International Journal of Industrial Organization 38 (2015) 44–51

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Refunds and returns in a vertically differentiated industry☆

Roman Inderst [a,b,*], Gilad Tirosh [a]

[a] Johann Wolfgang Goethe University Frankfurt, Germany
[b] Imperial College London, United Kingdom

ARTICLE INFO

Article history:
Received 14 January 2014
Received in revised form 20 October 2014
Accepted 10 December 2014
Available online 13 January 2015

Keywords:
Refunds
Cancelation terms
Vertical differentiation

ABSTRACT

Firms frequently offer refunds, both when physical products are returned and when service contracts are terminated prematurely. We show how refunds act as a "metering device" when consumers learn about their personal valuation while experimenting with the product or service. Our theory predicts that low-quality firms offer inefficiently strict terms for refunds, while high-quality firms offer inefficiently generous terms. This may help to explain the observed variety in contractual terms.

1. Introduction

Refunds and termination clauses are ubiquitous. For instance, the annual turnover of product return in the U.S. retail industry exceeds 100 billion U.S. dollars, of which 70% are due to reasons of taste and fit (see Posselt et al., 2008; Anderson et al., 2009). Likewise, contracts for services, including insurances, utilities, or subscriptions to health clubs, frequently involve cancelation clauses allowing customers to terminate prematurely. As we document below, however, there is wide variety in the use that firms make of such contracts. We present a theory based on firms' vertical differentiation that generates such heterogeneity. In our model, firms use refund and return policies to extract a higher share of the information rent of those consumers who have a stronger preference for the respective product. In equilibrium, high-quality firms offer excessively generous terms, while low-quality firms offer excessively strict terms.

Several studies find that higher-quality retailers, such as up-market stores or internet retailers with a higher customer rating, offer more generous terms.[1] When firm characteristics and product quality are unobservable to customers at the time of purchase, this relationship could be explained through signaling, similar to the theory of warranties put forward by Grossman (1981).[2] However, when the reported measures of quality, such as customer ratings, are readily observable by customers, signaling alone cannot explain the observed heterogeneity. In our model, high-quality firms extract more surplus from consumers by offering an excessively high refund, while the opposite holds for competing low-quality firms. These distortions "at both ends" also distinguish our theory from models that explain contract heterogeneity by an efficiency rationale, e.g., as goods differ in their salvage value to firms after they are returned.[3] Below we relate our findings to the predictions offered by other recent contributions to the literature on refunds and restocking fees.

In our model, consumers hold only privately observed prior beliefs about their valuation and they learn from experimenting with the product or service. With competition, two firms with a known high or low quality are in the market. Following the approach in Shaked and Sutton (1982), consumers who ultimately derive a higher utility from the product have also a higher marginal valuation for high quality. Prior to experimenting with a product or service, consumers have only imprecise knowledge of their utility (their "true type"). We characterize an equilibrium where the market is segmented as follows (albeit our key empirical prediction holds more generally, as we show as well). Customers who, ex-ante, have a lower expected utility turn to the low-quality firm, and customers who, ex-ante, have a higher expected utility turn to the high-quality firm. Consequently, for the low-quality firm its "marginal" customer has the highest ex-ante valuation among all its customers and, therefore, the lowest marginal valuation for a higher refund, given that he is less likely to return the product (or to terminate a contract prematurely). The opposite holds for

☆ The authors thank Pascal Courty, Florian Hoffmann, and Martin Peitz for helpful comments. The paper replaces an earlier working paper entitled "Refunds as a Metering Device".

* Corresponding author.
E-mail address: inderst@finance.uni-frankfurt.de (R. Inderst).

[1] See for some evidence Bonifield et al. (2010), where this relationship is found for non-consumable goods on the internet, or Heiman et al. (2002).

[2] Cf. also theories of moral hazard such as Mann and Wissink (1990).

[3] Davis et al. (1995) document a positive correlation between the salvage value and the refund.

http://dx.doi.org/10.1016/j.ijindorg.2014.12.002
0167-7187/

the high-quality firm. There, the "marginal" customer has the lowest ex-ante valuation and, therefore, the highest marginal valuation for a higher refund.

The optimal distortion of the refund away from the efficient level is now orthogonal for both firms, albeit explained by the same mechanism. By reducing the refund below the efficient level, which is equal to the salvage value or to the cost of continuing service, the low-quality firm can extract more of the consumer surplus of "inframarginal" customers, who have low ex-ante valuation. For the high-quality firm the opposite picture emerges. To extract relatively more consumer surplus from its "inframarginal" customers, who have a higher ex-ante valuation, the high-quality firm optimally offers a refund that is excessively generous.

In both cases, the results can be understood in terms of "metering" (cf. Schmalensee, 1981), namely when we interpret the ex-ante likelihood of a return as "metering" the extent of the usage that different customers will make of the refund option. Among the customers of the high-quality firm, it is the marginal customer for whom – from an ex-ante perspective – it will be most likely that the product is returned. Among the customers of the low-quality firm, the marginal customer is instead least likely to return the product. By increasing both the refund and the price so that the marginal consumer remains indifferent, the high-quality firm reduces the consumer surplus of infra-marginal customers, as the latter are less likely to return the product. The symmetric picture applies to the low-quality firm, where a less generous refund, accompanied by a price reduction that makes the marginal consumer still indifferent, reduces the consumer surplus for the respective infra-marginal customers, as these are ex-ante more likely to return the product. In both cases, the optimal choice trades off information rent extraction with total surplus maximization, as a distortion of the refund away from the first-best choice strictly reduces the latter.

To summarize, what is key in our setting is the combination of ex-ante private information by consumers,[4], the fact that this is correlated with the likely usage that they will make of the refund option, and vertical competition. We also only consider single-product firms. [5] The role of refunds to extract (more) consumer surplus when consumers have ex-ante private information follows the sequential screening literature (cf. Courty and Li, 2000), albeit – as we discuss below in much detail – we restrict consideration to a single contract offer rather than a menu. Our key contribution to this literature is the consideration of vertical competition and the thereby obtained different implications for low- and high-quality firms.

Notably in the marketing literature, recent contributions have explored other aspects and determinants of optimal refund policies on which this paper must remain silent. Shulman et al. (2009) consider a firm with two products, which can thus still capture sales when one product is returned, and they consider as well the option of pre-purchase information provision. The latter option is particularly relevant in their setting as ex-ante consumers may differ not only in their expected valuation but also in the knowledge about their valuation.[6] Further, while in our paper consumers can learn only about a firm specific utility component, while experimenting with one product a consumer may also learn about his valuation for a competing product. Shulman et al. (2011) explore both types of learning in a setting with horizontal differentiation. Notably they contrast the findings when two differentiated products are offered by a monopolist with those when they are offered by competing firms and find that as the perceived differentiation increases, the refunds become less generous with competition and more generous with a monopolist.[7] Finally, these as well as our paper typically consider variations in the refund terms, while Che (1996) considers risk averse consumers, and Davis et al. (1995) examine instead the use of full money-back return policies. Notably, these may also give rise to extreme consumer opportunism, as analyzed for instance in Hess et al. (1996), where consumers buy and return a product for the strict purpose of free renting.

The rest of this paper is organized as follows. Section 2 introduces the baseline model. In Sections 3 and 4 we derive the results under monopoly and competition, respectively. Section 5 extends the model to the case where consumers vary in the quality of information they possess ex-ante. Section 6 offers some concluding remarks.

2. The model

2.1. Utility

We model a market for differentiated goods. For a given consumer in this market, utility depends both on a common quality measure y that differs between firms and an intrinsic fit t that differs between consumers: $u(t, y) = yt$. Importantly, we postulate that t, which is not observed by the firms, is also only partially observed by consumers before they purchase. Specifically, we stipulate that $t = \theta + \varepsilon$, where θ is observed before purchase and ε only after purchase. Here, $\theta \in \underline{\theta}, \overline{\theta}] \in \Theta$ is distributed according to $G(\theta)$ with density $g(\theta) > 0$, while $\varepsilon \in \underline{\varepsilon}, \overline{\varepsilon}]$ is distributed according to $F(\varepsilon)$ with density $f(\varepsilon) > 0$. We assume that these are drawn independently, both for a given consumer and across consumers, but we will also comment below on how our results generalize.[8] There is a mass one of consumers in the market.

2.2. Firms and contracts

We consider a model of vertical differentiation. Precisely, at most two firms operate in the market, which we denote by $i = l, h$. Firm h's intrinsic quality is given by $y_h > 0$ and firm l's is given by y_l where $0 < y_l < y_h$. That is, firm h's intrinsic quality is better than that of firm l.

Firms have constant production costs $c > 0$ and offer the following contracts. A contract specifies a sales price p together with a refund r. When a good is returned, it has the salvage value s with $0 < s < c$. Note that the salvage value is independent of the firm's quality y. For instance, we may suppose that after early return the good is no longer suitable for its primary usage. As noted in the Introduction, our model also applies to service contracts. Then, $c - s$ is the cost of initiating a contract, s is the cost of continuing to service a customer who has not terminated earlier, $p - r$ is the price for initiating the contract, and r is the additional payment required to continue the service.

Firms offer a single contract (p, r). Such a restriction may be particularly realistic with physical products. Otherwise, a firm would have to ascertain that a customer who bought under a less generous refund policy does not claim a higher refund by returning a product that was bought by another customer under a more generous refund policy. That is, the firm must ensure that product–customer matches remain

[4] We do not consider the possibility of costly pre-purchase information acquisition, as in Matthews and Persico (2007).

[5] Somewhat related, we also do not consider at all vertical issues, which would arise when the manufacturer interacted with a (possibly also multi-product) retailer. Such a channel structure approach is taken notably in Shulman et al. (2010), which allows to combine problems of forward channel management (i.e., wholesale pricing) and reverse channel structures (i.e., accepting and salvaging returns).

[6] See also Inderst and Peitz (2012) for a model with such two-dimensional ex-ante heterogeneity, albeit there the focus is on two-dimensional screening with nonlinear (service) contracts, so that there is no relationship to the issue of a refund.

[7] Kuskov and Lin (2010) consider costly information provision by firms, e.g., through free sampling that resembles a full refund policy, in a vertically differentiated setting. However, the key difference is that consumers obtain this information "for free", i.e., before making any payment to the respective seller, while in our model, as well as typically also in other models of refunds and returns, a consumer must first pay an initial price before learning about the product. In this sense, the results of Kuskov and Lin (2010) are thus more closely related to those in the literature on pre-contractual (pre-sale) information provision, e.g., Lewis and Sappington (1994); Johnson and Myatt (2006).

[8] It should be noted that the additive structure of t is by itself not restrictive. In fact, the part of t that is unknown to a consumer ex-ante (the "error term") can always be defined as $\varepsilon = t - \theta$.