

Collegio Carlo Alberto



Robust Pricing with Refunds

Toomas Hinnosaar and Keiichi Kawai

No. 563

December 2018

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Robust Pricing with Refunds^{*}

Toomas Hinnosaar[†]

Keiichi Kawai[‡]

This version: January 19, 2019

First Draft: April 18, 2018

Abstract

We characterize a selling mechanism that is robust to the seller's uncertainty about the buyer's signal structure. We show that by offering a generous refund policy, the seller can significantly reduce this type of uncertainty and regain market power. A simple mechanism that combines a generous refund policy and random non-refundable discounts achieves the best guaranteed-profit among all possible mechanisms.

JEL: D82, C79, D42

Keywords: optimal pricing, robustness, return policies, refunds, monopoly, information design, mechanism design

1 Introduction

We analyze a seller's robust pricing problem in the face of uncertainty about the buyer's information and learning. Consumers learn about product characteristics from various sources, and firms are often not only uncertain about how much value buyers receive from consuming their products but also how well buyers know whether these products fit their needs. For example, an online shoe store does not know whether a shopper has been "showrooming" and already knows whether particular shoes fit well. Similarly, a booking agent does not know whether a consumer planning to book a flight has a specific travel itinerary in mind or not. Such uncertainty regarding the buyer's information and learning can severely limit the seller's ability to extract profits from trade.

What profit can the seller guarantee himself in the presence of such uncertainty? How should the seller set the price to achieve the best guaranteed-profit? We show that a simple

^{*}We are grateful to Pak Hung Au, Gary Biglaiser, Ralph Boleslavsky, Laura Doval, Juan Dubra, Jeff Ely, Dino Gerardi, Marit Hinnosaar, Johannes Hörner, Atsushi Kajii, Teddy Kim, Anton Kolotilin, Fei Li, Hongyi Li, Giorgio Martini, Ignacio Monzón, Alessandro Pavan, Karl Schlag, Roland Strausz, Balázs Szentes, and Alexander Wolitzky for the useful discussions and valuable comments. We would also like to thank seminar participants at SAET (Taipei), the International Conference on Economic Theory and Applications (Chengdu), the International Conference on Game Theory (Stony Brook), Midwest Economic Theory Conference (Nashville), Kyoto, Duke-UNC, and Miami

[†]Collegio Carlo Alberto, toomas@hinnosaar.net.

[‡]UNSW, Sydney, k.kawai@unsw.edu.au.

mechanism that combines a generous refund policy with random non-refundable offers is the answer.

More specifically, we analyze a bilateral trade model, where the buyer's valuation of the product is either high or low, and the exact value is unknown to everyone including the buyer. The seller and the buyer share a common prior about the buyer's valuation. The buyer observes a private signal of her valuation but learns her valuation only after purchasing the product. The seller can allow the buyer to return the product but incurs a restocking cost for each returned product.

The key feature of the model is the seller's uncertainty about the buyer's information structure (i.e., the distribution of signals). The seller neither knows the buyer's signal distribution nor has a prior belief about possible signal distributions. He only knows that the signals satisfy Bayes' rule.¹ For example, the support of the signal distribution may only consist of moderate signals, i.e., the buyer's information structure may be quite uninformative. Then a sufficiently high price would result in no sales. Alternatively, the buyer's signal distribution may generate either a sufficiently favorable or unfavorable signal, i.e., the buyer's information structure may be quite informative. Then, setting a low price would leave a lot of money on the table.

Our goal is to identify the seller's selling mechanism that performs as well as possible regardless of the buyer's signal distribution. We show that a simple mechanism that randomizes between an offer with a generous refund and log-uniformly distributed non-refundable discounted offers achieves the seller's best guaranteed-profit, i.e., the profit the seller can guarantee irrespective of the buyer's signal distribution.

To understand the intuition behind our result, we first note that the seller can reduce the importance of the buyer's signal in the purchasing decision by offering a generous refund. That is, for any given buyer's signal distribution, the buyer is more likely to buy the product at the same price with a refund than without. By making the buyer's demand less elastic this way, a generous refund enables the seller to charge a higher price without sacrificing the probability of trade. The seller, however, needs to be wary of costs associated with product return. For each returned product, the seller incurs the restocking cost. Thus, the seller must ensure that if an offer is refundable, it is only attractive to the buyer who is unlikely to return the product (i.e., the one with a sufficiently favorable signal). By doing so, the seller can ensure that the buyer who accepts the refundable offer always brings a positive profit in expectation. The seller achieves this goal by offering a generous but not full refund.

We also note that, loosely speaking, any pricing policy that brings a large profit when the buyer is relatively well-informed tends to perform poorly when the buyer is relatively uninformed, and vice versa. Therefore, to increase his guaranteed profit, the seller needs to hedge simultaneously against the buyer's signal distributions that are relatively informative and relatively uninformative.

With these observations at hand, we now explain why the seller can achieve the best guaranteed-profit through the mixture of a generous refund (that is only attractive to the buyer with a sufficiently favorable signal) and random non-refundable offers (that are attractive to the buyer only when her signal is moderate or favorable). The generous refund brings a large profit if and only if the buyer's signal distribution is likely to gener-

¹The environment where the seller chooses a price-refund pair without knowing the realized signal drawn from a commonly-known distribution is analyzed in Inderst and Tirosh (2015) and Kraemer and Strausz (2015).

ate sufficiently favorable signals. A non-refundable offer brings a large profit if and only if the buyer's signal distribution is likely to generate a signal that is equal to or slightly above the price offered by the seller. Furthermore, for any buyer's signal distribution under which the refundable offer brings a small profit, the randomized non-refundable offer brings a sufficiently large profit and vice versa.² In this sense, the refundable offer hedges against the event that the signal distribution is informative, and the randomization over non-refundable offers hedges against the less informative distributions. This observation also explains why the seller who is interested in maximizing the guaranteed-profit would find it optimal to always use a generous refundable offer with a positive probability no matter how high the restocking cost to the seller is.

We then use our findings to derive the buyer-optimal information structure, as well as the sharp upper-bound of the buyer's payoff with respect to the possible information structure. More specifically, we analyze the buyer-optimal information design problem, where the buyer first chooses an information structure, and the seller chooses an optimal pricing and refund policy knowing the buyer's information structure.³ We show that the buyer-optimal signal distribution is the worst-possible distribution to the seller (i.e., the distribution for which the seller's maximal profit is exactly the best guaranteed-profit). This finding implies that if we compare the highest payoffs the buyer can obtain (by choosing her information structure) when the seller can offer a refund and cannot offer a refund, then the former is strictly lower than the latter. In this sense, a refund policy increases the seller's market power against buyers who search information strategically.

Our findings are closely related to the results in Roesler and Szentes (2017) and Du (2018). Roesler and Szentes (2017) identify the information structure that maximizes the buyer's welfare when the seller best responds to the information structure via uniform pricing. Du (2018) shows that the information structure found in Roesler and Szentes (2017) indeed minimizes the profit the seller can obtain, and the seller can obtain the best guaranteed-profit by what he calls an exponential pricing regardless of the buyer's signal distribution.⁴

In both of these papers, the interaction between the seller and the buyer ends with the purchase decision. In contrast to these papers, we allow the interaction to continue after the purchase, i.e. after the buyer learns the product fit. With a refund policy, the seller can potentially increase his profit through two channels. The seller can indirectly control the buyer's learning by incentivizing the buyer to learn through purchase. The seller also can sequentially screen the buyer, first by the signal distribution and the signal; and then by the realized valuation for the product.⁵ Therefore, the selling mechanism that maximizes

²The Bayes-plausibility requires that the expected value of the posteriors induced by the buyer's signal distribution must be equal to the prior. Loosely speaking, this implies that when the buyer's signal distribution is likely to generate a sufficiently favorable signal, it also is likely to generate a sufficiently unfavorable signal.

³For example, the buyer may delegate the information gathering to a third party such as an algorithm or an employee to commit to a certain information structure.

⁴Libgober and Mu (2017) analyze a robust dynamic pricing problem where the product is durable and buyers learn about their value for the product over time.

⁵The literature on sequential screening and dynamic mechanism design has identified why and how advance sales to still-uninformed consumers can help the seller. See e.g., Gale and Holmes (1992, 1993); Courty and Li (2000); Eső and Szentes (2007); Nocke et al. (2011); Gallego and Sahin (2010); Ely et al. (2017). The closest to our paper is von Wangenheim (2017), who studies a model where the buyer learns the value over time and the seller can offer flexible contracts, which can be interpreted as a (costless) refund policy. He finds

the seller’s best guaranteed-profit may be complex. Nevertheless, we obtain results parallel to Roesler and Szentes (2017) and Du (2018): a generalization of the signal distribution identified in Roesler and Szentes (2017) is the worst possible buyer’s signal distribution for the seller; and an extension of the exponential pricing found in Du (2018) achieves the seller’s best guaranteed-profit against any possible buyer’s signal distributions.

Our findings offer a novel rationale for generous return policies. The literature has identified various reasons that companies may use return policies: e.g., as costly signals for product quality and product fit for the consumer (Grossman (1981); Moorthy and Srinivasan (1995); Inderst and Ottaviani (2013)); as insurance for risk-averse consumers (Che (1996)); and as a tool for price discrimination (Zhang (2013); Escobari and Jindapon (2014); Inderst and Tirosch (2015)).⁶ Among these, the closest to our paper is Inderst and Tirosch (2015). In an environment where the seller knows the buyer’s signal distribution, Inderst and Tirosch (2015) show that return policies work as “metering devices”, where refunds make different consumers more similar and thus allow the firm to capture more of the surplus by raising prices.⁷ Consequently, the seller sets the refund amount above the restocking cost. In contrast, our results show that the seller offers a generous refund (i.e., “almost” full refund) when the seller is uncertain of the buyer’s signal distribution.

Lastly, we note that the model analyzed in the present paper can be interpreted as a game between the seller, who aims to maximize his profit by indirectly controlling the buyer’s learning on product fit through the design of a price-refund pair, and the adversarial nature, who aims to minimize the seller’s profit by directly choosing the buyer’s signal distribution. Thus, the game is akin to Bayesian persuasion games with competing information designers. We therefore fully utilize the concavification technique (Aumann et al. (1995); Kamenica and Gentzkow (2011)); and the properties of equilibrium payoff functions in the competitive Bayesian-persuasion settings (Boleslavsky and Cotton (2018); Au and Kawai (2017a,b)) to derive our results.

2 Model

There is a (male) seller of a product, and a (female) buyer whose valuation for the product is $v \in \{v_l, v_h\}$, where $0 < v_l < v_h$. The buyer’s valuation v follows a commonly known distribution such that $\pi = \Pr(v = v_h)$. Neither the seller nor the buyer knows the realization of v . However, the buyer receives a signal of her valuation v prior to purchase. (Details will be explained below.) The seller can sell the product at price v_l to a third-party not specified within the model.⁸

We start our analysis with the case in which the seller’s (pure) strategy is a contract (p, r) that specifies a price p together with a refund r . The seller may use a mixed strategy $\Delta\{(p, r)\}$. We call the seller’s mixed strategy, i.e., a distribution over contracts, a *policy*; and the realized contract an *offer*. Based on the information she has, the buyer decides

that the seller obtains the static monopoly profit, which means capturing full surplus in our framework.

⁶Escobari and Jindapon (2014) also provide some empirical evidence on the use of refundable tickets by airlines. They show that a fully refundable ticket is typically about 50% more expensive than a non-refundable ticket, but the difference disappears in the last week before the departure. These facts fit well with our model predictions.

⁷Similar ideas have been studied in other contexts, such as overbooking by airlines, e.g., Ely et al. (2017).

⁸That is, we assume that the trade between the seller and the buyer is socially efficient.

whether to buy after observing the contract (p, r) . We analyze the more general setting in which the seller can offer any mechanism in Section 3.4.

If the buyer purchases the product, then she learns the realized value of v . If the realized offer is non-refundable (i.e. $r = 0$), then the game ends. If $r > 0$, then the buyer decides whether or not to return the product. When the product is returned, the value of the seller's outside option decreases by c . The term c , which we call the *restocking cost*, encapsulates the losses the seller needs to absorb when the product is returned.⁹ In the analysis below, instead of in terms of c , we state our results in terms of *the normalized restocking cost* $\gamma \equiv \frac{c}{v_h - v_l} / (1 + \frac{c}{v_h - v_l})$ for conciseness.¹⁰

If the buyer purchases and keeps the product, then her payoff is $v - p$ and the seller's profit is p . If the realized offer is refundable (i.e., $r > 0$) and she returns the product, then her payoff is $r - p$, and the seller's profit is $p - r + (v_l - c)$. If the buyer does not buy the product, then her payoff and the seller's profit are zero and v_l , respectively.

We can represent a buyer's signal as a posterior $q = \Pr(v = v_h)$ that is a random variable drawn from a distribution function $F \in \mathcal{F} \equiv \{F : \mathbb{E}_F[q] = \pi\}$.¹¹ For this reason, we use a distribution F over posteriors to represent the buyer's information structure, and call it a *signal distribution*. Analogously, by a *signal* q , we refer to the realization of the posterior from signal distribution F .

Take an arbitrary buyer's signal distribution F . We use $V((p, r) | F)$ to represent the seller's expected profit when the offer is (p, r) , and $\mathbb{E}_{\Delta\{(p, r)\}} V((p, r) | F)$ to represent the expected profit from policy $\Delta\{(p, r)\}$. We say policy $\Delta\{(p, r)\}$ *guarantees profit* V when

$$\mathbb{E}_{\Delta\{(p, r)\}} V((p, r) | F) \geq V \text{ for all } F \in \mathcal{F},$$

i.e., $\min_{F \in \mathcal{F}} \mathbb{E}_{\Delta\{(p, r)\}} V((p, r) | F) \geq V$. Our goal is to identify a strategy $\Delta\{(p, r)\}$ that brings the best guaranteed-profit defined by

$$V^* \equiv \sup_{\Delta\{(p, r)\}} \min_{F \in \mathcal{F}} \mathbb{E}_{\Delta\{(p, r)\}} V((p, r) | F).$$

The best guaranteed-profit is the supremum of the profits the seller can guarantee in the following game against a fictitious player called the (adversarial) nature whose objective is to minimize the seller's expected profit:

1. The seller chooses a policy $\Delta\{(p, r)\}$.
2. Nature chooses a signal distribution F after observing the seller's choice of $\Delta\{(p, r)\}$.
3. The buyer observes a signal $q \sim F$ and an offer $(p, r) \sim \Delta\{(p, r)\}$.

⁹If the seller wants to sell the returned product to the third party, he incurs the cost associated with processing returns, repackaging and restocking the merchandise, testing the products, and having to reduce the price from v_l for no longer being able to sell as new. Even if the seller is better off not selling the returned product to the third-party, he may still need to pay fees associated with credit card refund transactions, e.g., the provider may keep the interchange fee and/or charge a transaction fee, as well as shipping fees. Therefore, we do not exclude the cases where the "net salvage value" $v_l - c$ being negative.

¹⁰Observe that the normalized restocking cost γ takes a value in $(0, 1)$, strictly increasing in $\frac{c}{v_h - v_l}$, and $\gamma \rightarrow 1$ as $\frac{c}{v_h - v_l} \rightarrow \infty$. That is, we can capture possible values of $\frac{c}{v_h - v_l} \in (0, \infty)$ in terms of $\gamma \in (0, 1)$.

¹¹Notice that $F \in \mathcal{F}$ if and only if $\int_0^1 F(q) dq = 1 - \pi$.

4. The buyer decides whether or not to buy. If she does not buy, the game ends.
5. The buyer learns the value v if she buys. If $r = 0$, the game ends.
6. If $r > 0$, the buyer decides whether or not to return the product.
7. If the buyer returns the product, the seller refunds r to the buyer, and incurs the restocking cost c .

Without loss of generality, we normalize the monetary units so that $v_l = 0$ and $v_h = 1$. More precisely, when $(v_l, v_h, c) = (0, 1, \tilde{c})$, suppose that the best guaranteed-profit and the offer that achieves it are V^* and (p, r) , respectively. If $(v_l, v_h, c) = (\tilde{v}_l, \tilde{v}_h, (\tilde{v}_h - \tilde{v}_l) \tilde{c})$, then the offer $(f(p), \mathbb{1}_{r > v_l} f(r))$ achieves the best guaranteed-profit $f(V^*)$, where $f(x) \equiv \tilde{v}_l + x(\tilde{v}_h - \tilde{v}_l)$.¹²

Note that if the normalized restocking cost γ is zero, i.e., the restocking cost c is zero, then the seller can fully control the buyer's learning at no cost.¹³ Consequently, the seller can capture the full surplus from trade π regardless of the buyer's information structure. In the following, we study a more realistic case where the restocking cost is strictly positive.

3 Best Guaranteed-Profit

We characterize the best guaranteed-profit V^* by identifying its lower and upper bounds. We first identify the seller's best guaranteed-profit from a deterministic policy (Lemma 1). That is, we identify

$$\underline{V}^* \equiv \sup_{(p,r)} \min_{F \in \mathcal{F}} V((p, r) | F).$$

Next we derive a worst distribution $F_{\bar{V}^*}$ for the seller by analyzing an auxiliary game in which the seller chooses his pricing policy after nature chooses the buyer's signal distribution (Lemma 3). That is, we derive

$$F_{\bar{V}^*} \in \arg \min_{F \in \mathcal{F}} \sup_{(p,r)} V((p, r) | F) \text{ and } \bar{V}^* \equiv \sup_{(p,r)} V((p, r) | F_{\bar{V}^*}).$$

While $\underline{V}^* \leq V^* \leq \bar{V}^*$ by construction, there exists a threshold value $\bar{\gamma}$ such that if the (normalized) restocking cost $\gamma \in (0, \bar{\gamma}]$ then $\underline{V}^* = \bar{V}^*$ and therefore the best guaranteed-profit $V^* = \underline{V}^* = \bar{V}^*$ can be achieved by a deterministic policy. In contrast, if the normalized restocking cost is higher, i.e. $\gamma \in (\bar{\gamma}, 1)$, then $\underline{V}^* < \bar{V}^*$ and therefore there is room for improvement from a randomized policy. In this case we show that a randomization over a refundable offer with a generous refund and log-uniformly distributed non-refundable offers achieves the upper bound \bar{V}^* regardless of the signal structure F (Theorem 1). Lastly, we show that there exists no mechanism that guarantees the seller a higher profit than $V^* = \bar{V}^*$ (Theorem 2).

¹²For concreteness, observe that the normalized net salvage value $v_l - c$ is necessarily negative. This simply means that product return is costly to the seller. See also footnote 9.

¹³For any $\varepsilon > 0$, the seller can guarantee himself $\pi - \varepsilon$ by offering $(p, r) = (1, 1 - \frac{\varepsilon}{1-\pi})$, i.e., the seller's best guaranteed-profit is π .

3.1 Best Guaranteed-Profit with Deterministic Policies

We start our analysis by identifying the best guaranteed-profit by a deterministic policy. That is, we characterize $\underline{V}^* \equiv \sup_{(p,r)} \min_{F \in \mathcal{F}} V((p,r)|F)$, which is a lower bound for V^* .

Let $v(q|(p,r))$ denote the seller's expected profit when the offer is (p,r) and the buyer's signal is q . Then, $V((p,r)|F) = \mathbb{E}_F[v(q|(p,r))]$. Also for notational simplicity, we use $\underline{V}(p,r)$ to denote its minimum over F , i.e., $\underline{V}(p,r) \equiv \min_{F \in \mathcal{F}} V((p,r)|F)$, so the best guaranteed-profit with deterministic offers is $\underline{V}^* = \sup_{(p,r)} \underline{V}(p,r)$.

Suppose that the seller makes a non-refundable offer $(p,0)$. Then, the buyer with signal q buys if and only if $q \geq p$, i.e.,

$$v(q|(p,0)) = \begin{cases} 0 & \text{if } q \leq p, \\ p & \text{if } q > p. \end{cases}$$

The seller's profit $V((p,r)|F)$ is minimized when the probability of the buyer's signal being larger than p is minimized, i.e., when F minimizes $1 - F(p)$ subject to $\mathbb{E}_F[q] = \pi$. If $\pi > p$, then this occurs when the buyer's signal distribution induces two signals p (which results in no trade) and 1 (which results in trade) with probabilities $\frac{1-\pi}{1-p}$, and $\frac{\pi-p}{1-p}$, respectively. In contrast, if $\pi \leq p$, then this occurs when the buyer's signal does not disclose any additional information, i.e., induces signal π (which results in no trade) with probability one.




Figure 1: Profit of a non-refundable offer $(p, 0)$

More formally, deriving $\underline{V}(p, 0)$ is a canonical Bayesian persuasion problem in which the information designer (nature) chooses the distribution over signals F to maximize $-v(\cdot|(p, 0))$ subject to $F \in \mathcal{F}$. We thus can derive $\underline{V}(p, 0)$ by utilizing the concavification approach.¹⁴ Let $\text{con}[-v(\cdot|(p, 0))](q)$ be the value of concave closure of $-v(\cdot|(p, 0))$ at q , which is represented by the red-dotted line in Figure 1.¹⁵ By construction $\underline{V}(p, 0) = -\text{con}[-v(\cdot|(p, 0))](\pi)$. While a higher p results in a higher profit margin should trade occur, it leads to a lower probability of trade. The seller balances this trade-off by offering

¹⁴See Aumann et al. (1995) and Kamenica and Gentzkow (2011).

¹⁵The concave closure of function G is defined by $\text{con}[G](q) = \sup\{g|(q, g) \in \text{co}(G)\}$, where $\text{co}(G)$ is the convex hull of the graph of G .

$p = 1 - \sqrt{1 - \pi}$, so

$$(1) \quad \begin{aligned} \underline{V}(p, 0) &= -\text{con}[-v(\cdot | (p, 0))](\pi) = \begin{cases} 0 & \text{if } p \geq \pi \\ p \times \frac{\pi - p}{1 - p} & \text{if } p < \pi \end{cases} \\ &\leq \sup_p \underline{V}(p, 0) = \sup_{p \in [0, \pi]} p \frac{\pi - p}{1 - p} = (1 - \sqrt{1 - \pi})^2. \end{aligned}$$

Next, consider a refundable offer, i.e., (p, r) such that $r > 0$. Without loss of generality, we only consider the case where $p \geq r$.¹⁶ The payoff of the buyer with signal q from buying is $q \times 1 + (1 - q) \times r - p$. She thus buys only if her signal is above the marginal signal $\bar{q}(p, r) \equiv \frac{p-r}{1-r}$; and returns with probability $1 - q$ if she buys. The seller's profit from the buyer with signal q is thus

$$(2) \quad v(q | (p, r)) = \begin{cases} 0 & \text{if } q < \bar{q}(p, r), \\ \min\{0, \bar{v}(q; p, r)\} & \text{if } q = \bar{q}(p, r), \\ \bar{v}(q; p, r) & \text{if } q > \bar{q}(p, r), \end{cases}$$

where $\bar{v}(q; p, r) \equiv p - (1 - q)(c + r) = p - (1 - q)\left(\frac{\gamma}{1 - \gamma} + r\right)$.




Figure 2: Profit from a refundable offer (p, r)

Observe that if the marginal signal $\bar{q}(p, r)$ is low, i.e., if the refund is generous, then the buyer buys even when her signal q is low. The buyer with a low signal is likely to return the product, and thereby is likely to bring the seller an (ex-post) negative profit. More specifically, as captured by the blue-lines in Figures 2(a) to 2(c), the seller's profit from the buyer with signal q conditional on sales $\bar{v}(q; p, r)$ is increasing in q ; and the profit from the buyer with signal γ is negative, i.e., $\bar{v}(\gamma; p, r) < 0$ if and only if $\bar{q}(p, r) < \gamma$.

Therefore, if $\bar{q}(p, r) < \gamma$, then the seller's profit is minimized when the probability of the seller receiving a signal $\bar{q}(p, r)$ is maximized, as captured by the red dotted-line in Figure 2(a). The seller then can improve his guaranteed-profit by offering a less generous refund, and thereby increasing the marginal signal.

¹⁶Notice that if $p < r$, then the buyer buys irrespective of the value of the signal. Therefore, $\underline{V}(p, r) < \underline{V}(p, r - \varepsilon)$ for a sufficiently small $\varepsilon > 0$.

In contrast, if the marginal signal is sufficiently high, i.e., $\tilde{q}(p, r) > \gamma$, then the buyer who buys the product always brings a positive (ex-post) profit to the seller. The seller's profit is thus minimized when the probability of the seller receiving a signal above $\tilde{q}(p, r)$ is minimized, as captured by the red dotted line in Figure 2(b). The seller then can improve his guaranteed-profit by offering a more generous refund, and thereby by lowering the marginal signal.

Combining these observations, we conclude that if the seller were to offer a positive refund, then he would set the price p as close to one as possible, and r to $\frac{p-\gamma}{1-\gamma}$ so that the marginal signal is exactly at $\tilde{q}(p, r) = \gamma$. By doing so, the seller can guarantee himself $\frac{\pi-\gamma}{1-\gamma}$. Therefore, when $r > 0$,

$$\underline{V}(p, r) = -\text{con}[-v(\cdot | (p, r))] (\pi) \leq \sup_{p, r > 0} \underline{V}(p, r) = \max \left\{ \frac{\pi - \gamma}{1 - \gamma}, 0 \right\}.$$

Recall that the seller's best guaranteed-profit without any refund is $(1 - \sqrt{1 - \pi})^2$, as derived in (1). Thus, the seller's best guaranteed-profit with a deterministic policy is

$$(3) \quad \underline{V}^* = \sup_{p, r} \underline{V}(p, r) = \begin{cases} \frac{\pi - \gamma}{1 - \gamma} & \text{if } \gamma \leq \hat{\gamma}(\pi), \\ (1 - \sqrt{1 - \pi})^2 & \text{if } \gamma > \hat{\gamma}(\pi), \end{cases}$$

where $\hat{\gamma}(\pi) \equiv \frac{2(1 - \sqrt{1 - \pi})}{2 - \sqrt{1 - \pi}}$. The resulting best guaranteed-profit \underline{V}^* is depicted in Figure 3.

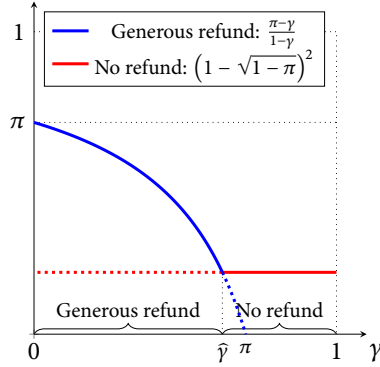


Figure 3: Best guaranteed-profit \underline{V}^* with a deterministic offer

Lemma 1. *Suppose that the seller can only use a deterministic pricing policy (p, r) . Then the seller's best guaranteed-profit is \underline{V}^* defined in (3). For any $\varepsilon > 0$, either a generous refund $(p, r) = \left(1 - \varepsilon, 1 - \frac{\varepsilon}{1-\gamma}\right)$ (when $\gamma \leq \hat{\gamma}(\pi)$), or no refund $(p, r) = (1 - \sqrt{1 - \pi}, 0)$ (when $\gamma \geq \hat{\gamma}(\pi)$) guarantees profit of $(1 - \varepsilon) \underline{V}^*$.*

3.2 Best Guaranteed-Profit with Policy-Independent Signals

We now consider an auxiliary game in which the seller chooses his pricing policy $\Delta\{(p, r)\}$ after observing the nature's choice of the buyer's signal distribution F . By analyzing this game, we identify a *worst distribution for the seller* $F_{\underline{V}^*}$, i.e.,

$$F_{\underline{V}^*} \in \arg \min_{F \in \mathcal{F}} \sup_{(p, r)} V((p, r) | F).$$

Then $\bar{V}^* \equiv \sup_{(p,r)} V((p,r) | F_{\bar{V}^*})$ defines an upper bound of V^* .

As we discussed in Section 3.1, for a given offer (p, r) , the marginal signal is $\tilde{q}(p, r) = \frac{p-r}{1-r}$. Therefore, when the buyer's signal distribution is F , the seller's profit from an offer (p, r) is

$$V((p, r) | F) = p(1 - F(\tilde{q}(p, r))) - \mathbf{1}_{r>0}(c + r) \int_{\tilde{q}(p, r)}^1 (1 - q) dF(q).$$

To facilitate the further discussion, for any small ε , we say $(p, r) = \left(1 - \varepsilon, 1 - \frac{\varepsilon}{1-\gamma}\right)$ is an offer with a generous refund. Notice that the marginal signal for an offer with a generous refund is γ , and hence the seller's profit from a generous refund is $(1 - \varepsilon) \int_{\gamma}^1 \frac{q-\gamma}{1-\gamma} dF(q) \geq 0$. Also, when the buyer's signal distribution is F , by the profit from a generous refund, which we denote by $V_R(F)$, we refer to the supremum of the profits the seller can achieve by some offer with a generous refund $V_R(F)$. That is,

$$(4) \quad V_R(F) \equiv \sup_{\varepsilon} V\left(\left(1 - \varepsilon, 1 - \frac{\varepsilon}{1-\gamma}\right) \middle| F\right) = \int_{\gamma}^1 \frac{q-\gamma}{1-\gamma} dF(q).$$

We derive $F_{\bar{V}^*}$ in the following steps. We first characterize the seller's highest profit from a non-refundable offer for a given buyer's signal distribution F , which we denote by $V_{NR}(F)$. We also identify a lower bound of the seller's profit from a generous refund when his profit from a non-refundable offer is V , which we denote by $\Phi(V)$. Lastly, we show that a distribution F that minimizes $\max\{\Phi(V_{NR}(F)), V_{NR}(F)\}$ is a worst distribution $F_{\bar{V}^*}$ for the seller, and $\bar{V}^* = \Phi(V_{NR}(F_{\bar{V}^*}))$.

Take an arbitrary signal distribution F . By applying an argument similar to Roesler and Szentes (2017), we find that the profit the seller can achieve by a non-refundable offer is

$$V_{NR}(F) \equiv \inf\{V : F(q) \geq G_V(q) \text{ for all } q\},$$

where

$$(5) \quad G_V(q) \equiv \begin{cases} 0 & \text{if } q \in [0, V), \\ 1 - \frac{V}{q} & \text{if } q \in [V, 1), \\ 1 & \text{if } q = 1. \end{cases}$$

To see this, notice that if the buyer's distribution is G_V , then for any non-refundable offer $(p, 0)$, $p \geq V$, the seller's profit is V because the buyer's demand is unit elastic. Therefore, if $F(q) = G_V(q)$ for some V and $q \geq V$, then the seller's profit from a non-refundable offer is V .¹⁷

Next, take an arbitrary buyer's distribution F such that $V_{NR}(F) = V$. We derive a lower-bound of the seller's profit from a generous refund, which we denote by $\Phi(V)$. We note that, by (4), the seller's profit from a generous refund is equal to the case where the buyer's signal distribution is $\tilde{F} \in \mathcal{F}$ such that $\tilde{F}(q) = F(q)$ for all $q \in [0, \gamma)$, and $\tilde{F}(q) = \tilde{F}(\gamma)$ for

¹⁷More formally, the supremum of the seller's profit from a non-refundable offer, i.e., $\sup_p V((p, 0) | F)$, is $\sup_p p(1 - F(p) + \Delta F(p))$, where $\Delta F(p)$ is the size of atom of F at p . Furthermore, since $F(q) \geq G_{V_{NR}(F)}(q)$ for all q , we have $p(1 - F(p) + \Delta F(p)) \leq p(1 - G_{V_{NR}(F)}(p))$ for all p . Since $p(1 - G_{V_{NR}(F)}(p)) = V_{NR}(F)$ for all $p \in [V_{NR}(F), 1]$, we have $\sup_p V((p, 0) | F) = V_{NR}(F)$.

all $q \in [\gamma, 1]$. Therefore, if the buyer's signal distribution is F , then the seller's profit from a generous refund is $1 - \int_{\gamma}^1 \frac{F(q)}{1-\gamma} dq = 1 - \bar{F}(\gamma)$.

The nature thus can lower the seller's profit from a generous refund (while keeping the highest profit from a non-refundable offer at V) by making the area under $F(q)$ on $[\gamma, 1]$, i.e., $\int_{\gamma}^1 F(q) dq$, as large as possible subject to the constraints $\int_0^1 F(q) dq = 1 - \pi$ and $F(q) \geq G_V(q)$ for all q . With this observation at hand, define

$$\Phi(V) \equiv \begin{cases} \frac{V(\ln \frac{V}{\gamma} - 1) + \pi}{1-\gamma} & \text{if } V < \gamma, \\ \frac{\pi - \gamma}{1-\gamma} & \text{if } V \geq \gamma. \end{cases}$$

so that $\int_0^{\gamma} G_V(q) dq + (1 - \gamma)(1 - \Phi(V)) = 1 - \pi$. Then, we can conclude that the seller's profit from a generous refund is at least $\Phi(V)$. Furthermore, the seller's profit from a generous refund is $\Phi(V)$ if $F(q) = G_V(q)$ for all $q \in [0, \gamma]$.

Let \widehat{V} be $V \in [0, \gamma]$ that minimizes $\max\{V, \Phi(V)\}$. Observe that $\Phi(V)$ is strictly decreasing in V on $[0, \gamma)$, and $\lim_{V \rightarrow \gamma} \Phi(V) = \frac{\pi - \gamma}{1-\gamma}$. Therefore, either (i) $\Phi(V) > V$ for all $V \in [0, \gamma)$ (when $\frac{\pi - \gamma}{1-\gamma} > \gamma$, or equivalently $\gamma < 1 - \sqrt{1 - \pi}$), or (ii) there exists unique $\widetilde{V} = \frac{\pi - \widetilde{V}}{1 - \gamma - \ln \frac{\widetilde{V}}{\gamma}} \in [0, \gamma]$ such that $\Phi(\widetilde{V}) = \widetilde{V}$ (when $\gamma \geq 1 - \sqrt{1 - \pi}$). Therefore,

$$\widehat{V} = \begin{cases} \gamma & \text{if } \gamma < 1 - \sqrt{1 - \pi}, \\ \widetilde{V} & \text{if } \gamma \geq 1 - \sqrt{1 - \pi}, \end{cases} \quad \text{and } \Phi(\widehat{V}) \geq \widehat{V}.$$

Now, define $F_{\widehat{V}^*}$ by

$$(6) \quad F_{\widehat{V}^*}(q) \equiv \begin{cases} 0 & \text{if } q \in [0, \Phi(\widehat{V})], \\ G_{\widehat{V}}(q) & \text{if } q \in [\Phi(\widehat{V}), \gamma], \\ 1 - \widehat{V} & \text{if } q \in [\max\{\Phi(\widehat{V}), \gamma\}, 1], \\ 1 & \text{if } q = 1. \end{cases}$$

Figures 4(a) and 4(b) illustrate $F_{\widehat{V}^*}$ when $\gamma < 1 - \sqrt{1 - \pi}$ and $\gamma \geq 1 - \sqrt{1 - \pi}$, respectively.

Lemma 2. *Suppose that the buyer's signal distribution is $F_{\widehat{V}^*}$. Then the seller's profit and profit from a generous refund are both $\Phi(\widehat{V})$, i.e., $\sup_{(p,r)} V((p,r) | F_{\widehat{V}^*}) = V_R(F_{\widehat{V}^*}) = \Phi(\widehat{V})$.*

Proof. In Appendix. □

Furthermore, for any distribution $F \neq F_{\widehat{V}^*}$, there exists a non-refundable offer or a generous refund that brings a profit higher than $\Phi(\widehat{V})$. If $F(q) < F_{\widehat{V}^*}(q)$ for some $q \in [0, \gamma)$, then the seller's profit from a non-refundable offer $(q, 0)$ is higher than \widehat{V} , i.e., $V_{NR}(F) > \widehat{V} = \Phi(\widehat{V})$. Next, suppose that $F(q) \geq F_{\widehat{V}^*}(q)$ for all $q \in [0, \gamma)$ so that $\int_{\gamma}^1 F(q) dq \leq \int_{\gamma}^1 F_{\widehat{V}^*}(q) dq$. Then the seller's profit from a generous refund $1 - \int_{\gamma}^1 \frac{F(q)}{1-\gamma} dq$ exceeds $\Phi(\widehat{V}) = 1 - \int_{\gamma}^1 \frac{F_{\widehat{V}^*}(q)}{1-\gamma} dq$. Therefore,

$$\sup_{(p,r)} V((p,r) | F) \geq \sup_{(p,r)} V((p,r) | F_{\widehat{V}^*}) \quad \text{for all } F \in \mathcal{F}.$$

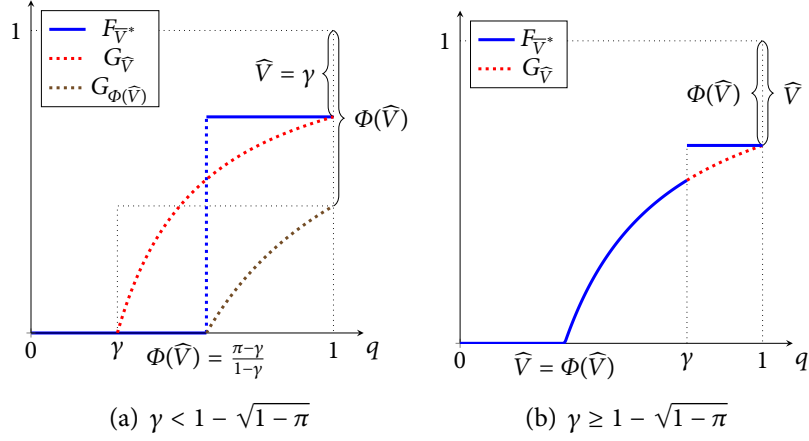


Figure 4: Worst-case distribution $F_{\bar{V}^*}$

This proves that $\bar{V}^* = \Phi(\hat{V})$, and¹⁸

$$(7) \quad \bar{V}^* = \begin{cases} \frac{\pi - \gamma}{1 - \gamma} & \text{if } \gamma < 1 - \sqrt{1 - \pi}, \\ \frac{\pi - \bar{V}^*}{1 - \gamma - \ln \frac{\bar{V}^*}{\gamma}} = \frac{-\pi}{W_{-1}\left(-\frac{\pi}{\gamma} e^{\gamma - 2}\right)} & \text{if } \gamma \geq 1 - \sqrt{1 - \pi}. \end{cases}$$

Lemma 3. *The distribution $F_{\bar{V}^*}$, defined by (6), is a seller's worst distribution. If the buyer's signal distribution is $F_{\bar{V}^*}$, then the highest profit the seller can achieve is \bar{V}^* defined by (7), and it bounds the guaranteed-profit V^* from above.*

3.3 An Optimal Stochastic Policy

We now identify the seller's best guaranteed-profit V^* using the lower bound \underline{V}^* (the red dotted line in Figure 5) and the upper bound \bar{V}^* (the blue solid line in Figure 5). For each $\varepsilon > 0$, we show that there exists a strategy $\Delta\{(p, r)\}$ such that

$$(1 - \varepsilon) \bar{V}^* \leq \mathbb{E}_{\Delta\{(p, r)\}} V((p, r) | F) \text{ for all } F \in \mathcal{F}.$$

This establishes that the seller's best guaranteed-profit is \bar{V}^* , i.e., $V^* = \bar{V}^*$.

By Lemmas 1 and 3, $\bar{V}^* = \underline{V}^*$ whenever the normalized restocking cost $\gamma \leq \bar{\gamma} = 1 - \sqrt{1 - \pi}$; and $\bar{V}^* > \underline{V}^*$ if $\gamma > \bar{\gamma}$. Therefore, if $\gamma \leq \bar{\gamma}$, then by the analysis in Section 3.1, for any $\varepsilon > 0$, there exists a generous refund that achieves $(1 - \varepsilon) \bar{V}^*$. If $\gamma \geq \bar{\gamma}$, however, then for any deterministic pricing policy (p, r) , there exists a signal distribution F such that $\bar{V}^* > V((p, r) | F)$. This implies that the seller cannot guarantee profits close to \bar{V}^* by a deterministic policy. Nevertheless, for any $\varepsilon > 0$, there exists a pricing policy that guarantees $(1 - \varepsilon) \bar{V}^*$.

To see this, we say a strategy is a *mixture of log-uniform discounts and a generous refund* if, for some $\varepsilon > 0$, the seller makes:¹⁹

¹⁸The function $W_{-1}(\cdot) \leq -1$ denotes the lower branch of the Lambert W function, i.e. function $W_{-1}(x) = z$ is defined as the smaller of the two real solutions to the equation $ze^z = x$ for $z < 0$.

¹⁹This is a variant of the exponential price auction found in Du (2018).




Figure 5: \underline{V}^* (Lemma 1) and \overline{V}^* (Lemma 3)

1. non-refundable discounted offers $(p, 0)$, $p \in [\overline{V}^*, \gamma)$, with density

$$s_0^\varepsilon(p) = \begin{cases} 0 & \text{if } \gamma \leq \overline{\gamma}, \\ \frac{1}{p(1-\gamma-\ln \frac{\overline{V}^*}{\gamma})} & \text{if } \gamma > \overline{\gamma}, \end{cases}$$

2. a generous refund $(p, r) = (1 - \varepsilon, 1 - \frac{\varepsilon}{1-\gamma})$ with probability

$$s_r^\varepsilon = \begin{cases} 1 & \text{if } \gamma \leq \overline{\gamma} \\ 1 - \int_{\overline{V}^*}^{\gamma} s_0^\varepsilon(p) dp = \frac{1-\gamma}{1-\gamma-\ln \frac{\overline{V}^*}{\gamma}} & \text{if } \gamma > \overline{\gamma} \end{cases} .$$

Before proceeding, we reemphasize that the restocking cost c is finite if and only if the normalized restocking cost γ is strictly below 1. Consequently, no matter how high the restocking cost c is, the probability of a generous refund under a mixture of log-uniform discounts and a generous refund s_r^ε is always positive, even though s_r^ε monotonically converges to zero as the restocking cost c goes to infinity (or equivalently, γ goes to 1).

Below we show that the seller can guarantee himself the upper bound \overline{V}^* identified in Section 3.2 by a mixture of log-uniform discounts and a generous refund. More specifically, the supremum of the guaranteed profit from all mixtures of log-uniform discounts and a generous refund is \overline{V}^* . Loosely speaking, any pricing policy that brings a large profit when the buyer is relatively well-informed tends to perform poorly when the buyer is relatively uninformed, and vice versa. Therefore, to increase his guaranteed profit, the seller needs to hedge simultaneously against the buyer's signal distributions that are relatively informative and relatively uninformative. A mixture of log-uniform discounts and a generous refund achieves this goal: the randomization over non-refundable offers works as a hedge against the distributions that are not informative while the refundable offer works as a hedge against the distributions that are informative.

As we formally show in the proof of Theorem 1, when he uses a mixture of log-uniform discounts and a generous refund, and the buyer's signal distribution turns out to be F , the seller's (expected) profit from some non-refundable offer $\widetilde{V}_{NR}(F)$ and from the generous

refund $\tilde{V}_R(F)$ are, respectively,

$$\tilde{V}_{NR}(F) = \frac{\gamma - \bar{V}^* - \int_{\bar{V}^*}^{\gamma} F(q) dq}{1 - \gamma - \ln \frac{\bar{V}^*}{\gamma}} \text{ and } \tilde{V}_R(F) = \frac{\pi - \gamma + \int_0^{\gamma} F(q) dq}{1 - \gamma - \ln \frac{\bar{V}^*}{\gamma}} + o_{\varepsilon},$$

The seller's profit from non-refundable offers, $\tilde{V}_{NR}(F)$, is decreasing in $\int_{\bar{V}^*}^{\gamma} F(q) dq$. Since $\int_0^1 F(q) dq = 1 - \pi$ and F is increasing, if $\int_{\bar{V}^*}^{\gamma} F(q) dq$ is small, then $\int_{\gamma}^1 F(q) dq$ is large. Loosely speaking, this occurs when $F(q)$ is relatively large and flat on $[\gamma, 1]$, i.e., F is relatively uninformative. In contrast, $\int_{\bar{V}^*}^{\gamma} F(q) dq$ is large, i.e., $\tilde{V}_{NR}(F)$ is small, when $F(q)$ is relatively small and flat on $[\gamma, 1]$, i.e., F is relatively informative. In this sense, the randomization over non-refundable offers (in the mixture of log-uniform discounts and a generous refund) brings a large profit when the buyer's signal distribution is uninformative, and a small profit when the buyer's signal distribution is informative. A similar argument shows that in clear contrast to $\tilde{V}_{NR}(F)$, the seller's profit from the generous refund, $\tilde{V}_R(F)$, brings a large profit when the buyer's signal distribution is informative, and a small profit when the buyer's signal distribution is uninformative.

Theorem 1. *For any ε , there exists a mixture of log-uniform discounts and a generous refund that guarantees profit of $(1 - \varepsilon) \bar{V}^*$. That is, the seller's best guaranteed-profit V^* is \bar{V}^* .*

Proof. In Appendix. □

Observe that the seller's best guaranteed-profit \bar{V}^* is strictly higher than the best profit the seller can guarantee without refunds $\sup_p \min_{F \in \mathcal{F}} V((p, 0) | F)$ (i.e., the brown dotted-line in Figure 5). The reason is that without the opportunity of making refundable offers, which work as a hedge against informative signal distributions, the seller is forced to use a pricing policy that induces lower prices more frequently. Then, however, the resulting profit would be lower when the buyer's signal distribution turned out to be not so informative.²⁰

3.4 Best Guaranteed-Profit with General Mechanisms

We have identified the best guaranteed-profit V^* that the seller can achieve using a simple randomized pricing and refund policy. One may wonder if the seller can improve his best guaranteed-profit by using a more intricate mechanism that screens the buyer based on her signal distribution F and her realized signal q . Below we show that there exists no such mechanism. More specifically, for any (indirect or direct) mechanism, the seller's profit is bounded from above by V^* for some buyer's signal distribution.

To show this, we analyze an auxiliary game in which the nature chooses the buyer's signal distribution F , and then the seller chooses a mechanism after observing the nature's choice. We show that if the buyer's signal distribution is $F_{\bar{V}^*}$, then the seller's profit

²⁰The result reported here is closely related to the finding in Du (2018) that analyzes the environment where the seller cannot make refundable offers (or equivalently the restocking cost is infinitely large). Du (2018) shows that the distribution $F_{\bar{V}_{NR}^*}$, where $\bar{V}_{NR}^* \in (0, \bar{V}^*)$ that uniquely solves $\bar{V}_{NR}^*(1 - \ln \bar{V}_{NR}^*) = \pi$, is the worst possible buyer's information structure for the seller; and that the log-uniformly randomized prices over $[\bar{V}_{NR}^*, 1]$ achieves the seller's best guaranteed-profit regardless of the buyer's signal distribution. We also note that Roesler and Szentes (2017) show that this distribution $F_{\bar{V}_{NR}^*}$ maximizes the buyer's welfare when the seller best responds to this information structure via uniform pricing.

is bounded from above by V^* . Since the seller's best guaranteed-profit is bounded from below by V^* , the seller's equilibrium profit of this auxiliary game is V^* . Furthermore, the seller's equilibrium profit of this auxiliary game defines an upper bound of the seller's best guaranteed-profit.

For a (possibly indirect) mechanism Ψ , denote the seller's profit from the buyer with signal q by $v(q|\Psi)$. We say that an outcome (induced by mechanism Ψ) is undominated if there exists no mechanism $\tilde{\Psi}$ such that $v(q|\tilde{\Psi}) \geq v(q|\Psi)$ for all q , and $v(q|\tilde{\Psi}) > v(q|\Psi)$ for some q . The lemma below establishes that for any undominated outcome, there exists an outcome-equivalent simple static direct mechanism with refunds.

Definition 1. We say a mechanism $M \equiv \{p(q), \{\alpha_0(q), \alpha_r(q)\}_{q \in [0,1]}\}$ is a direct mechanism with refunds if, for each buyer's report $q \in [0, 1]$, the mechanism specifies (i) $p(q)$: the transfer from the buyer to the seller; (ii) $\alpha_0(q) \in [0, 1]$: the probability that the buyer receives the product without an option to return; and (iii) $\alpha_r(q) \in [0, 1 - \alpha_0(q)]$: the probability that the buyer receives the product with an option to return with refund $r = 1$.

Lemma 4. Take an arbitrary subgame following the nature's move. For any undominated outcome the seller can induce by an indirect mechanism, there exists an outcome-equivalent direct mechanism with refunds that is individually rational and incentive compatible.

Proof. In the appendix. □

Let F be the nature's choice. Under a given direct mechanism with refunds $M = \{p(q), \{\alpha_0(q), \alpha_r(q)\}_{q \in [0,1]}\}$, if the buyer with signal q reports q' , then her payoff is

$$\begin{aligned} U(q'; q|M) &\equiv (\alpha_0(q') + \alpha_r(q'))q + \alpha_r(q')(1 - q) - p(q') \\ &= q\alpha_0(q') - (p(q') - \alpha_r(q')). \end{aligned}$$

The seller's profit is $\mathbb{E}_F[v(q|M)]$, where $v(q|M)$ is his profit from the buyer with signal q , i.e.,

$$v(q|M) \equiv p(q) - \alpha_r(q)(1 - q)(c + 1) = p(q) - \alpha_r(q) \frac{1 - q}{1 - \gamma}.$$

Let \mathcal{M} be the set of all direct mechanisms with refunds M that satisfy the following two conditions:

$$\begin{aligned} \text{(IC)} \quad & U(q; q|M) \geq U(q'; q|M) \text{ for all } q' \text{ and } q, \\ \text{(IR)} \quad & U(q; q|M) \geq 0 \text{ for all } q. \end{aligned}$$

By adopting the standard argument, we can simplify the seller's problem to the one in which he chooses an increasing function $\alpha_0(\cdot)$ (instead of M , which is a triplet of functions, that satisfies IC and IR). More formally, the incentive compatibility condition (IC) is equivalent to

$$\alpha_0(q) \text{ is increasing in } q \text{ and } U(q; q|M) = \int_0^q \alpha_0(\tilde{q}) d\tilde{q}.$$

Therefore, if $M^* = \{p^*(q), \{\alpha_0^*(q), \alpha_r^*(q)\}_{q \in [0,1]}\}$ is the seller's best-response to the nature's choice F , then we can represent $v(q|M^*)$, i.e., the seller's profit from buyer with signal q , independent of $p^*(q)$ and $\alpha_r^*(q)$.

Lemma 5. Suppose that $M^* = \{p^*(q), \{\alpha_0^*(q), \alpha_r^*(q)\}_{q \in [0,1]}\}$ is a seller's best-response to the nature's choice F , then

$$\alpha_0^*(q) \text{ is increasing, } \alpha_r^*(q) = \begin{cases} 0 & \text{if } q < \gamma, \\ 1 - \alpha_0^*(q) & \text{if } q \geq \gamma, \end{cases}$$

and, $v(q|M^*) = v(q|\alpha_0^*)$, where

$$(8) \quad v(q|\alpha_0^*) \equiv \begin{cases} q\alpha_0^*(q) - \int_0^q \alpha_0^*(\bar{q}) d\bar{q} & \text{if } q < \gamma, \\ q\alpha_0^*(q) + \frac{q-\gamma}{1-\gamma} (1 - \alpha_0^*(q)) - \int_0^q \alpha_0^*(\bar{q}) d\bar{q} & \text{if } q \geq \gamma. \end{cases}$$

Proof. In Appendix. □

In the subgame where the nature has chosen F_{V^*} , the seller's payoff is V^* .

Lemma 6. Suppose that $M^* = \{p^*(q), \{\alpha_0^*(q), \alpha_r^*(q)\}_{q \in [0,1]}\}$ is a seller's best-response to the nature's choice F_{V^*} . Then $\mathbb{E}_{F_{V^*}} [v(q|M^*)] = \mathbb{E}_{F_{V^*}} [v(q|\alpha_0^*)] = V^*$.

Proof. In the appendix. □

Therefore, the seller's equilibrium profit in this auxiliary game is V^* . Furthermore, the seller's equilibrium profit defines an upper bound of the seller's best-guaranteed profit, which is bounded from below by V^* . We thus have the required result.

Theorem 2. Suppose that the seller can guarantee himself V using some mechanism. Then, for any $\varepsilon > 0$, there exists a mixture of log-uniform discounts and a generous refund that guarantees $(1 - \varepsilon)V$.

4 Buyer-Optimal Outcomes

In this section, we identify the buyer-optimal information structure and the outcomes that can be supported by some information structure. We analyze the following game:

1. The buyer chooses a signal distribution $F \in \mathcal{F}$.
2. The seller observes the signal structure and chooses a mechanism. By the arguments above, we can without loss in generality focus on (deterministic) policies (p, r) .
3. The buyer observes a signal generated by F , and then decides whether to purchase or not.

The key assumption we impose is that the buyer commits to a signal distribution at the beginning of the game. In other words, the buyer commits not to learn any information that is not contained in the signal generated using distribution F .²¹ Note that if the buyer lacks such a commitment power, then she would choose to acquire a fully informative signal. Knowing this, the seller would always set the price to 1 and can capture the full surplus.

²¹An interpretation of our buyer-optimality result is as follows: the buyer ex-ante delegates the information gathering to a third-party, such as an agent or an algorithm, knowing that the seller best responds to the signal structure by choosing a pricing policy.

Given our focus on the buyer-optimal outcome, we assume that the seller makes an offer that induces a higher buyer's payoff in case he is indifferent between two or more offers. Therefore, the assumption that the seller only uses a deterministic policy is without loss. More formally, let $U((p, r) | F)$ be the buyer's payoff when her signal distribution $F \in \mathcal{F}$, and the offer is (p, r) . Also, let $V((p, r) | F)$ be the corresponding seller's profit. Let $(p^*, r^*)_F$ be an element in $S^*(F) \equiv \arg \max_{(p, r)} V((p, r) | F)$ such that

$$U((p^*, r^*)_F | F) = \max_{(p^*, r^*) \in S^*(F)} U((p^*, r^*) | F).$$

Our goal is to characterize

$$F^* \in \arg \max_{F \in \mathcal{F}} U((p^*, r^*)_F | F), \text{ and } U^* \equiv U((p^*, r^*)_{F^*} | F^*).$$

Note that the seller's attempt to maximize his profit can cause inefficiency through two channels. First, the seller may make an offer that results in no sales even if the buyer's expected value is strictly positive. Second, even when the buyer accepts the seller's offer, the product may be returned. Nevertheless, there exists a buyer-optimal outcome that is efficient.

The total surplus from trade is bounded from above by π . Furthermore, the seller can guarantee a profit of V^* irrespective of the choice the buyer makes (Theorem 1). Therefore, the buyer's payoff under the buyer-optimal information structure U^* is bounded from above $\pi - V^*$. However, the preceding analysis implies that this bound is sharp, i.e., $U^* = \pi - V^*$, and the worst distribution F_{V^*} for the seller identified in Section 3.2 is the buyer-optimal signal distribution. This is because when the buyer's signal distribution is F_{V^*} the seller's profit is maximized by offering $(p, r) = (V^*, 0)$. More precisely,

$$\sup_{(p, r)} V((p, r) | F_{V^*}) = V^* = V((V^*, 0) | F_{V^*}).$$

Furthermore, when the seller offers $(V^*, 0)$, the trade occurs with probability one and there are no returns. Consequently, $U((V^*, 0) | F) = \pi - V^*$.

Theorem 3. *The seller's worst signal distribution F_{V^*} is a buyer-optimal signal distribution. Under the corresponding buyer-optimal outcome, the seller's profit and the buyer's payoff are V^* and $\pi - V^*$, respectively.*

As an immediate corollary, we can characterize pairs of buyer's payoff and seller's profit that can be supported by some signal structure. More precisely, we say (\hat{U}, \hat{V}) is a *feasible outcome* supported by a F if (\hat{p}, \hat{r}) maximizes $V((p, r) | F)$ with respect to (p, r) ; $\hat{V} = V((\hat{p}, \hat{r}) | F)$; and $\hat{U} = U((\hat{p}, \hat{r}) | F)$. We denote the set of all feasible outcomes by \mathcal{O} . It follows from the previous discussion that $\mathcal{O} = \{(U, V) : V \in [V^*, \pi - U], U \in [0, \pi - V^*]\}$.

5 Discussion

We analyzed the seller's robust pricing problem with uncertainty about the buyer's information and learning, and showed that a simple mechanism that utilizes a generous refund achieves the best guaranteed-profit.

Our model hinges on the key, but restrictive, assumption of a binary buyer valuation, which significantly simplifies the analysis. This assumption is reasonable if buyers' primary concerns are of the exact product match: e.g., shoes and clothes either do or don't fit, a gadget is either compatible with the buyer's use or not, and a business traveler either needs to be in a particular location on a specific date or not. We note that the implication of this assumption is twofold. First, there exists a one-to-one mapping between signals and expected willingness to pay for the product (in the absence of a refund policy). Second, a buyer's return decision is independent of the amount of refund.

Those two properties fail to hold if more than two levels of product fit are present. Thus, the assumption of binary buyer valuation, whilst restrictive, is important for the tractability of the model. More precisely, in the absence of a refund policy, the buyer buys only if her signal q , i.e., a posterior over possible levels of product fit v , satisfies $\mathbb{E}_q[v] \geq p$. That is, the seller's profit depends on the buyer's signal q only through the value of expected willingness to pay it induces, i.e., $\mathbb{E}_q[v]$. In this sense, if the seller cannot utilize a refund policy, then we still can represent the seller's uncertainty as the uncertainty over distributions over expected willingness to pay, which is a one-dimensional random variable.²² In the presence of a refund policy (p, r) , however, the buyer with signal q buys only if the right-tail of q is sufficiently fat, i.e., $\Pr(v > r|q) \mathbb{E}_q[v|v > r] + \Pr(v \leq r|q) r \geq p$. Therefore, two signals with an identical willingness to pay can result in different outcomes for the seller. Consequently, we are no longer able to capture the seller's uncertainty in terms of the uncertainty over distributions on expected willingness to pay, rendering the generalization into this direction not straightforward.

Having said that, we conjecture that the relaxation of the binary buyer valuation assumption yields qualitatively similar results: for any given restocking cost, there exists a generous refund policy (or a randomization over refund policies) that guarantees the seller a non-negative ex-post (expected) profit; and hence the seller can strictly improve his profit against the worst signal distribution by utilizing such a refund policy.

Another aspect that we did not address is the seller's learning of demand through pricing. Our insight that a well-designed refund policy limits the significance of buyer learning on the seller's profit should apply even in a dynamic environment. However, in a dynamic environment, carefully designed dynamic pricing and resulting buyer's purchasing and return decisions can also be used to learn about what/how buyers learn about the product fit, and hence to improve future pricing decisions. Investigating how the seller's learning motive would shape the intertemporal pricing with refunds would be an interesting venue for future research.

Another possible venue is the application to platform designs. A platform, such as eBay or Airbnb, could choose information that is revealed to the buyer. The seller chooses a pricing and return policy, and the buyer makes a purchasing and return decision. In the absence of competition among platforms, it is natural to expect that one of the Pareto-optimal outcomes identified in the previous section would arise. Analyzing how competition affects the welfare and equilibrium information structure is an interesting question.

²²More precisely, the distributions (over signals) that minimize the seller's profit from a deterministic non-refundable offer; and the seller's worst distribution over signals when the seller cannot utilize a refund policy can both be characterized in terms of distributions over expected willingness to pay.

References

- AU, P. H. AND K. KAWAI (2017a): “Competitive Disclosure of Correlated Information,” *Working Paper, UNSW*.
- (2017b): “Competitive Information Disclosure by Multiple Senders,” *Working Paper, UNSW*.
- AUMANN, R., M. MASCHLER, AND R. STEARNS (1995): *Repeated Games with Incomplete Information*, MIT Press.
- BOLES LAVSKY, R. AND C. COTTON (2018): “Limited Capacity in Project Selection: Competition through Evidence Production,” *Economic Theory*, 65, 385–421.
- CHE, Y.-K. (1996): “Customer Return Policies for Experience Goods,” *Journal of Industrial Economics*, 44, 17–24.
- COURTY, P. AND H. LI (2000): “Sequential Screening,” *Review of Economic Studies*, 67, 697–717.
- DU, S. (2018): “Robust Mechanisms Under Common Valuation,” *Econometrica*, forthcoming.
- ELY, J. C., D. GARRETT, AND T. HINNOSAAR (2017): “Overbooking,” *Journal of the European Economic Association*, 15, 1258–1301.
- ESCOBARI, D. AND P. JINDAPON (2014): “Price Discrimination through Refund Contracts in Airlines,” *International Journal of Industrial Organization*, 34, 1–8.
- ESÓ, P. AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74, 705–731.
- GALE, I. L. AND T. J. HOLMES (1992): “The Efficiency of Advance-Purchase Discounts in the Presence of Aggregate Demand Uncertainty,” *International Journal of Industrial Organization*, 10, 413–437.
- (1993): “Advance-Purchase Discounts and Monopoly Allocation of Capacity,” *American Economic Review*, 83, 135–146.
- GALLEGO, G. AND O. SAHIN (2010): “Revenue Management with Partially Refundable Fares,” *Operations Research*, 58, 817–833.
- GROSSMAN, S. J. (1981): “The Informational Role of Warranties and Private Disclosure about Product Quality,” *Journal of Law & Economics*, 24, 461–483.
- INDERST, R. AND M. OTTAVIANI (2013): “Sales Talk, Cancellation Terms and the Role of Consumer Protection,” *Review of Economic Studies*, 80, 1002–1026.
- INDERST, R. AND G. TIROSH (2015): “Refunds and Returns in a Vertically Differentiated Industry,” *International Journal of Industrial Organization*, 38, 44–51.

- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KRAHMER, D. AND R. STRAUZ (2015): “Optimal Sales Contracts with Withdrawal Rights,” *The Review of Economic Studies*, 82, 762–790.
- LIBGOBER, J. AND X. MU (2017): “Informational Robustness in Intertemporal Pricing,” *Working Paper, Harvard University*.
- MOORTHY, S. AND K. SRINIVASAN (1995): “Signaling Quality with a Money-Back Guarantee: The Role of Transaction Costs,” *Marketing Science*, 14, 442–466.
- NOCKE, V., M. PEITZ, AND F. ROSAR (2011): “Advance-Purchase Discounts as a Price Discrimination Device,” *Journal of Economic Theory*, 146, 141–162.
- ROESLER, A.-K. AND B. SZENTES (2017): “Buyer-Optimal Learning and Monopoly Pricing,” *American Economic Review*, 107, 2072–80.
- VON WANGENHEIM, J. (2017): “Consumer-Optimal Information Design,” *Rationality and Competition Discussion Paper Series*, 53.
- ZHANG, J. (2013): “Revenue Maximizing with Return Policy When Buyers Have Uncertain Valuations,” *International Journal of Industrial Organization*, 31, 452–461.

A Proofs for Section 3 (Best Guaranteed-Profit)

Proof of Lemma 2: First, suppose that $\gamma < 1 - \sqrt{1 - \pi}$. Then $V_{NR}(F_{\widehat{V}^*}) = \Phi(\widehat{V})$. Consider a refundable offer (p, r) . If $\tilde{q}(p, r) > \Phi(\widehat{V})$, then

$$V((p, r) | F_{\widehat{V}^*}) = \gamma p \leq \gamma < \frac{\pi - \gamma}{1 - \gamma} = \Phi(\widehat{V}).$$

If $\tilde{q}(p, r) \leq \Phi(\widehat{V})$, then

$$\begin{aligned} V((p, r) | F_{\widehat{V}^*}) &= p - (1 - \pi)(r + c) \\ &= p - (1 - \gamma)(1 - \Phi(\widehat{V})) \left(\frac{p - \tilde{q}(p, r)}{1 - \tilde{q}(p, r)} + \frac{\gamma}{1 - \gamma} \right) \text{ (increasing in } \tilde{q}(p, r), p) \\ &\leq 1 - (1 - \gamma)(1 - \Phi(\widehat{V})) \left(\frac{1 - \Phi(\widehat{V})}{1 - \Phi(\widehat{V})} + \frac{\gamma}{1 - \gamma} \right) = \Phi(\widehat{V}). \end{aligned}$$

Next, suppose that $\gamma \geq 1 - \sqrt{1 - \pi}$. Then $V_{NR}(F_{\widehat{V}^*}) = \widehat{V} = \Phi(\widehat{V})$. Consider a refundable offer (p, r) . If $\tilde{q}(p, r) > \gamma$, then $V((p, r) | F_{\widehat{V}^*}) \leq \widehat{V} = \Phi(\widehat{V})$. If $\tilde{q}(p, r) \leq \gamma$, then $V((p, r) | F_{\widehat{V}^*}) < \Phi(\widehat{V})$ because the seller's profit from the buyer with signal $q \in [\tilde{q}(p, r), \gamma]$ is negative.

Proof of Theorem 1: We limit our attention to the case where $\gamma > \bar{\gamma}$. Fix arbitrary signal structure F . The seller's payoff of using policy with ε is

$$\mathbb{E}[V((p, r) | F)] = \underbrace{\widetilde{V}_{NR}(F)}_{\text{profit from discounts}} + \underbrace{\widetilde{V}_R^\varepsilon(F)}_{\text{profit from a generous refund}},$$

where

$$\begin{aligned} \widetilde{V}_{NR}(F) &\equiv \int_{\widehat{V}^*}^{\gamma} p(1 - F(p)) s_0^\varepsilon(p) dp, \\ \widetilde{V}_R^\varepsilon(F) &\equiv s_r^\varepsilon \int_{\gamma}^1 \left((1 - \varepsilon) - (1 - q) \left(\left(1 - \frac{\varepsilon}{1 - \gamma} \right) + c \right) \right) dF(q). \end{aligned}$$

Observe that

$$\begin{aligned} \widetilde{V}_{NR}(F) &= \int_{\widehat{V}^*}^{\gamma} p(1 - F(p)) s_0^\varepsilon(p) dp = \frac{1}{1 - \gamma - \ln \frac{\widehat{V}^*}{\gamma}} \int_{\widehat{V}^*}^{\gamma} (1 - F(p)) dp \\ &= \frac{\gamma - \widehat{V}^* - \int_{\widehat{V}^*}^{\gamma} F(p) dp}{1 - \gamma - \ln \frac{\widehat{V}^*}{\gamma}} \geq \frac{\gamma - \widehat{V}^* - \int_0^{\gamma} F(p) dp}{1 - \gamma - \ln \frac{\widehat{V}^*}{\gamma}}, \end{aligned}$$

and

$$\begin{aligned} \widetilde{V}_R^\varepsilon(F) &= s_r^\varepsilon \int_{\gamma}^1 \left(1 - \varepsilon - (1 - q) \left(1 - \frac{\varepsilon}{1 - \gamma} + \frac{\gamma}{1 - \gamma} \right) \right) dF(q) \\ &= \frac{1 - \gamma}{1 - \gamma - \ln \frac{\widehat{V}^*}{\gamma}} \left(1 - (1 + c) \int_{\gamma}^1 F(q) dq - \varepsilon \int_{\gamma}^1 \left(1 + \frac{1 - q}{1 - \gamma} \right) dF(q) \right) \\ &= \frac{\pi - \gamma + \int_0^{\gamma} F(p) dp}{1 - \gamma - \ln \frac{\widehat{V}^*}{\gamma}} - \frac{\varepsilon(1 - \gamma)}{1 - \gamma - \ln \frac{\widehat{V}^*}{\gamma}} \int_{\gamma}^1 \left(1 + \frac{1 - q}{1 - \gamma} \right) dF(q). \end{aligned}$$

Recall that $\frac{\pi - \bar{V}^*}{1 - \gamma - \ln \frac{\bar{V}^*}{\gamma}} = \bar{V}^*$ by (7). We thus have

$$\begin{aligned} \mathbb{E} [V((p, r) | F)] &= \widetilde{V}_{NR}(F) + \widetilde{V}_R^\varepsilon(F) \\ &\geq \frac{\pi - \bar{V}^*}{1 - \gamma - \ln \frac{\bar{V}^*}{\gamma}} - \frac{\varepsilon(1 - \gamma)}{1 - \gamma - \ln \frac{\bar{V}^*}{\gamma}} \int_{\gamma}^1 \left(1 + \frac{1 - q}{1 - \gamma}\right) dF(q) \\ &= \bar{V}^* - \frac{\varepsilon(1 - \gamma)}{1 - \gamma - \ln \frac{\bar{V}^*}{\gamma}} \int_{\gamma}^1 \left(1 + \frac{1 - q}{1 - \gamma}\right) dF(q). \end{aligned}$$

Proof of Lemma 4: We say a direct mechanism $\Psi = \left\{ \alpha_q, p_q, \left\{ (\kappa_q^v, \tau_q^v) \right\}_{v \in \{0,1\}} \right\}_{q \in [0,1]}$ is a two-step mechanism if it specifies, for each reported signal $q \in [0, 1]$, (i) α_q : the probability that the buyer receives the product; (ii) p_q : the transfer from the buyer to the seller; and (iii) $\left\{ \kappa_q^v, \tau_q^v \right\}_{v \in \{0,1\}}$: the direct mechanism that specifies, for each buyer's reported realized valuation $v \in \{0, 1\}$, the pair of probability the buyer keeps the product κ_q^v and the transfer from the seller to the buyer τ_q^v such that $\kappa_q^v v + \tau_q^v \geq \max \{v, \kappa_q^{v'} v + \tau_q^{v'}\}$ for $v \in \{0, 1\}$ and $v' \neq v$.

If the buyer with signal q reports q' under two-step mechanism Ψ , then her payoff is

$$u(q'; q | \Psi) \equiv \alpha_{q'} q (\kappa_{q'}^1 + \tau_{q'}^1) + \alpha_{q'} (1 - q) \tau_{q'}^0 - p_{q'}.$$

We say a two-step mechanism Ψ is an IR-IC two-step mechanism if $u(q, q | \Psi) \geq 0$ for all q and $u(q; q | \Psi) \geq u(q'; q | \Psi)$ for all q and q' . The seller's profit from the buyer with signal q under an IR-IC two-step mechanism Ψ is

$$\begin{aligned} v(q | \Psi) &\equiv p_q - \alpha_q q \left((1 - \kappa_q^1) c + \tau_q^1 \right) - \alpha_q (1 - q) \left((1 - \kappa_q^0) c + \tau_q^0 \right) \\ &= -u(q; q | \Psi) - \alpha_q q \left((1 - \kappa_q^1) c - \kappa_q^1 \right) - \alpha_q (1 - q) \left(1 - \kappa_q^0 \right) c. \end{aligned}$$

Take an arbitrary *undominated* outcome. By the revelation principle, there exists an outcome-equivalent IR-IC two-step mechanism $\widetilde{\Psi} = \left\{ \widetilde{\alpha}_q, \widetilde{p}_q, \left\{ (\widetilde{\kappa}_q^v, \widetilde{\tau}_q^v) \right\}_{v \in \{0,1\}} \right\}_{q \in [0,1]}$. Define a two-step mechanism Ψ as

- (i) $\alpha_q = \widetilde{\alpha}_q$;
- (ii) $p_q = \widetilde{p}_q + \widetilde{\alpha}_q \left(1 - (\widetilde{\kappa}_q^1 + \widetilde{\tau}_q^1) \right)$;
- (iii) $(\kappa_q^0, \tau_q^0) = (\widetilde{\kappa}_q^1 + \widetilde{\tau}_q^1 - \widetilde{\tau}_q^0, 1 + \widetilde{\tau}_q^0 - (\widetilde{\kappa}_q^1 + \widetilde{\tau}_q^1))$; and
- (iv) $(\kappa_q^1, \tau_q^1) = (1, 0)$.

We first show that $\Psi \equiv \left\{ \alpha_q, p_q, \left\{ (\kappa_q^v, \tau_q^v) \right\}_{v \in \{0,1\}} \right\}_{q \in [0,1]}$ is an IR-IC two-step mechanism, and outcome-equivalent to $\widetilde{\Psi}$. We then show that Ψ can be implemented by a direct mechanism with refunds.

To verify that Ψ is a two-step mechanism, we show that $\kappa_q^v \in [0, 1]$, $\tau_q^v \in [0, 1]$, $\kappa_q^1 + \tau_q^1 \geq \max \{1, \kappa_q^0 + \tau_q^0\}$, and $\tau_q^0 \geq \max \{0, \tau_q^1\}$. Observe that $\kappa_q^0 + \tau_q^0 = \kappa_q^1 + \tau_q^1 = 1$, and $\tau_q^1 = 0$.

Thus, we are done if we show that $\tau_q^0 \in [0, 1]$. Since $\tilde{\Psi}$ is a two-step mechanism, $\tilde{\tau}_q^0 \geq \tilde{\tau}_q^1$ and $\tilde{\kappa}_q^1 + \tilde{\tau}_q^1 \geq \tilde{\kappa}_q^0 + \tilde{\tau}_q^0$, which imply

$$\begin{aligned}\tau_q^0 &= 1 + \tilde{\tau}_q^0 - (\tilde{\kappa}_q^1 + \tilde{\tau}_q^1) \geq 1 + \tilde{\tau}_q^0 - (\tilde{\kappa}_q^1 + \tilde{\tau}_q^0) = 1 - \tilde{\kappa}_q^1 \geq 0; \text{ and} \\ \tau_q^0 &= 1 + \tilde{\tau}_q^0 - (\tilde{\kappa}_q^1 + \tilde{\tau}_q^1) \leq 1 + \tilde{\tau}_q^0 - (\tilde{\kappa}_q^0 + \tilde{\tau}_q^0) = 1 - \tilde{\kappa}_q^0 \leq 1.\end{aligned}$$

Ψ is an IR-IC two-step mechanism because $\tilde{\Psi}$ is an IR-IC two-step mechanism, and for all q and q' ,

$$\begin{aligned}u(q'; q|\Psi) - u(q'; q|\tilde{\Psi}) &= \alpha_{q'} q (1 - (\tilde{\kappa}_{q'}^1 + \tilde{\tau}_{q'}^1)) + \alpha_{q'} (1 - q) (1 - (\tilde{\kappa}_{q'}^1 + \tilde{\tau}_{q'}^1)) + \tilde{p}_q - p_q \\ &= \alpha_{q'} (1 - (\tilde{\kappa}_{q'}^1 + \tilde{\tau}_{q'}^1)) - \alpha_{q'} (1 - (\tilde{\kappa}_q^1 + \tilde{\tau}_q^1)) = 0.\end{aligned}$$

IR-IC two-step mechanisms Ψ and $\tilde{\Psi}$ are outcome-equivalent if $v(q|\Psi) = v(q|\tilde{\Psi})$ for all q . Observe that

$$\begin{aligned}v(q|\Psi) - v(q|\tilde{\Psi}) &= \alpha_q q (1 - \tilde{\kappa}_q^1) (1 + c) + \alpha_q (1 - q) ((\tilde{\kappa}_q^1 + \tilde{\tau}_q^1) - (\tilde{\kappa}_q^0 + \tilde{\tau}_q^0)) c \\ &\geq 0.\end{aligned}$$

However, since the outcome of $\tilde{\Psi}$ is undominated by assumption, the above inequality must hold with equality for all q . Thus, Ψ is an IR-IC two-step mechanism that is outcome-equivalent to $\tilde{\Psi}$.

Using Ψ , we define a direct mechanism with refunds $M = \{p(q), \{\alpha_0(q), \alpha_r(q)\}\}_{q \in [0, 1]}$ by $\alpha_0(q) = \alpha_q \kappa_q^0$, $\alpha_r(q) = \alpha_q (1 - \kappa_q^0)$, and $p(q) = p_q$. Under M , if the buyer with signal q reports q' , then her payoff is

$$\begin{aligned}\alpha_0(q') q + \alpha_r(q') - p(q') &= \alpha_{q'} \kappa_{q'}^0 q + \alpha_{q'} (1 - \kappa_{q'}^0) - p_{q'} \\ &= \alpha_{q'} q + \alpha_{q'} (1 - q) (1 - \kappa_{q'}^0) - p_{q'} \\ &= u(q'; q|\Psi).\end{aligned}$$

The seller's profit from the buyer with signal q (who truthfully reports q) is thus

$$\begin{aligned}p(q) - \alpha_r(q) (1 - q) (1 + c) &= p_q - \alpha_q (1 - q) (1 - \kappa_q^0) (1 + c) \\ &= v(q|\Psi).\end{aligned}$$

This establishes that M implements Ψ .

Proof of Lemma 5: Since $U(q; q|M) = \int_0^q \alpha_0(\bar{q}) d\bar{q}$,

$$\begin{aligned}v(q|M) &= p(q) - a_r(q) \frac{1 - q}{1 - \gamma} = q \alpha_0(q) + \alpha_r(q) - \int_0^q \alpha_0(\bar{q}) d\bar{q} - a_r(q) \frac{1 - q}{1 - \gamma} \\ &= q \alpha_0(q) + \frac{q - \gamma}{1 - \gamma} \alpha_r(q) - \int_0^q \alpha_0(\bar{q}) d\bar{q}.\end{aligned}$$

Observe that $v(0|M) \leq 0$. Therefore, if M^* is a solution and $q = 0$, then the seller chooses $\alpha_0^*(q) = \alpha_r^*(q) = 0$. Next, if $q \in (0, \gamma)$, then since $\frac{q - \gamma}{1 - \gamma} < 0$, $\alpha_r^*(q) = 0$. If $q = \gamma$, $v(\gamma|M^*)$

does not depend on $\alpha_r^*(q)$. Therefore, $\alpha_r^*(q) = 1 - \alpha_0^*(q)$. Similarly, if $q \in (\gamma, 1]$, then since $\frac{q-\gamma}{1-\gamma} > 0$, $\alpha_r^*(q) = 1 - \alpha_0^*(q)$.

Proof of Lemma 6: If $\gamma < 1 - \sqrt{1 - \pi}$, then the support of $F_{\bar{V}^*}$ is $\{\bar{V}^*, 1\}$. Thus, by (8), we can conclude that $\alpha_0^*(q) = 0$ on $[0, \bar{V}^*)$, and $\alpha_0^*(q) = \alpha_0^*(\bar{V}^*)$ on $(\bar{V}^*, 1)$. Since, $F_{\bar{V}^*}$ induces \bar{V}^* and 1 with probability $1 - \gamma$ and γ , respectively, the seller's profit from using M^* is

$$\begin{aligned} \mathbb{E}_{F_{\bar{V}^*}} [v(q|\alpha_0^*)] &= (1 - \gamma) \underbrace{\left[\bar{V}^* \alpha_0^*(\bar{V}^*) + \frac{\bar{V}^* - \gamma}{1 - \gamma} (1 - \alpha_0^*(\bar{V}^*)) \right]}_{=v(\bar{V}^*|\alpha_0^*)} + \gamma \underbrace{[1 - \alpha_0^*(\bar{V}^*) (1 - \bar{V}^*)]}_{=v(1|\alpha_0^*)} \\ &= \bar{V}^*. \end{aligned}$$

If $\gamma \geq 1 - \sqrt{1 - \pi}$, then the support of $F_{\bar{V}^*}$ is $[\bar{V}^*, \gamma] \cup \{1\}$. Therefore, $\alpha_0^*(q) = 0$ on $[0, \bar{V}^*)$, and $\alpha_0^*(q) = \alpha_0^*(\gamma)$ on $(\gamma, 1)$. Furthermore, $F_{\bar{V}^*}$ has density $\frac{\bar{V}^*}{q^2}$ over the interval $[\bar{V}^*, \gamma)$, and two mass points, γ (with probability $\frac{\bar{V}^*}{\gamma} - \bar{V}^*$) and 1 (with probability \bar{V}^*). Thus, the seller's profit from using M^* is

$$\begin{aligned} \mathbb{E}_{F_{\bar{V}^*}} [v(q|\alpha_0^*)] &= \int_{\bar{V}^*}^{\gamma} \underbrace{\left[q\alpha_0^*(q) - \int_{\bar{V}^*}^q \alpha_0^*(\tilde{q})d\tilde{q} \right]}_{=v(q|\alpha_0^*), q \in [0, \gamma)} \frac{\bar{V}^*}{q^2} dq + \left(\frac{\bar{V}^*}{\gamma} - \bar{V}^* \right) \underbrace{\left[\gamma\alpha_0^*(\gamma) - \int_{\bar{V}^*}^{\gamma} \alpha_0^*(q)dq \right]}_{=v(\gamma|\alpha_0^*)} \\ &\quad + \bar{V}^* \underbrace{\left[1 - \int_{\bar{V}^*}^{\gamma} \alpha_0^*(\tilde{q})d\tilde{q} - \alpha_0^*(\gamma)(1 - \gamma) \right]}_{=v(1|\alpha_0^*)}. \end{aligned}$$

Since $\int_{\bar{V}^*}^{\gamma} (q\alpha_0^*(q) - \int_{\bar{V}^*}^q \alpha_0^*(\tilde{q})d\tilde{q}) \frac{\bar{V}^*}{q^2} dq = \frac{\bar{V}^*}{\gamma} \int_{\bar{V}^*}^{\gamma} \alpha_0^*(q)dq$, we have $\mathbb{E}_{F_{\bar{V}^*}} [v(q|\alpha_0^*)] = \bar{V}^*$.