^h - \lambda^l(\tilde{f}(\ell + j) - \tilde{f}(\ell)) \qquad (49) \\ = \quad & \frac{r\lambda^l}{r + \lambda^l}\left(H(\ell^*) - \frac{r + \lambda^l}{r}\Pi^h\right) - \lambda^h \Pi^h e^\ell + \gamma e^{\frac{r+\lambda^l}{\lambda^l - \alpha\lambda^h}\ell}, \qquad (50) \end{aligned}$$

for some $\gamma \in \mathbf{R}$ whose expression we omit. Because $\frac{r+\lambda^l}{\lambda^l - \alpha\lambda^h} \geq 1$, the expression in (13) is concave in $x := e^\ell$ if $\gamma < 0$, and hence, the expression in (12) is positive if it is positive for $\ell = -\infty, \ell^*$. If $\gamma > 0$, so that (13) is convex in x, then algebra shows the derivative at (13), evaluated at $\ell = \ell^*$ is negative, so it suffices then to check at ℓ^*. Hence, we are left with checking the condition at $\ell = -\infty, \ell^*$. At $\ell = \ell^*$, note that $f(\ell^*) = \tilde{f}(\ell^*)$, as well as $H(\ell) = H(\ell^*)$ (our bounds are tight), and so (12) holds with equality, by definition of ℓ^*. At $\ell = -\infty$, (13) is positive as well, since

[43] More formally, for $\ell < \ell^*$,

$$f(\ell) - \tilde{f}(\ell) = \mathbf{E}[e^{-r\tau(\ell)}(\Pi(\hat{\mathcal{U}}(\ell_{\tau(\ell)})) - \Pi(\hat{\mathcal{U}}(\ell^*)))],$$

where $\tau(\ell)$ is the first time that the l-type visits the buyer before ℓ^* is reached (if he does), starting from belief ℓ. Note that the inequality is trivially satisfied when $\ell + j \geq \ell^*$.

$H(\ell^*) = \Pi(\hat{\mathcal{U}}(\ell^*)) + \frac{\lambda^l}{r}\Pi^l > \frac{r+\lambda^l}{r}\Pi^h$ is positive.

For $\ell \geq \ell^*$, we must check that deviating to $\check{\mathcal{U}}(\ell)$ is suboptimal, that is, that it holds that

$$\lambda^l\left(\check{\Pi}(\ell) + \frac{\lambda^l}{r}\Pi^l - f(\ell)\right) + \big(\lambda^l - \alpha\lambda^h\big)f'(\ell) \leq rf(\ell),$$

where $\check{\Pi}(\ell) := \Pi(\check{\mathcal{U}}(\ell))$, or equivalently,

$$\lambda^l\left(\check{\Pi}(\ell) + \frac{\lambda^l}{r}\Pi^l\right) \leq (r+\lambda^l)f(\ell) - (\lambda^l - \alpha\lambda^h)f'(\ell). \tag{51}$$

Since $\check{\mathcal{U}}(\ell)$ is non-decreasing in ℓ, $\check{\Pi}(\ell)$ is non-increasing in ℓ, and so is the LHS of the inequality. Because

$$\frac{\mathrm{d}}{\mathrm{d}\ell}\big((r+\lambda^l)f(\ell) - (\lambda^l - \alpha\lambda^h)f'(\ell)\big) = \big(e^{\ell-\ell^*} - e^{-\kappa(\ell-\ell^*)}\big)(r+\alpha\lambda^h)\frac{\lambda^h\Pi^h}{r}e^{\ell^*} \geq 0,$$

(using the definition of ℓ^* to eliminate the constant $\Pi(\hat{\mathcal{U}}(\ell^*))$ that appears in the definition of f), the right-hand side of (51) is increasing in ℓ, all $\ell \geq \ell^*$. Hence, it suffices to check the inequality (51) for $\ell = \ell^*$. Inserting the formula for f and $\ell = \ell^*$ yields an equality (as should be expected, by construction).

It remains to show that deviating to an unacceptable offer is unattractive. This requires showing that, for all ℓ,

$$g(\ell) := (r+\alpha\lambda^l)f(\ell) + \alpha\big(\lambda^h - \lambda^l\big)f'(\ell) \geq 0. \tag{52}$$

For $\ell \geq \ell^*$, we compute

$$r\frac{\mathrm{d}\big(e^{\kappa\ell}g(\ell)\big)}{\mathrm{d}\ell} = \Pi^h e^{\kappa\ell}\big(\lambda^l(r+\alpha\lambda^l)\kappa + \lambda^h(r+\alpha\lambda^h)(1+\kappa)e^{\ell}\big) \geq 0,$$

so it suffices to show that $g(\ell^*) \geq 0$. Explicitly (using the definition of ℓ^* to eliminate $\Pi(\hat{\mathcal{U}}(\ell)))$,

$$\frac{rg(\ell^*)}{\Pi^h} = \lambda^l(r+\alpha\lambda^l) + \frac{(1+\kappa)\lambda^h(r+\alpha\lambda^h)\big(r+\alpha\lambda^l+\kappa(\lambda^l-\alpha\lambda^h)\big)e^{\ell^*}}{\kappa(r+\lambda^l+\kappa(\lambda^l-\alpha\lambda^h))} \geq 0.$$

Consider then $\ell < \ell^*$. Recall that on this interval f solves

$$(r+\lambda^l)f(\ell) = \lambda^l\left(\hat{\Pi}(\ell) + \frac{\lambda^l}{r}\Pi^l\right) + (\lambda^l - \alpha\lambda^h)f'(\ell).$$

Solving for f', we can simplify (52) to

$$\frac{(1-\alpha)(r+\alpha\lambda^h)}{\alpha(\lambda^h-\lambda^l)}f(\ell) \geq \hat{\Pi}(\ell) + \frac{\lambda^l}{r}\Pi^l. \tag{53}$$

Since

$$\lim_{\ell \to -\infty} \frac{f(\ell)}{\frac{\lambda^l}{\lambda^l + r}\Big(\hat{\Pi}(\ell) + \frac{\lambda^l}{r}\Pi^l\Big)} = 1,$$

(53) is also satisfied as $\ell \to -\infty$, since $\frac{(1-\alpha)(r+\alpha\lambda^h)}{\alpha(\lambda^h - \lambda^l)} > \frac{\lambda^l}{\lambda^l + r}$. Now, computing $g'(\ell)$ and solving $g'(\ell) = 0$ for the integral term that appears in the definition of $f(\ell)$ when $\ell < \ell^*$; re-inserting into $g(\ell)$ gives

$$g(\ell)\Big|_{\substack{g'(\ell)=0 \\ \ell < \ell^*}} > 0 \Leftrightarrow (r + \alpha\lambda^l)H(\ell) + \alpha(\lambda^h - \lambda^l)H'(\ell) > 0,$$

where $H(\ell) := \Pi(\hat{\mathcal{U}}(\ell)) + \frac{\lambda^l}{r}\Pi^l$ is the seller's continuation payoff when the buyer visits her. By assumption, however, $\ell < \ell^* \Rightarrow (r + \alpha\lambda^l)H(\ell) + \alpha(\lambda^h - \lambda^l)H'(\ell) > 0$, and so $g(\cdot)$ can only admit positive local minima on $\{\ell < \ell^*\}$.

B.3 Proof of Theorems 3 and 5

Recall that these two theorems assume $r < (1-\alpha)\lambda^l$. Recall also that the seller is supposed to set $\mathcal{U}(\ell) = \mathcal{U}^h$ for all ℓ, which both types accept. Suppose for sake of contradiction that there exists an offer $\mathcal{U} < \mathcal{U}^h$ that the l-type accepts. This requires

$$\mathcal{U} + \frac{\lambda^l}{r}\mathcal{U}^l \geq \alpha\big(\mathcal{U}^0 + Z^{o,l}\big) + (1-\alpha)\frac{\lambda^l}{r}\mathcal{U}^h, \text{ or } \mathcal{U} \geq \hat{\mathcal{U}},$$

where $\hat{\mathcal{U}}$ is given by (17). As explained in Section 2, if $r < (1-\alpha)\lambda^l$, $\hat{\mathcal{U}} > \mathcal{U}^h$, a contradiction. Hence, the seller cannot screen the l-type. Her choice is to either offer $\mathcal{U}^h$, and have both types accept, or make an unacceptable offer. Given that $\Pi^h > 0$, the former dominates the latter.

Bibliography

Athanasopoulou, P. (2009). "Relationship quality: a critical literature review and research agenda," *European journal of marketing*, **43**, 583–610.

Baik, S.A.A., and N. Larson (2023). "Price discrimination in the information age: Prices, poaching, and privacy with personalized targeted discounts." *The Review of Economic Studies,* **90**, 2085–2115.

Battaglini, M. (2005). "Long-Term Contracting with Markovian Consumers," *The American Economic Review*, **95**, 637–658.

Bergemann, D., and J. Välimäki (2002). "Strategic Buyers and Privately Observed Prices," *Journal of Economic Theory*, **105**, 469–482.

Chevalier, J.A., and A.K. Kashyap (2019). "Best prices: Price discrimination and consumer substitution," *Amer-*

ican Economic Journal: Economic Policy, **11**, 126–159.

Crémer, J. (1984). "On the economics of repeat buying," *Rand Journal of Economics*, **15**, 396–403.

Crémer, J., and R.P. McLean (1988). "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, **56**, 1247–1257.

Ehrenberg, A. (2000). "Repeat-buying: facts, theory and applications," *Journal of Empirical Generalisations in Marketing Science*, **5**, 392–770.

Fudenberg, D., and J.M. Villas-Boas (2006). "Behavior-Based Price Discrimination and Customer Recognition," in *Handbook on Economics and Information Systems*, edited by T. Hendershott, Elsevier, Bingley, UK.

Gallo, A. (2014). "The Value of Keeping the Right Customers." *Harvard Business Review*. Retrieved from: `https://hbr.org/2014/10/the-value-of-keeping-the-right-customers`.

Gerstner, E., and J.D. Hess (1991). "A theory of channel price promotions." *The American Economic Review*, **81**, 872–886.

Gerstner, E., Hess, J.D., and D.M. Holthausen (1994). "Price discrimination through a distribution channel: Theory and evidence," *The American Economic Review*, **84**, 1437–1445.

Hart, O., and J. Tirole (1988). "Contract Renegotiation and Coasian Dynamics," *The Review of Economic Studies*, **55**, 509–540.

Hosios, A., and M. Peters (1993). "Self-Selection and Monitoring in Dynamic Incentive Problems with Incomplete Contracts," *The Review of Economic Studies*, **60**, 149–174.

Jullien, B. (2000). "Participation Constraints in Adverse Selection Models," *Journal of Economic Theory*, **93**, 1–47.

Keller, G., and S. Rady (2001). "Strategic experimentation with Poisson bandits," *Theoretical Economics*, **10**, 275–311.

Kennan, J. (2001). "Repeated Bargaining with Persistent Private Information," *The Review of Economic Studies*, **68**, 719–755.

Klemperer, P. (1995). "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *The Review of Economic Studies*, **62**, 515–539.

Mussa, M., and S. Rosen (1978). "Monopoly and Product Quality," *Journal of Economic Theory*, **18**, 301–317.

Nevo, A., and C. Wolfram (2002). "Why do manufacturers issue coupons? An empirical analysis of breakfast cereals," *RAND Journal of Economics*, **33**, 319–339.

Schmalensee, R. (1982). "Differentiation Advantages of Pioneering Brands," *The American Economic Review*, **72**, 349–365.

Villas-Boas, J.M. (2004). "Consumer Learning, Brand Loyalty, and Competition," *Marketing Science*, **23**, 134–145.