Sales Talk, Return Policies, and the Role of Consumer Protection*

Roman Inderst† Marco Ottaviani‡

April 2009

Abstract

This paper analyzes the role of product return and contract cancellation policies in markets in which sellers advise buyers about the suitability of the products sold. By granting buyers the right to return at favorable terms or to cancel early, the seller's "cheap talk" at the point of sale becomes more credible. When all buyers are wary of the seller's incentives, equilibrium contractual provisions are second-best efficient, but involve excessive purchases (*ex ante* inefficiency) and excessive returns (*interim* inefficiency). Imposition of a minimum refund standard (even if not binding) improves welfare by reducing the seller's incentives to target credulous buyers.

*We thank seminar participants at York University and Ohio State University.

†Johann Wolfgang Goethe University Frankfurt (IMFS) and Imperial College London. E-mail: inderst@finance.uni-frankfurt.de.

‡Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208-2013, USA. E-mail: m-ottaviani@northwestern.edu.

1 Introduction

Many specialized products and long-term service plans, such as subscriptions to magazines and utility contracts, are "not bought but sold". Sellers or their representatives actively market these products by eliciting interest often through unsolicited telephone calls ("cold calling") or visits at the buyer's doorstep. While buyers may often have little acquaintance with the product, the seller may be in a better position to quickly sound out the likelihood that the product is suitable for the buyer's specific preferences and needs.

In these circumstances, there is a serious concern that buyers might end up purchasing products (or signing for services) they do not really need or want. In an attempt to protect consumers from aggressive marketing techniques, many countries have put in place legislation that allows buyers to return products at no penalty within a specified "cooling-off period".[1] Similarly, regulators sometimes limit the penalties sellers can impose on buyers who cancel long-term contracts. However, sellers also have an incentive to assure buyers through the contractual terms of their offer.[2] It is then natural to wonder: What are the properties of the contractual arrangements for product returns and contract cancellations that we should expect sellers to offer in the absence of policy intervention? When is policy intervention to protect consumers justified? What form should consumer protection policies take?

To address these questions, we first analyze the baseline outcome we expect to arise in an ideal world with fully rational buyers, in the absence of policy intervention. After observing a noisy signal about the suitability of product characteristics for the buyer's specific needs, the seller advises the buyer regarding purchase or, likewise, makes unverifiable claims that render the product more (or less) valuable to the buyer. After purchasing the product, the buyer becomes better informed about the product's suitability, and then

[1] In the U.S., for example, the Federal Trade Commission 'Cooling-Off Rule' gives buyers three days to cancel purchases of $25 or more when these purchases are made away from the seller's premises, with the exception of real estate, insurance, securities, motor vehicles sold at temporary locations, and arts or crafts sold at fairs. In the E.U., the 'Doorstep Selling' Directive 85/577/EEC protects consumers who purchase goods or services during an unsolicited visit by a seller at their doorstep (or otherwise away from the seller's business premises). This regulation provides a cooling-off period of seven days enabling the buyer to cancel the contract within that period and making the contract unenforceable if the buyer is not informed in writing of this right. Similar regulations are in place in most industrialized countries (see Office of Fair Trading 2004, Annex E).

[2] For example, sellers often offer generous return policies, thereby bearing losses when products are returned. In addition, regulations on cooling-off periods often originate from self-regulatory provisions set in place by sellers and their associations (such as the Direct Selling Association).

decides whether or not to return the product, according to the contractual terms offered by the seller (and possibly affected by the regulator). In this setting, a first-best efficient return (or cancellation) policy would ensure that the product is ultimately consumed (or that the contract is served until maturity) only if the buyer's expected utility, given the information available at that stage, exceeds the salvage value of the product (or the ensuing cost savings for the seller from early cancellation of a long-term contract).

While the seller may want to trick buyers into purchasing unsuitable products, rational buyers—being wary of this conflict of interest—should not be susceptible to such cheap talk. Our first result is that the seller can make the product recommendation more trustworthy by granting buyers a generous refund (or the right of early contract cancellation at a low penalty). The refund policy thus becomes an effective commitment device to improve the quality of information transmitted to the buyer. Thus, our baseline model with fully rational buyers offers a simple explanation for why sellers would voluntarily grant buyers beneficial terms for refunds (or cancellations).[3]

The equilibrium refund policy involves two types of inefficiencies compared to the first-best benchmark. First, some buyers purchase the product even though the seller *knows* that the expected joint surplus is negative. Second, some buyers end up canceling or returning the product even though, at that stage, it would be efficient not to do so. The seller's optimal policy involves thus too many refunds (or early cancellations) *both* because too many buyers sign up initially *and* because even those buyers for whom it was efficient and privately optimal to purchase the product end up returning too often. However, when all buyers are wary of the seller's incentives, we demonstrate that the seller's optimal pricing policy is second-best efficient.

Taking the level of ultimately dissatisfied buyers (or the high level of cancellation requests) as an indication of market failure and thus as a justification for policy intervention is misleading. In particular, a consumer protection policy that is designed to protect those remaining buyers who, without their knowledge, still sign a contract even though it is against their own best interest would reduce overall efficiency, because it would induce yet more inefficient cancellations or returns. However, there is a role for consumer protection policy when not all buyers are wary of the seller's incentives.

[3] Our model thus offers a simple explanation for the "excess-refund puzzle" (cf. the discussion in Matthews and Persico 2007) for sales that require initial advice.

While wary buyers see through the seller's cheap talk and use rational expectations to correctly back out the conditional distribution of the product's value for them, credulous buyers are willing to sign up under inflated perceptions. The seller can exploit these buyers by offering a contract with less generous cancellation rights. We show that, as long as the firm still sells to both wary and credulous buyers, a mandatory requirement that all buyers can later cancel a contract or seek a refund under the *most beneficial terms* granted to any buyer will still lead to the (second-best) efficient outcome. credulous buyers are then sufficiently protected by the presence of wary buyers. However, for products and services that are targeted only at credulous buyers there is scope for beneficial policy intervention that prescribes a *minimum standard* for buyers' right to seek a refund or to cancel early.[4] Such a minimum standard becomes effective even though it may prove not to be binding, because sellers offer more generous return terms than required by the regulator.

Our model thus suggests a rethinking of minimum statutory provisions, restricting them to products and sales channels for which it is reasonable to expect a predominance of credulous buyers. In addition, we find competition to be instrumental in mitigating or even resolving inefficiencies, because sellers whose margin are constrained by competition have lower incentives to provide unsuitable advice.

Literature. Our model embeds a simple game of strategic information transmission (Crawford and Sobel 1982 and Green and Stokey 2007) into a trading environment. Whereas in other analyses of strategic communication the conflict of interest between the sender (seller) and the receiver (buyer) is exogenously given (as in Pitchik and Schotter 1987), in our model the degree of preference alignment is endogenously determined through the contractual provisions for cancellation and refund offered *ex ante* by the seller.[5] In our setting, sorting results from the correlation of the seller's information with the information the buyer obtains by experiencing the product after purchase. In equilibrium, interests are not perfectly aligned, so that the seller is willing to induce some *ex ante* inefficiency at the contracting stage to reduce the inefficiently high return costs incurred at the *interim*

[4] Such statutory provisions are pervasive in many countries and often apply equally to doorstep sales as well as to distance sales through catalogue orders or the internet. Similar to the provision of information in our model, also with such distant purchases sellers should out of self-interest allow customers the right to return products or to cancel contracts, as otherwise wary customers may not believe the asserted features or qualities.

[5] In Inderst and Ottaviani (forthcoming), instead, advice is provided by a seller's employee, whose preferences depend on the incentives set by the seller.

stage.

Several papers, such as Davis et al. (1995), have analyzed the role of refunds to boost demand in a setting in which buyers are *ex ante* uncertain about a product's suitability. In particular, Che (1996) analyzes the role of buyer risk aversion, while Matthews and Persico (2007) show how refunds can lower buyers' cost of learning. Courty and Li (2000) show, instead, how a menu of refunds can be used to price discriminate across *ex ante* private informed buyers.[6] There is also a literature on the role of warranties, which generally are conditioned on some verifiable events, such as non-performance. The exercise of cancellation or return in our setting depends instead on *personal* dissatisfaction.[7]

The rationale for policy intervention this paper offers is different from that suggested by behavioral models that build on buyers' projection bias (Loewenstein, O'Donoghue, and Rabin 2003). That literature suggests that buyers who are unaware of their *upwards biased* perception at the time of purchasing must be protected "from themselves". In our model, instead, only credulous buyers must be protected from sellers' cheap talk—and that this is only necessary in certain circumstances and may often be simply achieved through a "non-discrimination" (also known as "most-favored customer") clause. While our model takes as exogenous the presence of a fraction of credulous buyers who are overly susceptible to the seller's influence, such a behavioral trait may be related to lack of experience dealing with self-interested salespeople.[8]

The rest of the paper is organized as follows. Section 2 introduces the model. Sections 3 and 4 analyze the benchmark case in which all customers are wary of the seller's incentives when providing advice. Section 5 introduces the possibility that some customers are instead credulous. Section 6 offers some concluding remarks. Proofs that are not shown in the main text are contained in the Appendix.

[6] In case (perfect) price discrimination is not feasible or too costly, Inderst and Tirosh (2009) show the seller is able to extract more consumer surplus on average by setting an inefficiently lenient refund policy that induces a "rotation" in the demand for its product. See also Lewis and Sappington (1994) and Johnson and Myatt (2006) on a seller's incentives to improve buyers' information about their willingness to pay for the product.

[7] The commitment role of return policies is also key in Hendel and Lizzeri's (2002) and Johnson and Waldman's (2003) models of leasing under asymmetric information. While in those models the redemption price set by the seller affects the quality of products return and therefore the informational efficiency in the second-hand market, in the pesent model the refund (or price for continuing service) offered by the seller affects the seller's own incentives to report information.

[8] See Spence (1977) for an early analysis of market outcomes when consumers misperceive quality. In our setting, misperceptions by customers are induced by the firm, rather than being exogenous.

2 Model

To be specific, consider the sale of a physical product. The key feature of our model is that at an initial encounter between a seller and a potential buyer, the seller has better information about the suitability of the product for the buyer's specific needs and preferences. After a purchase is made but before the product or contract is fully consumed, the buyer can additionally learn about the ultimate consumption value. At this stage, the buyer is allowed to ask for a refund, according to the terms specified by the contract initially offered by the seller.

The timing is as follows. At time $t = 0$, the seller decides on the sales price p and a refund q at which the buyer can potentially return the product.[9] At time $t = 1$, the seller observes signal s and then advises the buyer whether or not to purchase the product at price p. The seller incurs a production cost equal to $c > 0$ if the product is sold. After purchasing the product, at $t = 2$ the buyer observes signal b, capturing superior information about the expected utility obtained from using the product or trying the service contract.[10] At that stage, the buyer may choose to return the product and ask for a refund q, in which case the seller derives the salvage value v from the returned product. When the product is not returned, it is then consumed at $t = 3$, resulting in consumption benefits equal to u for the buyer.[11] There is no discounting, and utilities are additively separable in money.

From an *ex ante* perspective, the buyer's utility from the product, u, is distributed with distribution $F_u(u)$ over $U := [\underline{u}, \overline{u}]$, where $0 \leq \underline{u} < \overline{u}$ and where $f_u(u) > 0$ holds for all $u \in U$. We assume that that the salvage value is less than the cost of the product and that under full information it is efficient to trade the product for sufficiently high realizations of u, whereas trade is inefficient for sufficiently low realizations:

$$\underline{u} < v < c < \overline{u}. \tag{1}$$

While we present our results for the case of a purchase and return of a physical prod-

[9] By restricting attention to a simple offer of (p, q) we implicitly rule out stochastic refund policies as well as more general mechanisms. We expect that our key result that an excessive refund policy serves as a commitment to enhance the credibility of advice to hold more generally.

[10] As noted in the introduction, this may be the case when the buyer learns from early experimentation after delivery or after the service contract is set in place.

[11] This assumption is made for convenience, but our results would also hold if part of the consumption benefits were accrued at period $t = 2$.

uct, our setting also covers long-term service agreements, such as phone plans or utility contracts. In that case, all our results continue to hold with the following adoption of notation. Then, the initial cost of setting up the service is $c - v$ and the cost of service continuation is v, while the contract involves an initial payment of $p-q$ and a continuation payment of q.

Information. While the buyer may not have much prior knowledge of the product, the seller may be able to make some observations "at the doorstep" about the product's suitability for the buyer's preferences and needs. Thus, we assume that at $t = 1$, before the buyer purchases, the seller privately observes a first, informative signal $s \in S := [\underline{s}, \overline{s}]$. After a purchase is made but before a return decision must be made, at $t = 2$ the buyer privately observes a more informative signal $b \in B := [\underline{b}, \overline{b}]$.

To capture these specifications in a convenient way, the primitives of our model are the distributions that statistically relate the final consumption value, u, to the buyer's signal, b, and to the seller's signal, s. We stipulate that the more informative signal b is generated from the continuous distribution $G_b(b|u)$, while the less informative signal s is generated from the continuous distribution $H(s|b)$, with both families of distributions satisfying the Monotone Likelihood Ratio Property (MLRP). These key properties ensure that higher signals b and s are both good news about the ultimate realization of u. Note also that b is a sufficient statistics for s.

For convenience, we assume that both b and the noisier signal s are unboundedly informative. Precisely, we stipulate that $\frac{g_b(b|u)}{g_b(b|\underline{u})} \to 0$ for all $u > \underline{u}$ as $b \to \underline{b}$, $\frac{g_b(b|u)}{g_b(b|\overline{u})} \to 0$ for all $u < \underline{u}$ as $b \to \overline{b}$, $\frac{h(s|b)}{h(s|\underline{b})} \to 0$ for all $b > \underline{b}$ as $s \to \underline{s}$, and $\frac{h(s|b)}{h(s|\overline{b})} \to 0$ for all $b < \overline{b}$ as $s \to \overline{s}$. In words, after observing very high signals b or s it is almost sure that the highest utility obtains, whereas after observing very low signals b or s it is almost sure that the lowest utility obtains.[12]

The joint specification of $F_u(u)$, $G_b(b|u)$, and $H(s|b)$ exhausts the description of the information technology, given that together they pin down the informativeness of the seller's initial signal and of the buyer's subsequent signal. However, it is convenient to introduction some additional notation to shorten the exposition of the analysis that follows.

In analogy to the distribution $G_b(b|u)$, we can define for signal s the distribution

[12] These specifications allow to rule out corner solutions and case distinctions, but they are not necessary to obtain our main results.

$G_s(s;u)$ with density $g_s(s|u) := \int_B h(s|b)g_b(b|u)du$. This describes the distribution of the signal conditional on the ultimately realized utility u. Using $G_b(b|u)$ and $G_s(s;u)$, we can from Bayes' rule calculate the posterior distributions $\Phi_s(u|s)$ and $\Phi_b(u|b)$ given the observation of s and b, which have the respective densities

$$\phi_s(u|s) = \frac{g_s(s|u)f_u(u)}{\int_U g_s(s|\tilde{u})f_u(\tilde{u})\,d\tilde{u}} \qquad \text{and} \qquad \phi_b(u|b) = \frac{g_b(b|u)f_u(u)}{\int_U g_b(b|\tilde{u})f_u(\tilde{u})\,d\tilde{u}}.$$

Note that by sufficiency, the posterior distributions conditional on both b and s is the same as the posterior conditional on b alone.[13]

Next, from an *ex ante* perspective, s and b are distributed according to the distributions $F_b(b)$ and $F_s(s)$. Given the distribution of utilities $F_u(u)$ and the conditional distributions $G_b(b|u)$ and $G_s(s;u)$ for the respective signals, the densities of these distributions are $f_b(b) := \int_U g_b(b|\tilde{u})f_u(\tilde{u})\,d\tilde{u}$ and $f_s(s) := \int_U g_s(s|\tilde{u})f_u(\tilde{u})\,d\tilde{u}$, respectively. As a final bit of notation, we introduce the distribution $\Psi(b|s)$ of b conditional on s, with density

$$\psi(b|s) = \frac{h(s|b)\,f_b(b)}{\int_B h(s|\tilde{b})f_b(\tilde{b})d\tilde{b}}.$$

The distribution $\Psi(b|s)$ captures the seller's beliefs about the signal that the buyer will observe after purchasing the product. This distribution will be instrumental in expressing the seller's decision problem at the advice stage.

Efficient Decision Rules. We begin by deriving the efficient decision rules, at periods $t = 1$ and $t = 2$. Starting backwards, once a purchase is made in $t = 1$, at $t = 2$ it is efficient to consume the product if the expected utility conditional on the (sufficient) signal b does not fall short of the salvage value v:

$$E[u|b] := \int_U u\phi_b(u|b)du \geq v. \tag{2}$$

Otherwise, it is efficient to return the product. Note that $E[u|b]$ is strictly increasing in b by the MLRP of $G_b(b|u)$, which implies strict First-Order Stochastic Dominance (FOSD) of the posterior distribution $\Phi_b(u|b)$.

Given this *interim* efficient decision rule, at $t = 1$ it is *ex ante* efficient to purchase whenever

$$\int_B \max\{E[u|b], v\}\,\psi(b|s)db \geq c, \tag{3}$$

[13] Given that b is a sufficient statistics for s, the posterior distribution satisfies $\phi(u|b,s) = \frac{g_b(b|u)h(s\ |b)f_u(u)}{\int_U g_b(b|u)h(s|b)f_u(u)} = \frac{g_b(b|u)f_u(u)}{\int_U g_b(b|u)f_u(u)} = \phi_b(u|b)$.

which uses the conditional distribution $\Psi(b|s)$ derived from Bayes' rule. The left-hand side of (3) is strictly increasing in s, which follows because $\psi(b|s)$ satisfies FOSD, which in turn holds by the MLRP of $H(s|b)$.

Together with continuity and the assumed informativeness properties of the signal distributions at the boundaries of the respective supports B and S, the first-best efficient decision rules (2) and (3) pin down unique and interior *interim* and *ex ante* cutoffs, $\underline{b} < b_{FB} < \overline{b}$ and $\underline{s} < s_{FB} < \overline{s}$. The initial purchase should be made only if $s \geq s_{FB}$, while subsequently the product should be returned only if $b < b_{FB}$ and consumed if $b \geq b_{FB}$.[14]

For the following analysis, it also convenient to extend the concept of the *ex ante* efficient decision rule to the case in which the subsequent, *interim* decision rule is not efficient. Suppose thus that the *interim* decision rule is characterized by any cutoff b^*, which may be different to b_{FB}. In what follows, the relevant case satisfies $b^* > b_{FB}$. For this case, there exists a unique cutoff $\widetilde{s}_{FB}(b^*)$ such that, given the subsequently applied cutoff b^*, an initial purchase is efficient if and only if $s \geq \widetilde{s}_{FB}(b^*)$. If interior, $\widetilde{s}_{FB}(b^*)$ is given by

$$\Psi(b^*|\widetilde{s}_{FB}(b^*))v + \int_{b^*}^{\overline{b}} E[u|b]\psi(b|\widetilde{s}_{FB}(b^*))db = c. \tag{4}$$

Note that for the existence of the cutoff we use: (i) the property that an increase in signal s induces a first-order stochastic shift in the conditional distribution of $b|s$ and (ii) the fact that, when the interim return decision is made according to the cutoff $b^* > b_{FB}$, the interim social surplus conditional on b increases in b.[15] In addition, when $b^* > b_{FB}$ holds, $\widetilde{s}_{FB}(b^*)$ is strictly increasing in b^*. Thus, when the product is returned more often than it is efficient at the *interim* stage, it is *ex ante* efficient to apply a higher cutoff. This is intuitive because the application of an inefficiently high *interim* cutoff b^* implies a reduction in the maximum surplus that can be realized, for given *ex ante* signal s.[16]

3 Returns and Advice in Equilibrium

To characterize the equilibrium without policy intervention, we proceed by backward induction. In Section 3.1 we analyze the buyer's return decision at $t = 2$ following purchase.

[14] Note that each of the signal realizations $s = s_{FB}$ and $b = b_{FB}$ has zero probability.

[15] To see that the conditional social surplus increases in b when $b^* \geq b_{FB}$, note that it is equal to v for $b < b^*$ and equal to $E[u|b]$ for $b \geq b^*$, where $E[u|b^*] > v$.

[16] This comparative statics result is obtained from implicit differentiation of (4), using $E[u|b^*] - v > 0$ and that $\Psi(b|s)$ shifts in the FOSD order.

In Section 3.2 we turn to the seller's advice and the buyer's purchase decision at $t = 1$. In Section 3.3 we then solve for the price p that is set in equilibrium at period $t = 0$, for any given refund q. Finally, in Section 4 we derive the level of refund q set by the seller in equilibrium.

3.1 Returns

After purchasing at $t = 1$, at $t = 2$ the buyer optimally chooses to keep the product whenever the expected utility given the observed signal b is not below the level of the refund, i.e., whenever $E[u|b] \geq q$. When $E[u|\underline{b}] = \underline{u} < q$ and $E[u|\overline{b}] = \overline{u} > q$ hold, then this decision rule gives rise to a unique cutoff rule

$$E[u|b^*] = q, \tag{5}$$

with $\underline{b} < b^* < \overline{b}$. For $q \geq \overline{u}$ the product would always be returned, which from (1) would not allow to make positive profits. For $q \leq \underline{u}$ the product would never be returned. This case (to be ruled out below) is captured by setting $b^* = \underline{b}$.

Note that the cutoff b^* in (5) is strictly increasing in q. Also, the buyer's privately optimal decision whether to return the product is only *interim* efficient in case the refund q just matches the salvage value v. Instead, for $q > v$ the product would be returned too frequently ($b^* > b_{FB}$), while for $q < v$ it would be returned too infrequently ($b^* < b_{FB}$).

3.2 Advice

Turn now to the buyer's purchasing decision in $t = 1$, where the seller privately observes s. At this stage, our model gives rise to a game of "cheap talk". Given that the buyer's decision is binary, we can restrict consideration to a binary message set for the seller, according to whether the seller advises the buyer to purchase or not the product. In what follows, we restrict consideration to the informative equilibrium.[17] In equilibrium, the buyer follows the seller's advice.

Given a price p and under the assumption that the buyer follows the seller's advice, the seller prefers to advise the buyer to purchase if

$$\pi(s) := (p - c) + \Psi(b^*|s)(v - q) \geq 0. \tag{6}$$

[17] As is well known, in any cheap-talk game there is always a "babbling" equilibrium in which the seller's message has no information content.

Here, $\Psi(b^*; s)$ denotes the probability that the buyer will subsequently return the product, as assessed by the seller conditional on observing a signal realization equal to s.

When $v = q$ holds, which would lead to the *interim* efficient return policy with $b^* = b_{FB}$, then as long as the up-front margin is positive, $p - c > 0$, the seller would *indiscriminately* want to advise the buyer to purchase. When instead the refund paid to buyers following product return lies above the salvage value, the seller faces the following trade-off when advising customers. When a purchase is made, the seller earns the up-front margin $p - c$, but risks making a subsequent loss in case the buyer returns the product and obtains a refund exceeding the seller's salvage value. Note that when the seller observes a low realization of the signal s, it becomes more likely that the buyer will observe a signal $b < b^*$ and so will ask for a refund. Consequently, for $q > v$ the seller has an incentive to advise a purchase only when privately observing sufficiently high values of s.[18]

Proposition 1 *Suppose the buyer follows the seller's advice and that the seller enjoys a positive sale margin, $p - c > 0$. If the refund is set below the salvage value, $q \leq v$, then the seller always advises to purchase. If instead $q > v$, then there exists an interior cutoff $\underline{s} < s^* < \overline{s}$, characterized by*

$$\pi(s^*) = 0, \tag{7}$$

at which the seller only advises the buyer to purchase when $s \geq s^*$ and not to purchase otherwise. In addition, s^* is strictly increasing in q and strictly decreasing in p.

As the refund q increases, it becomes more costly for the seller to advise the buyer to purchase. In particular, when the seller observes a lower signal realization s, this implies a higher probability that also the buyer's subsequent observation of b will be sufficiently low so as to make the buyer return the product. On the other hand, the higher is the initial price p, the more the seller is tempted into advising the buyer to purchase the product even after observing lower signal realizations s.

3.3 Pricing Equilibrium

While Proposition 1 conducts a comparative analysis for different levels of the initial price p, in equilibrium this price is chosen at $t = 0$. When determining the optimal offer (p, q)

[18] Proposition 1 restricts consideration to the case in which $p > c$. For the following characterization, this is indeed the only relevant case. In particular, it is straightforward to rule out the case in which the seller initially makes a loss, $p < c$.

the seller's objective function is

$$\Pi := \int_{s^*}^{\overline{s}} \left[(p-c) + \Psi(b^*|s)(v-q)\right] f_s(s)ds. \tag{8}$$

Although the seller initially designs the offer (p, q), for what follows it is useful to take a slightly different perspective. We first take some refund q as given and solve for the resulting pricing equilibrium. This is described by a prevailing price p together with the prevailing cutoff s^* that the seller applies at the advice stage—all as a function of q.

Note that the buyer (weakly) prefers to follow the seller's advice to purchase, instead of choosing not to purchase, whenever

$$\int_{s^*}^{\overline{s}} \left[\int_B \max\left\{E[u|b], q\right\} \psi(b|s)db\right] \left(\frac{f_s(s)}{1-F_s(s^*)}\right) ds \geq p, \tag{9}$$

which defines the buyer's participation constraint. Note that (9) uses the conditional beliefs held by the buyer when advised to purchase. In the present analysis, the buyer is thus assumed to see through the seller's incentives at the advice stage, as captured by the applied cutoff s^*. When choosing the product's price p for given refund level q, the seller sets p at the highest possible level consistent with the buyer's participation constraint, so that (9) must be binding.

Second-Best Efficiency. As long as we restrict consideration to offers with a positive initial margin $p > c$, the program is thus to maximize (8) subject to the restrictions that b^* is given by (5), that s^* is as characterized in Proposition 1, and that p satisfies the participation constraint (9) with equality for given s^*. Note that, after substitution of (9) into (8), we have that the seller's objective is to maximize the *ex ante* surplus:

$$\Pi = \int_{s^*}^{\overline{s}} \left[\int_{b^*}^{\overline{b}} E[u|b]\psi(b|s)db + \Psi(b^*|s)v - c\right] f_s(s)ds. \tag{10}$$

From this observation it follows immediately that the seller's privately optimal offer is also second-best efficient.

Uniqueness of Pricing Equilibrium. For given refund q and cutoff s^*, (9) pins down a unique price from the buyers' participation constraint. For any given $p > c$, in turn, the cutoff s^* is uniquely determined from Proposition 1—i.e., either $s^* = \underline{s}$ or $s^* > \underline{s}$ solves (7). When a solution p with $p > c$ exists for these two conditions, this solution

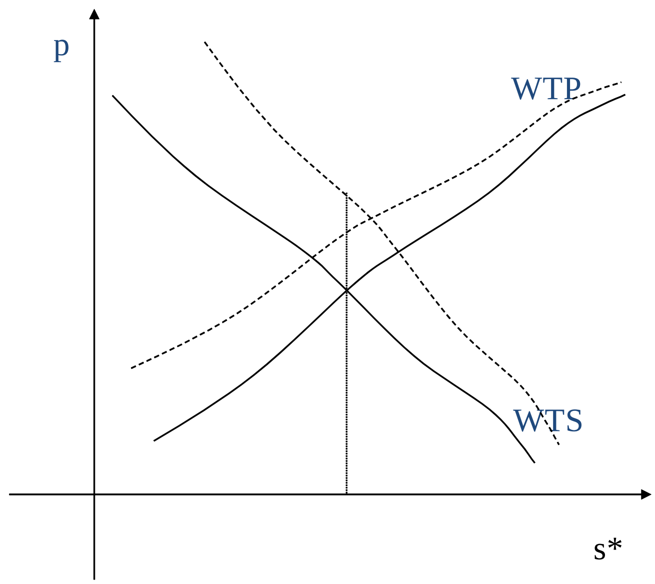

Figure 1: Pricing equilibrium and effect of higher refund q

is unique. To see this, note first that, holding q fixed, (9) defines a strictly increasing mapping from s^* to p. The higher is the seller's cutoff, the more the buyer is willing to pay when advised to purchase the product—as depicted by the upward-sloping willingness-to-pay (WTP) curves in Figure 1 for two levels of q. On the other hand, from condition (7) in Proposition 1, the seller's optimality condition defines a decreasing mapping from p to s^*—the willingness-to-sell (WTS) curves in Figure 1 for two levels of q. Holding fixed any given q, the equilibrium contract offered by the firm is characterized by the crossing of the WTP and the WTS curve.

To further streamline the exposition, we focus on the case in which it is not possible for the seller to sell with probability one and still make profits:

$$\int_{\underline{s}}^{\overline{s}} \left[\int_B \max \left\{ E[u|b], v \right\} \psi(b; s) db \right] f_s(s) ds \leq c, \tag{11}$$

Note here that the term on the left-hand side of (11) captures the maximum expected surplus the seller extracts from the buyer when selling with probability one.

Proposition 2 *Suppose condition (11) holds. Then for any refund $q > v$ there is a unique cutoff $\underline{s} < s^* < \overline{s}$ and a unique price $p > c$ at which both (7) and (9) hold as equalities: The price p corresponds to customers' willingness-to-pay given s^*, while given p the seller advises customers to purchase if $s \geq s^*$. Instead, there is no trade with positive probability* *when $q < v$, while for $q = v$ there can only be trade with positive probability when $p = c$, in which case the seller makes zero profits.*

When $q = v$, we know that the seller would like to advise all customers to purchase as long as the up-front margin is positive, $p > c$. However, the buyer would then obtain a negative expected surplus according to (11), so that this scenario is not compatible with equilibrium.[19] When both $q = v$ and $p = c$, the seller is indifferent between advising customers to purchase or not for all observed signals s. However, this case is not relevant because the seller can realize strictly positive profits in equilibrium by choosing a different refund level, as we show below.

When (11) does not hold, then a strictly positive surplus is realized even when the seller advises all customers to purchase, and when this advice is indeed followed. Then, the case with $q = v$ would be compatible with trade (with probability one) and strictly positive profits, given $p > v$. Still, the subsequent results on the role of refunds would survive.

4 Refunds as Commitment

Proposition 2 still takes the refund level as given, determining the resulting pricing equilibrium, together with the prevailing cutoff for the seller's advice. Interestingly, as the seller increases the price to fully capture customers' willingness-to-pay, the seller's willingness-to-sell also increases, thereby pushing down s^*. Proposition 3 shows how the seller can use the refund to commit to apply a higher cutoff, while still extracting all consumer surplus.

Proposition 3 *Take some refund q that from Proposition 1 gives rise to a pricing equilibrium with an interior cutoff $\underline{s} < s^* < \overline{s}$ and a price p that satisfies (9) with equality. As the refund is now increased to $\widetilde{q} > q$, a new, unique pricing equilibrium results, characterized by a strictly higher cutoff $\widetilde{s}^* > s^*$ and a strictly higher price $\widetilde{p} > p$ satisfying (9) again with equality.*

The higher is the refund level, q, the higher is, ceteris paribus, the buyer's willingness to pay, given that the buyer's option of returning the product then becomes more valuable.

[19] The case with $q < v$ is even more immediate, because then the seller would gain more from a sale after observing a lower value s. Thus, in this case there would be no trade in equilibrium, given that the resulting expected surplus would be strictly negative by (11).

This allows the seller to charge a higher price, p. Holding now q fixed, a higher price in turn induces the seller to apply a strictly lower cutoff s^*, given that a sale becomes more profitable. But the lower is the cutoff that the seller applies, the lower will now be the customers' willingness-to-pay. Proposition 3 claims that the *joint* effect from the increase in q, which pushes s^* up, and the increase in p, which pushes s^* down, leads unambiguously to a higher cutoff: $\widetilde{s}^* > s^*$.

The intuition for this result is as follows. The buyer's higher willingness to pay is determined by the *expected* use that the buyer will make of the (higher) refund, where this expectation is taken conditional to the information available to the buyer when purchasing, $s \geq s^*$. For the determination of s^*, instead, the seller takes into account the expected costs at the (higher) refund, computed based on the information available to the seller when advising the *marginal* buyer to purchase the product, $s = s^*$. These costs increase by more than the buyer's willingness to pay, given that following a lower signal s also a lower realization of b becomes more likely. In total, for the seller the expected cost that a higher refund generates at the marginal signal s^* is not fully compensated by the price increase as the latter only reflects buyers' expected benefits from the higher refund over all higher signals $s \geq s^*$. After the increase in the refund, at the previous cutoff signal s^* the seller then strictly prefers to advise the buyer *not* to purchase.[20]

Figure 1 depicts the resulting shifts in the willingness-to-sell and in the willingness-to-pay curves—the dashed curves correspond to a higher level of q than the continuous curves. After an increase in the refund, both the willingness-to-pay (WTP) and the willingness-to-sell (WTS) curves shift upwards, intersecting now at a higher price p, but also at a higher cutoff s^*.

4.1 Optimal Refund Policy

We are now in a position to characterize the seller's optimal offer at $t = 0$.

Proposition 4 *The optimal offer (p, q) specifies $q > v$ and leads to two types of inefficiencies: (i) From $b^* > b_{FB}$ there is an inefficiently high level of returns at the interim stage;*

[20]The proof of Proposition 3 reveals that there is an additional effect at work that goes in the same direction. When $q > v$ is further increased, this further exacerbates the *interim* inefficiency. Holding s^* constant and adjusting p so as to make the buyer indifferent, the resulting loss in surplus (for any given $s \geq s^*$) is borne by the seller, which further induces the seller to reduce s^*. (This effect, however, vanishes as $q \to v$, while the effect discussed in the main text survives also then.)

(ii) From $s^* < \widetilde{s}_{FB}(b^*)$, given the subsequently applied cutoff b^*, there is an inefficiently high level of initial purchases. The equilibrium outcome is second-best efficient.

The intuition for Proposition 4 is as follows. When $q = v$ holds, we know that the seller cannot make positive profits. This is because the seller would then want to indiscriminately advise the buyer to purchase for any price $p > c$, according to Proposition 1. But in this case the buyer's willingness to pay, given by the left-hand side of (11), is in fact strictly below the seller's cost. Take now the strictly interior signal $\underline{s} < \tilde{s} < \underline{s}$ at which

$$\int_{\tilde{s}}^{\overline{s}} \left[\int_B \max\left\{E[u|b], v\right\} \psi(b|s)db \right] f_s(s)ds = c$$

holds. When $s^* = \tilde{s}$, setting p equal to the buyer's willingness to pay results in $p = c$. When the seller now slightly increases the refund to $q > v$ and adjusts p so as to still satisfy the buyer's participation constraint (9), we know from Proposition 3 that s^* and p both (marginally) increase, ensuring that $p > c$. However, this observation does not yet imply that the seller makes positive profits, given that the seller now makes a loss on any returned product. The seller realizes strictly positive profits only if the *ex ante* expected surplus in (10) is strictly positive. To see that this is indeed the case, note that the first-order effect that a marginal increase in $q > v$, and thus a marginal increase in b^*, has on social surplus is *zero* at $q = v$. On the other hand, from $\tilde{s} < s_{FB}$ the resulting increase in the *ex ante* cutoff s^*, which follows from Proposition 3, has a strictly positive first-order effect on social surplus.[21] Taken together, the joint effect of the (inefficient) increase in the *interim* cutoff b^* and the (efficient) increase in the *ex ante* cutoff s^* is thus to increase total surplus.

In principle, it would be possible to further raise the refund (and, consequently, also the price) until the *ex ante* cutoff becomes (first-best) efficient. At that point, the seller would advise customers to purchase if and only if this is indeed efficient, given the subsequently applied *interim* cutoff b^*: $s^* = \widetilde{s}_{FB}(b^*)$. However, this is not optimal for the firm. Starting from an offer that induces $s^* = \widetilde{s}_{FB}(b^*)$, the firm can realize strictly higher profits by decreasing q. To see this, note that now the first-order effect on total surplus from a marginal change of s^* is zero, while the increase in b^* has a strictly positive effect on *interim* efficiency, given that $q > v$.

[21] Note here that s_{FB} is derived for the interim efficient decision rule, which obtains at $q = v$.

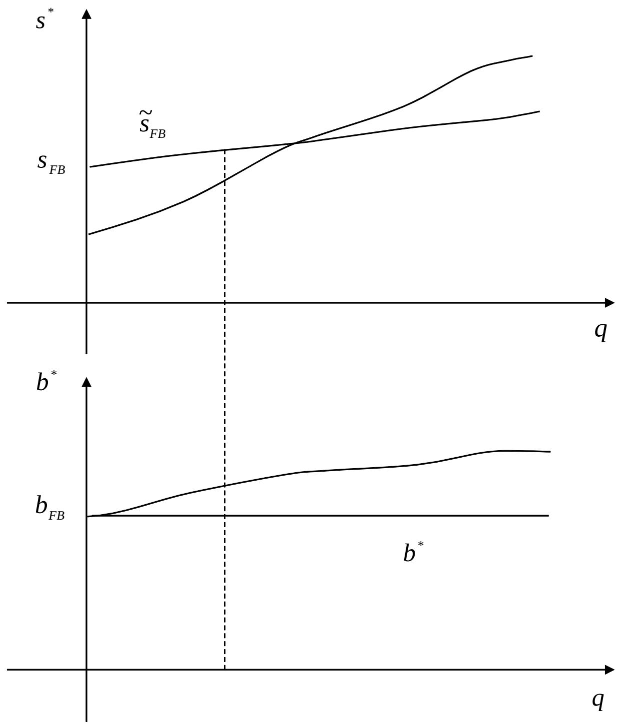

Figure 2: Tradeoff between *ex ante* and *interim* inefficiencies

Altogether, at the optimal offer the firm optimally trades off *interim* with *ex ante* efficiency. As remarked above, the firm's choice of contract is also second-best efficient because the firm extracts all consumer surplus. Figure 2 illustrates the resulting trade-off. As the refund increases, given the simultaneous adjustment in the price, the *ex ante* cutoff s^* increases (top panel) and the *interim* cutoff b^* increases (bottom panel). The top panel also depicts the second-best efficient cutoff $\widetilde{s}_{FB}(b^*)$, which is a strictly increasing function of q, given that an increase in q leads to an even less efficient *interim* cutoff b^*. At the optimal offer (p, q), we have both $s^* < \widetilde{s}_{FB}(b^*)$ and $b^* > b_{FB}$.

4.2 Discussion

At the offer that is optimally made by the seller there will be too many refunds for two reasons. Too many buyers sign up initially *and* even those buyers for whom purchase is *ex ante* efficient end up returning too often. The high level of returns is, however, not an indication of market failure and thus a justification for policy intervention.

Note next that given that buyers are initially just indifferent between purchasing or not, after observing their realized signal b, with positive probability they will regret the initial purchase. (Clearly, it always holds that $p > q$.) Moreover, if buyers could have direct access to the signal that is privately observed by the seller, s, for a strictly positive interval $s > s^*$ they would refuse to purchase and thus not follow the seller's advice. In other words, some buyers end up buying even though, given all information that is available, their expected utility is negative.[22] This set is also strictly larger than the set of buyers who bought even though total expected surplus was negative: $s \in [s^*, \widetilde{s}_{FB}(b^*))$.[23]

When policy intervention would prescribe a mandatory refund level, then if this deviates from the seller's optimal choice, total efficiency is strictly lower. Moreover, as long as the seller still has all pricing power, he will adjust the price so that customers' expected surplus is still zero. While such policy intervention may thus affect the observed number of returns or while it may have an impact on the number of customers who are ultimately dissatisfied with their purchase, it has no effect on overall consumer surplus and reduces social efficiency. In Section 5 we reexamine the role of consumer protection policies when some buyers are not sufficiently wary to see through the seller's incentives at the advice stage.

4.3 Competition

For brevity's sake we consider the following short-cut model of competition in the spirit of a contestable market. We suppose that the seller only attracts a customer when offering expected utility, U, not below some reservation value $\overline{U} \geq 0$.[24] Hence, generalizing the participation constraint (9), it must hold that

$$U := \int_{s^*}^{\overline{s}} \left[\int_B \max\{E[u|b], q\} \psi(b|s) db \right] \left(\frac{f_s(s)}{1 - F_s(s^*)} \right) ds - p \geq \overline{U}. \tag{12}$$

[22] The latter feature would also survive in a slightly enriched model, in which customers differ in some privately observed characteristics. Rather than extracting all consumer surplus (as in our model in which the price p is pinned down by the binding participation constraint (9), the seller would then solve a standard monopoly pricing problem with downward-sloping demand, for any given q.

[23] This is immediate because at $\widetilde{s}_{FB}(b^*)$ the total surplus from purchasing is zero, while from $\widetilde{s}_{FB}(b^*) > s^*$ the seller's expected profits are strictly positive.

[24] Note that this formulation implies that a given customer purchases at most one product (instead of experimenting with several products that he returns after observing low signals). Also, a buyer who is advised not to purchase does not turn to a different seller, but realizes zero utility.

The feasible set of values U is bounded from above by the maximum utility that the seller can promise the buyer, $\overline{U}_{\max}$. We find that $\overline{U}_{\max}$ equals the maximum feasible surplus:

$$\overline{U}_{\max} = \int_{s_{FB}}^{\bar{s}} \left[\int_B \max\left\{E[u|b], v\right\} \psi(b|s) db \right] f_s(s) ds - c.$$

Consider now first how an increase in $\overline{U}$ affects the pricing equilibrium for given refund q, as derived in Section 3.3. While $\overline{U}$ clearly does not affect the firm's willingness-to-sell (i.e., the cutoff that the firm optimally applies for a given price), it affects the maximum price that the firm can charge for given s^*, namely through a (parallel) downward shift of the willingness-to-pay function in Figure 1. The resulting pricing equilibrium must thus result in a strictly lower price p and a strictly higher cutoff s^*.

Recall next that in equilibrium the seller's offer maximizes social surplus. After substituting from the resulting pricing equilibrium, the seller's remaining choice variable is the level of refund. This directly determines the buyer's *interim* cutoff (cf. (5)) and, through the pricing equilibrium, it also determines the seller's *ex ante* cutoff. For the purpose of the present argument only denote this by $\widehat{s}^*(q, \overline{U})$. As $\overline{U}$ increases, $\widehat{s}^*(q, \overline{U})$ increases, leading to strictly higher social surplus, for given q.[25] Because this property holds for all q and the seller's objective is to maximize social surplus, an increase in $\overline{U}$ leads to strictly higher social surplus also *at* the seller's (adjusted) optimal offer. Intuitively, the lower price that a higher reservation value $\overline{U}$ implies makes the commitment problem of the seller less severe. While this leads to higher social surplus, the seller is, however, clearly strictly worse off.

Proposition 5 *Suppose that in a contestable market the seller's offer must leave customers at least with expected utility $\overline{U} \in [0, \overline{U}_{\max}]$. Then the higher is the customers' reservation value $\overline{U}$, the lower are both the price p and the refund q. Consumer surplus and social surplus (efficiency) are strictly increasing in $\overline{U}$. As $\overline{U}$ approaches the maximum feasible value $\overline{U}_{\max}$, the refund becomes efficient, $q = v$, while the price equals the firm's cost of production, $p = c$.*

Note that Proposition 5 does not assert monotonicity of the seller's offer in $\overline{U}$. While it may at first be intuitive that as $\overline{U}$ increases, the refund and the price both gradually

[25] We use here that we can safely ignore all offers that would result in too conservative advice for the seller, $s^* > \widetilde{s}_{FB}(b^*)$.

decrease, leading to a gradual reduction of b^* and a gradual increase of s^*, this may not hold generally as the respective first-order condition for the seller's program depends on local properties of the distribution functions. Even if such monotonicity does not prevail, when $\overline{U}$ is sufficiently high such that buyers can extract almost all of the maximum feasible social surplus, competition provides an increasingly tight bound on the maximum *ex ante* and interim distortions, $|s^* - s_{FB}|$ and $|b^* - b_{FB}|$. For the comparison of the monopolistic outcome with this limiting case we can thus generally observe that with competition the refund policy is *less* beneficial for buyers. While this may at first be surprising, recall that in the present model with only wary customers the role of the refund is to credibly commit the seller to offer less biased advice. When the margin for a sale is reduced by competition, the seller's commitment problem at the advice stage is less extreme, thereby reducing the need for granting a high refund.

5 Credulous Buyers

Not all buyers may be in a position to see through the seller's strategic talk. From now, suppose that a fraction α of buyers is credulous and blindly believes the seller's cheap talk. The remaining buyers are wary and fully understand the strategic incentives at play, as in the baseline model. It immediately follows that the seller will claim that $s = \overline{s}$ when wishing to advise in favor of a purchase. While this claim is taken at face value by credulous buyers, wary buyers correctly infer that this only credibly contains the information $s \geq s^*$. While admittedly very simplistic (and confined to a setting where S has indeed an upper bound $\overline{s}$), our modelling specification incorporates the key distinction between the two types of buyers in a tractable way.[26]

Because credulous buyers base their willingness-to-pay on the seller's inflated claim, they purchase whenever

$$\int_B \max\left\{E[u|b], q\right\} \psi(b|\overline{s})db \geq p. \tag{13}$$

5.1 Market Outcome

Serving Both Wary and Credulous Customers. When the seller offers a contract (p, q) to attract both credulous and wary buyers, then the outcome is unaffected by

[26] This approach to model naivete in strategic information transmission games follows Kartik, Ottaviani, and Squintani (2007) and Bolton, Freixas, and Shapiro (2009).

the presence of credulous buyers. The firm's pricing power is constrained by the lower willingness-to-pay of wary buyers, which consequently will enter into the firm's objective function, implying that the firm's optimal offer solves the same trade-off as before, irrespective of α.

Proposition 6 *When the firm offers a contract (p, q) that is directed to both credulous and wary buyers, the outcome is identical to the one characterized by Proposition 4 and does not depend on the fraction α of credulous buyers.*

Note that when the price satisfies the wary customers' participation constraint in (9), credulous customers expect to realize a strictly positive expected utility according to (13), even though in reality their expected consumer surplus is equal to zero, given that the monopoly position enjoyed by the seller. The presence of wary customers thus protects credulous customers from exploitation.[27]

Serving only Credulous Customers. Suppose now the seller only targets credulous customers. Then, it is optimal for him to raise the price p so that he extracts, for a given refund level, their full willingness-to-pay, as given in (13). This offer is then indeed no longer acceptable to wary customers.

When targeting only credulous customers, the level of the refund no longer serves as a commitment device, which renders the seller's cheap talk costly, but instead becomes an instrument to extract more surplus from credulous buyers. Thus, we find that the refund is below the salvage value, $q < v$. The intuition for this result is that while the seller recommends a purchase to all $s \geq s^*$, credulous buyers believe that $s = \overline{s}$. As the probability with which the product is subsequently returned, $\Psi(b^*|s)$, is strictly decreasing in s, credulous buyers thus put a strictly lower value on an increase in the refund q than what such an increase actually costs to the seller.[28]

Proposition 7 *If the seller targets only credulous buyers, then $q < v$ and the seller advises all credulous buyers to purchase, irrespective of the observed signal s. Credulous buyers realize negative expected surplus.*

[27] This finding is somewhat similar to the general tenet in the literature dealing with search or shopping cost that the presence of customers with lower such costs brings down prices, which benefits all customers.

[28] The logic driving this result is similar to that of models of contracting with heterogeneous priors (e.g., Eliaz and Spiegler 2006).

Intuitively, which of the two cases of Propositions 6 and 7 apply depends on the composition of the market. Note that the seller's profits from serving both types of customers do not depend on α, whereas those from serving only credulous buyers clearly increase proportionally with α. Thus, there exits an interior cutoff for the fraction of credulous buyers α such that the seller targets only credulous buyers if and only if α does not fall below this threshold.

5.2 Consumer Protection

Proposition 6 and our preceding results imply that there is no scope for policy intervention when the seller finds it optimal to serve both types of buyers. This is, instead, no longer the case when the seller chooses to serve only credulous buyers.

When the seller only targets credulous buyers, then from Proposition 7 the seller offers an inefficiently low refund and, in addition, advises all credulous buyers to purchase. In this case, it is straightforward to show that efficiency can be increased by prescribing a minimum statutory requirement $q \geq \overline{q} > v$.[29] As the seller will still raise p until the credulous buyers' participation constraint (13) is satisfied, as long as the seller will still target only credulous buyers', their expected surplus will remain negative. Still, the higher is $\overline{q}$, the less they will be exploited, such that policy intervention now leads to strictly higher consumer surplus.[30]

What is more, through setting a minimum standard $\overline{q}$, policy intervention may be able to ensure that the seller no longer targets only credulous buyers. To see this, suppose for a moment that the second-best efficient refund level with wary buyers is unique and given by q_{SB}. Then by setting a minimum standard $\overline{q} \leq q_{SB}$ the seller's profits from serving all buyers are unaffected, while those from serving only credulous buyers are strictly lower. Moreover, while at $\overline{q} = q_{SB}$ the first-order effect on the former profits from an increase in $\overline{q}$ is still zero, it is clearly strictly negative for the latter profits with only credulous buyers. However, for still higher $\overline{q}$ it can no longer be ensured that a further increase in $\overline{q}$ reduces profits with only credulous buyers by more than those with all buyers.

[29] Note that the second-best efficient choice of $\overline{q}$ is, however, generally different from that characterized in Proposition 4. This follows as, for given q, the price that the seller can charge gullible buyers is strictly higher than the price that he can charge wary buyers, implying a strictly lower cutoff s^*.

[30] Note that we abstract from the possibility that policy intervention targets directly the firm's price p (e.g., through a cap).

Proposition 8 *In the presence of credulous buyers, the imposition of a minimum refund level $\overline{q} \leq q_{SB}$ is beneficial for the following reasons:*
i) By reducing the seller's profits only when he targets credulous buyers, a higher minimum refund level may induce the seller to instead serve all buyers, in which case social efficiency and consumer surplus are highest (with $q = q_{SB}$).
ii) If the seller still targets only credulous buyers, then an increase in $\overline{q}$ reduces the expected loss that credulous buyers incur. As long as $\overline{q} \leq v$, it also unambiguously increases welfare.

Note that when the imposition of a minimum standard induces the seller to serve all buyers (assertion i), then for this purpose it need not be binding. Sellers may then voluntarily offer a strictly higher refund (namely the second-best refund q_{SB}). Still, the minimum standard is effective because without it the seller would switch to offering a lower standard and only target credulous buyers.

For assertion ii) recall that without policy intervention the seller advises all credulous buyers to purchase. This still holds as long as $\overline{q} \leq v$ and $p > c$. The expected surplus is then from (11) strictly negative. Still, by imposing a higher standard $\overline{q} \leq v$ policy intervention ensures at least that *interim* efficiency improves.

Discriminating Offers. When the seller still serves both buyer groups, we presumed that this is done with a single, uniform offer (p, q). Even when the seller can not observe a buyer's characteristics, offering a menu of contracts to practice indirect price discrimination may be more profitable. Suppose thus that the seller's offer specifies (q_W, p_W) for wary buyers and a different contract (q_G, p_G) for credulous buyers. From our preceding observations we already know that in case the offer is separating, then incentive compatibility implies that $q_G < q_W$.

To be specific, we stipulate that the seller first provides advice and that the buyer subsequently chooses from the offered menu of contracts. Note also that an implicit assumption in this set-up is that credulous buyers do not learn about their own credulity when seeing the menu. That is, they do not ask why given the seller's advice other buyers may, instead, prefer the alternative option, (q_W, p_W). If this was not the case, then the seller would, as previously assumed, be still restricted to making a uniform offer.

For brevity's sake we do not spell out in detail the seller's program when designing a discriminatory pair of contracts.[31] Instead, for the purpose of the present discussion it is sufficient to note that, as can be easily shown, the offer has the following characteristics: i) Wary buyers realize zero surplus in expectation; ii) credulous buyers realize strictly negative surplus in expectation; iii) social surplus is strictly lower than in the case where a single (second-best efficient) offer is made to all buyers.

Policy intervention may now include the imposition of a nondiscriminatory requirement. Put differently, it may prescribe that all buyers subsequently have access to the most beneficial terms of return and cancellation that are offered to *any* buyers. As is immediate, provided that the seller still serves all buyers, the seller will then optimally offer the second-best efficient contract.

Proposition 9 *When the seller serves all buyers, then consumer surplus and social welfare are both strictly higher in case the seller is restricted to a uniform offer (e.g., as all buyers have the statutory right to return the product under the most beneficial terms that the seller offers to any buyer).*

Propositions 8 and 9 together thus provide a rationale for policy intervention in the presence of credulous buyers, advocating the possible use of a minimum statutory refund level together with a "non-discrimination" requirement.

6 Conclusion

Firms and their agents who sell to buyers who, at least when signing the contract, are still ill-informed about how the characteristics of a product or service will suit their particular needs, face a problem of credibility when they try to convince buyers that it is also in their own best interest to purchase. Such sale advice must remain cheap talk when buyers are wary of sellers' own motives and when unsuitable advice is not costly to the seller. We show that the sell can and want to induce ex post costs of unsuitable advice, and thus make the advice more credible, by granting buyers a generous right of refund or early cancellation. After a purchase is made but before the product is fully consumed or before the service agreement expires, buyers become familiar with the product and are in a

[31] The program is not standard to the extent that both offers (q_W, p_W) and (q_G, p_G) affect the seller's cutoff decision, s^*, at the advice stage. As the seller faces a commitment problem vis-à-vis wary customers, we find that both offers are distorted compared to the case in which either all customers are wary or in which the seller targets only credulous customers.

position to experiment and obtain additional information about the product's attributes. the costs associated to refunds that exceeds the seller's salvage value (or the margin lost from early cancellation) discipline the seller to initially advise on a purchase only if the seller is sufficiently confident about the product's suitability, given the seller's superior information at the point of sale.

A key insight of the baseline analysis with only wary buyers is that though the seller's optimal offer still leads to excessive purchases (*ex ante inefficiency)* and excessive returns (*interim inefficiency)*, it is still second-best efficient. Policy intervention that would prescribe a different refund policy would only reduce social welfare, while having no effect on consumer surplus. We showed, however, how social efficiency and consumer surplus both increase with more competition. The seller's lower margin, which is reduced by competition, dampens the incentives to provide unsuitable advice in order to increase sales.

A role for policy intervention emerges when a sufficiently large fraction of buyers is credulous, instead of wary, and thus believes in the seller's cheap talk. The seller is then tempted to either target only credulous buyers, who given their (on average) inflated expectations have a higher willingness-to-pay, or to make (self-selecting) discriminatory offers. In the offer that is targeted to credulous buyers, the terms of cancellation or refund no longer play a role as a commitment device, but they are instrumental so as to better exploit buyers' inflated beliefs. As a result, the prevailing refund is then below the (*interim*) efficient level. Credulous buyers can be protected by a combination of a minimum statutory refund, even though in equilibrium the seller will offer more beneficial terms such that it actually does not bind, and a non-discriminatory requirement, which allows all buyers to return the product under the most favorable conditions that the seller offered to any buyer.

The model and results are framed in terms of the contractually stipulated level of refund. An alternative contractual variable is the length of time over which buyers can return a product or cancel a contract without penalty. Extending this period allows buyers to obtain more precise information, while resulting in a deterioration of the salvage value of the product. Our analysis suggests that market contracts will stipulate the second-best efficient duration when buyers are wary, even in the absence of policy intervention. Firms would instead offer inefficiently short trial periods when targeting credulous buyers, so as to exploit these buyers' expectations inflated by the unsuitable advice provided.

Appendix: Proofs

Proof of Proposition 1. Note first that given the specifications on the informativeness of the signals at the boundaries, we have that $\pi(\underline{s}) = p - c + v - q$ and $\pi(\overline{s}) = p - c$. Moreover, note that $\pi(s)$ is strictly increasing and continuous when $q > v$ and constant when $q = v$. This implies that $\pi(s) > 0$ for all $s \in S$ when $q \geq v$ and $p > c$. When $q > v$ and $p > c$, existence of a strictly interior cutoff s^* follows then from FOSD of $\Psi(b^*|s)$ in s. For the comparative statics results, implicit differentiation of (7) and FOSD of $\Psi(b^*|s)$ give

$$\frac{ds^*}{dq} = \frac{1}{v-q} \frac{\Psi(b^*|s^*)}{\left.\frac{d\Psi(b^*|s)}{ds}\right|_{s=s^*}} > 0, \qquad \frac{ds^*}{dp} = -\frac{1}{v-q} \frac{1}{\left.\frac{d\Psi(b^*|s)}{ds}\right|_{s=s^*}} < 0,$$

as claimed. **Q.E.D.**

Proof of Proposition 2. For all $q > v$ we show first existence of a unique pair (p, s^*) such that (9) is satisfied with equality and (7) holds with $\underline{s} < s^* < \overline{s}$. The binding constraint (9) defines a continuous and strictly increasing mapping $\widetilde{p}(s^*)$, with $\widetilde{p}(\underline{s}) < c + q - v$ by (11) and $\widetilde{p}(s^*) = \overline{u} > c$ by (1). Restricting the domain to $p \in [c, c + q - v]$, (7) defines a continuous and strictly decreasing mapping $\widetilde{s}^*(p)$, with $\widetilde{s}^*(c) = \underline{s}$ and $\widetilde{s}^*(c + q - v) = \overline{s}$. These boundary conditions together with monotonicity and continuity of the two mappings guarantee existence and uniqueness of (p, s^*) with $\underline{s} < s^* < \overline{s}$. Finally, for the case with $p > c$ and $q \leq v$, we have that $s^* = \underline{s}$ by Proposition 1, so that this case is not consistent with the participation constraint (9) under condition (11). **Q.E.D.**

Proof of Proposition 3. For this and the following proofs it is convenient to write out the binding constraint (9) more explicitly as

$$\int_{s^*}^{\overline{s}} \left[\Psi(b^*|s)q + \int_{b^*}^{\overline{b}} E[u|b]\psi(b|s)db \right] \left(\frac{f_s(s)}{1 - F_s(s^*)} \right) ds = p. \tag{14}$$

Instead of using the implicit function theorem on the system of equations (7) and (14), it is more convenient to prove the result indirectly by arguing to a contradiction. Suppose thus that it holds that $ds^* < 0$ as $dq > 0$ at the margin. From total differentiation of (14) in (q, p, s^*, b^*), we have that

$$dp < dq \left[\int_{s^*}^{\overline{s}} \Psi(b^*|s) \left(\frac{f_s(s)}{1 - F_s(s^*)} \right) ds \right],$$

so that for *given* s^* we have from (7) that

$$d\pi(s^*) < dq \left[\int_{s^*}^{\overline{s}} \Psi(b^*|s) \left(\frac{f_s(s)}{1 - F_s(s^*)} \right) ds - \Psi(b^*|s^*) \right] - dq(v-q) \frac{\psi(b^*|s^*)}{dE[u|b^*]/db^*}, \qquad (15)$$

where we used implicit differentiation of $E[u|b^*] = q$ in (5). The right-hand side of (15) is strictly negative because $q > v$, $dE[u|b^*]/db^* > 0$ (by FOSD of $\Phi_b(u|b)$), and

$$\Psi(b^*|s^*) > \int_{s^*}^{\overline{s}} \Psi(b^*|s) \left(\frac{f_s(s)}{1 - F_s(s^*)} \right) ds$$

(by FOSD of $\Psi(b|s)$), implying from strict monotonicity of $\pi(.)$ that s^* must increase rather than decrease, as stipulated originally—a contradiction. **Q.E.D.**

Proof of Proposition 4. The result follows immediately from the fact that the seller wants to maximize *ex ante* surplus, as given by (8). To see this note first that $v > q$ holds at an optimal offer once we show that the resulting profits Π, which equal *ex ante* surplus, are then strictly positive (while from Proposition 2 profits are zero otherwise). From Proposition 3 and differentiability we have further that $ds^*/dq > 0$, while from (5) we have $db^*/dq > 0$. Hence, maximizing Π with respect to q, we have the derivative

$$\begin{aligned} \frac{d\Pi}{dq} &= \frac{ds^*}{dq} f_s(s^*) \left[\Psi(b^*|s^*)v + \int_{b^*}^{\overline{b}} E[u|b]\psi(b|s^*)db - c \right] \\ &\quad + \frac{db^*}{dq} \int_{s^*}^{\overline{s}} \psi(b^*|s) \left[v - E[u|b^*] \right] f_s(s)ds. \end{aligned} \qquad (16)$$

Given that $v < E[u|b^*]$ from $v < q$, the first-order condition $d\Pi/dq = 0$ requires that

$$\Psi(b^*|s^*)v + \int_{b^*}^{\overline{b}} E[u|b]\psi(b|s^*)db - c < 0,$$

which combined with (4) and FOSD of $\Psi(b|s)$ implies that $s^* < \widetilde{s}_{FB}(b^*)$. **Q.E.D.**

Proof of Proposition 7. The seller's profits when only serving credulous buyers are

$$\Pi_G := \alpha \int_{s^*}^{\overline{s}} \pi(s) f_s(s)ds,$$

where now p is substituted from (13). Note that p does not depend on s^*. After substitution, we obtain

$$\Pi_G = \alpha \int_{s^*}^{\overline{s}} \left[\int_{b^*}^{\overline{b}} E[u \mid b]\psi(b|\overline{s})db + \Psi(b^*; \overline{s})q + \Psi(b^*|s)(v-q) - c \right] f_s(s)ds.$$

At the optimal choice of q, it thus holds that

$$\int_{s^*}^{\overline{s}} \left[\Psi(b^*|s) - \Psi(b^*|\overline{s})\right] f_s(s) ds = -(v-q)\frac{db^*}{dq} \int_{s^*}^{\overline{s}} \left.\frac{d\Psi(b|s)}{db}\right|_{b=b^*} f_s(s) ds.$$

From FOSD of $\Psi(b|s)$ and $db^*/dq > 0$ it follows that $v > q$, from which we have in turn that $s^* = \underline{s}$. **Q.E.D.**

Proof of Proposition 8. Assertion i) follows from the argument in the main text. Next, for assertion ii) it only remains to show that credulous buyers' (true) expected surplus is a strictly increasing function of the refund q. Given that $s^* = \underline{s}$, a credulous buyer's expected loss is equal to

$$L := \int_{s^*}^{\overline{s}} \left[\left[\Psi(b^*|\overline{s}) - \Psi(b^*|s)\right] q + \int_{b^*}^{\overline{b}} E[u|b] \left[\psi(b|\overline{s}) - \psi(b|s)\right] db \right] f_s(s) ds,$$

so that

$$\frac{dL}{dq} = \int_{s^*}^{\overline{s}} \left[\Psi(b^*|\overline{s}) - \Psi(b^*|s)\right] f_s(s) ds < 0$$

follows from FOSD of $\Psi(b^*|s)$. **Q.E.D.**

7 References

Bolton, Patrick, Xavier Freixas, and Joel Shapiro. 2009. "The Credit Rating Game." NBER Working Paper 14712.

Che, Yeon-Koo. 1996. "Customer Return Policies for Experience Goods." *Journal of Industrial Economics*, 44(1), 17–24.

Courty, Pascal, and Hao Li. 2000. "Sequential Screening." *Review of Economic Studies*, 67(4), 697–718.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica*, 50(6), 1431–1451.

Davis, Scott, Eitan Gerstner, and Michael Hagerty. 1995. "Money Back Guarantees in Retailing: Matching Products to Consumer Tastes." *Journal of Retailing*, 71(1), 7–22.

Eliaz, Kfir, and Ran Spiegler. 2006. "Contracting with Diversely Naive Agents." *Review of Economic Studies*, 73(3), 689–714.

Green, Jerry R., and Nancy L. Stokey. 2007. "A Two-Person Game of Information Transmission." *Journal of Economic Theory*, 135(1), 90–104.

Hendel, Igal, and Alessandro Lizzeri. 2002. "The Role of Leasing under Adverse Selection." *Journal of Political Economy*, 110(1),113–143.

Inderst, Roman, and Marco Ottaviani. Forthcoming. "Misselling through Agents." *American Economic Review*.

Inderst, Roman, and Gilat Tirosh. 2009. "Rotating Demand: A Theory of Refunds and Warranties." Mimeo, University of Frankfurt.

Johnson, Justin P. and David P. Myatt. 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review*, 96(3), 756–784.

Johnson, Justin P. and Michael Waldman. 2003. "Leasing, Lemons, and Buybacks." *RAND Journal of Economics*, 34(2), 247–265.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani. 2007. "Credulity, Lies, and Costly Talk." *Journal of Economics Theory*, 134(1), 93–116.

Lewis, Tracy R., and David E.M. Sappington. 1994. "Supplying Information to Facilitate Price Discrimination." *International Economic Review*, 35 (2), 309–327.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics*, 118(4), 1209–1248.

Matthews, Steven A., and Nicola Persico. 2007. "Information Acquisition and Refunds for Returns." Mimeo, University of Pennsylvania.

Office of Fair Trading. 2004. "Doorstep Selling." http://www.oft.gov.uk/shared_oft /reports/consumer_protection/oft716.pdf

Pitchik, Carolyn, and Andrew Schotter. 1987. "Honesty in a Model of Strategic Information Transmission." *American Economic Review*, 77(5), 1032–1036.

Spence, Michael. 1977. "Consumer Misperceptions, Product Failure and Producer Liability." *Review of Economic Studies*, 44(3), 561–572.