Money-Back Guarantees*

Bruce McWilliams† Tridib Sharma‡ Levent Ülkü§

This version: November 2014

Abstract

We provide a framework to evaluate whether or not a seller can increase his revenue in interacting with a privately informed buyer by using money-back guarantees (MBGs). The buyer's value for the good exhibits *fit risk* and his type is multidimensional giving the probability of fit as well as the value in case of fit. We restrict attention to mechanisms that do not offer partial MBGs. We reformulate the optimal mechanism design problem and show that typically the optimal mechanism offers MBGs to some subset of types. Furthermore, choosing the optimal mechanism is tantamaount to choosing two prices: (i) a discount price at which no MBG is offered and (ii) a regular (higher) price which comes with a MBG. We also analyze two limit scenarios where private information is one-dimensional. If the seller knows the probability of fit but not its value, then MBGs are not useful. If, on the other hand, the value of fit is commonly known but its probability is buyer's private information, then MBGs can be used to extract full surplus from the buyer.

*We thank Rich McLean for useful comments.

†Department of Business Administration, ITAM, Mexico City, bruce@itam.mx
‡CIE and Department of Economics, ITAM, Mexico City, sharma@itam.mx
§CIE and Department of Economics, ITAM; Mexico City, levent.ulku@itam.mx

1 Introduction

A seller is interested in selling a product to a buyer. If the product meets the requirement of the buyer then he gets a valuation v. Otherwise, the product is of no use to him. Before purchase, he believes that the product will meet his requirement with some probability p. In such situations, the buyer faces *fit risk*. These risks are widely prevalent. An agent who has a flight from one location to another through a connecting hub, may not be able to make it to a meeting on time due to bad weather in the city where the hub is located. Ex ante, the flight is of value to him only if he is able to make it to the meeting on time. A retailer, buying inventory from a whole-seller, is able to generate profits only if he is able to sell. For him, *fit risk* arises due to uncertainty in local demand. A piece of furniture which looks good in a store may not look good in one's living room due to different lighting conditions. Buying a gift for one's spouse generates pleasure only if the spouse were to like the gift. In this paper we study how money-back guarantees (MBGs) facilitate transactions in the presence of fit risk.

An MBG is a guarantee that refunds the full price to the buyer if the buyer were to return the product.[1] Heiman et al. (2002) explicitly observe that retailers and manufacturers provide MBGs to help resolve fit risk. Product returns are an enormous phenomenon in the US market, exceeding \$100 billion annually in the US (Stock et al. 2002). Over 95% of retailers in a survey in Illinois offering some form of MBG (Sales and Marketing Management, 1994). The impact of MBGs on purchase decisions can vary greatly by product category, consumer type and distribution channel. For example, Anderson et al. (2009) found that for one catalog retailer, average product return rates were 23%, 14% and 29% for women's tops, men's tops and women's footwear respectively, and that offering MBGs increased demand by 16%, 9% and 53% respectively in these same categories. For the computer electronics industry, where between 11% and 20% of products are returned, Sprague et al. (2007) found that only 5% of returned products were defective. Two-thirds of customers returned their computers because "they did not meet expectations," while a quarter of the computers were returned because of "buyer's remorse."

Our concern in this paper is not for the 5% of returns made due to product defects–

[1] Even though partial money-back guarantees are easy to conceptualize and potentially interesting to analyze, our restriction to full money-back guarantees is rooted in FTC guidelines. In some cases, retailers charge a small restocking fee which we shall ignore.

these are generally covered by warranties and consumes have some legal protection for these occurrences. Rather, our focus is on the remaining 95% of returns which are subject to voluntary retailer discretion in the US. For example, in California retailers can avoid accepting returns of non-defective items from customers if they explicitly notify customers in a prominent location that products cannot be returned, or if they indicate "all sales final" on items not covered by MBGs. In our context, MBGs reflect the fact that in offering this guarantee the retailer publicly agrees to fully refund the purchase price to a dissatisfied customer even when the product is not defective.

A common theoretical explanation for MBGs is that they help signal quality. In these models, consumers are uncertain about quality, and the better quality firm gains consumer trust by offering costly MBGs as a show of confidence in their own product (e.g., Mann and Wissink, 1990; Moorthy and Srinavasan, 1995). The other traditional explanation for MBGs is that it provides insurance to risk averse consumers that do not know whether they will be satisfied with the product or not. In this case a risk neutral seller can gain from offering the MBG as insurance to the risk averse consumer (e.g., Heal, 1977; Che, 1996). More recently, researchers have suggested that MBGs can arise when the seller has a higher salvage value for the returned product than does the dissatisfied customer (Davis et al., 1995 and McWilliams, 2012). In this framework, sellers only offer MBGs if the difference in the salvage value between the seller and customer is greater than the sum of the costs they incur in returning the product. This can explain the fact that some products with low marginal costs (and therefore low salvage value for the seller), such as computer software, have shorter MBGs (in terms of the time period under which it is to be returned) than products with higher marginal costs such as appliances.

We propose that MBGs help screen buyer types. To eliminate salvage value as a possible explanation, we assume zero production and zero return costs, while modeling consumers as risk neutral eliminates the standard insurance motivation. Whereas the signaling models assume that the seller has private information (about the quality of the good) and consumers are partially informed, we assume the reverse. In our model, consumers are privately informed while the seller is partially informed.

We provide three results regarding the use of MBGs as part of an optimal mechanism. The first result is derived in an environment where the buyer is privately informed about both him valuation v, in the case of a fit, as well as the probability p of fit. In this environment, under very mild restrictions, we show that it is optimal for a seller to always offer some MBG. In the majority of product purchasing environments where retailers face a heterogeneous customer base, neither the probability of fit nor the valuation of the good is common knowledge. Here the standard practice in the US is for retailers to offer customers MBGs on a regular basis, and eliminate the MBG during sales events where prices are reduced. Big retailers, on the other hand, almost always offer MBGs. Our first result sheds light on both these business practices. The second result pertains to an environment where v is commonly known and p is the buyer's private information. Interactions between whole-sellers and retailers could probably be captured by such models. Many whole-sellers determine both retail and whole-sale prices. Thus, the profit per unit of the product would be commonly known. The retailer, however, may be better informed about local demand fluctuations. In this environment, we show that it is optimal for the seller to charge one single price and offer full MBG.[2] In the third scenario we show that, when p is commonly known but v is privately known to the buyer then it is optimal not to offer MBGs. This environment could well exist in the sale of customized products and services, where MBGs are usually not offered.[3]

There are several papers which, like us, highlight the pure screening role of MBGs. While all of them take a mechanism design approach, they differ in the details of the environments, for example in allocation sets and the buyer type spaces. In essence then, these papers address different kinds of MBGs appropriate for different product markets. All these papers, including ours, are modifications of the well-known problem introduced by Myerson (1981).

In Myerson's model an uninformed seller has to decide on whether or not to provide a good to a buyer who is privately informed about him valuation (a one dimensional type). An allocation indicates whether the good changes hands. The seller also has to decide on a price. Both these decisions are contingent on the buyer's type. Since type is unknown, the seller's decision needs to be such that it is in the interest of the buyer to reveal him true type. That is, the sellers decision has to satisfy incentive compatibility. The decision rule also needs to ensure that the buyer willingly trades. That is, the decision rule has

[2] That manufacturers commonly provide MBGs in such contexts is documented in Padmanabhan and Png (1997). They propose that offering MBGs to retailers ensure that sufficient inventory is carried.

[3] In standard services provided by hotels, fit risk should be common knowledge. Though MBGs are usually not provided by hotels, Hampton Inn hotels is an exception. They explicitly state that a customer can ask for their money back if they are unsatisfied with the hotel experience for whatever reason.

to satisfy individual rationality. In the class of all such rules, the optimal rule turns out to be quite simple. All types above a cut-off are offered the good and types below aren't. All types who are offered the good pays a price equal to the cut-off and types who are not offered the good pay nothing.

Matthews and Moore (1987), though related, is more about warranties than MBGs. The risk averse buyer's type is one dimensional. Provision of the good along with different quality warranties make the allocation set multi-dimensional. Warranties serve as insurance and different kinds of warranties are used to screen the types.

Courty and Li (2000), consider a model where the buyer does not know the valuation of the product till after it is bought. Buyer's types are distributions over an interval of valuations. Thus types are unidimensional but complex. They allow buyers to have only two types. One type's distribution stochastically dominates the other and hence is more informed. Alternately, both distributions have the same mean but different variances. The allocation set consists of the provision of the good and different amounts of money refund. The more informed buyer type gets a refund equal to the sellers fixed cost of production (zero in our case), while the less informed buyer can get a refund greater than, equal to or less than the cost (depending on the functional form of the distribution). In the second part of the paper, like us, they consider a continuum of types. The allocation set, however, does not have any refunds. To reduce complexity, they only consider a class of examples where the distribution and the realized value are a function of a common parameter. Thus, they essentially reduce their analyses to one dimensional (albeit complex) types. Screening is achieved by providing or not providing the good (cut-offs) and charging different prices.

Matthews and Persico (2005) show that refunds (above cost) could only arise if the seller is a monopolist. That too, only in particular situations. There are two types of buyers, those who know their valuations for sure and those who do not. The latter type, like in Courty and Li have a distribution over their valuation. But this distribution is common knowledge. Buyers who know their valuation, of course do not need refunds, the latter type may. The relative proportion of these two types then determine whether or not refunds are provided.

We believe that our paper is closest to that of Myerson. Unlike in Myerson, in our first environment, the buyer is privately informed about his two dimensional type. Our type space is constructed to capture fit risk. If the product fits, the buyer gets a value v.

Otherwise he gets a value 0. The product fits with probability p. The tuple $(v, p) \in [0, 1]^2$ is privately known to the buyer. The good is deterministically provided (or not) and a full refund (MBG) is deterministically offered (or not) as a function of reported types. Payments, as usual, are a function of reported types. Incentive compatibility, individual rationality and the restriction that refunds equal the price paid, partitions our buyer's type space into three parts. Buyers with low valuations and high fit risk do not receive the good and pay no price. Buyers with high valuations and low fit risk receive the good with no MBG at some constant price. Buyers with high valuation and high fit risk receive both the good and MBG. The price is equal to the MBG and is (weakly) higher than the price charged to buyers who are not offered the MBG.

Explicitly deriving the optimal mechanism in two (or more) dimensional type spaces is non-trivial (Armstrong, 1996; Rochet and Chone, 1998; Manelli and Vincent, 2007). However, we show under very mild restrictions on the seller's prior, that *in general* the optimal mechanism will always offer an MBG to some set of buyer types. We also offer an analytically tractable reformulation of the optimal mechanism design problem, in which the seller chooses, depending on his prior regarding type distribution, two parameters which define the mechanism. The two parameters are essentially prices: a discount price which goes without a MBG and a high price with which a full MBG is offered. Using this reformulation, we provide examples of optimal mechanisms when the seller's prior over the buyer's type space is uniform. We then go on to show that restricting the dimensionality of the type space to two is without loss of generality. Our qualitative results survive when the buyer is privately informed about his fit risk p, his valuation when there is a fit v_H and his valuation when there is no fit v_L. We only require that v_H is not less than v_L.

As mentioned earlier, we provide two additional results in the two dimensional setting. When fit risk is common knowledge but the buyer is privately informed about him valuation, then the optimal mechanism is the same as that in Myerson. When fit risk is privately known and the buyer's valuation is common knowledge then the optimal mechanism extracts all the surplus from the buyer. This result is interesting in its own right. To our knowledge, no other model achieves full surplus extraction in the case of a single buyer.

2 Environment and main results

A seller of an indivisible good is interacting with a privately informed buyer. The seller has no cost. The good has a random value to the buyer in the sense of exhibiting the following form of fit risk: with probability p it fits the buyer's needs and its value is $v > 0$, otherwise, with probability $1 - p$, there is no fit and its value is 0. Whether the good fits or not is only observed by the buyer after purchase. However the pair (p, v) is buyer's private information at the time of interaction, i.e., his type. The type takes values in $T = (0, 1)^2$ with a strictly positive density f.

The interaction results in an *outcome* (α, γ, π) where $\alpha \in \{0, 1\}$ indicates sale or no sale, $\gamma \geq 0$ is a money-back guarantee (MBG), and $\pi \geq 0$ is the price. If the good does not change hands, i.e., if $\alpha = 0$, then there is no payment and no MBG, i.e., $\gamma = \pi = 0$. Hence the outcome belongs to the set

$$C = \{(\alpha, \gamma, \pi) \in \{0, 1\} \times \Re_+ \times \Re_+ : \alpha = 0 \Rightarrow \gamma = \pi = 0\}.$$

Given $(\alpha, \gamma, \pi) \in C$, the payoffs are determined as follows. Both parties receive zero payoff if $\alpha = 0$. If $\alpha = 1$, the buyer pays the seller π. Next he observes the fit of the good and returns the good if and only if the MBG exceeds its value. Returning the good is costless. The resulting expected payoff of the buyer, depending on the outcome (α, γ, π) and type (p, v), is $\alpha(p \max\{v, \gamma\} + (1-p)\gamma) - \pi$. The seller receives the payment π and pays the MBG γ if the good is returned. Hence his expected payoff is given by $\pi - (1 - p)\alpha\gamma$ if $\gamma \leq v$ and $\pi - \alpha\gamma$ if $\gamma > v$. Note that if there is no MBG, i.e., if $\gamma = 0$, we specialize to the textbook scenario with the buyer value equal to the expectation pv.

We are interested in using mechanism design to analyze the role of MBGs in the optimal mechanism for the seller. A *mechanism* maps types (p, v) into contracting outcomes (α, γ, π). Abusing notation, we will denote a mechanism by a triple of functions $(\alpha, \gamma, \pi) : T \to C$. The following two conditions are the classical incentive constraints in mechanism design.

F1: Incentive compatibility (IC): for any two types (p, v) and (p', v')

$$\begin{aligned} &\alpha(p, v)(p \max\{v, \gamma(p, v)\} + (1 - p)\gamma(p, v)) - \pi(p, v) \\ &\qquad \geq \alpha(p', v')(p \max\{v, \gamma(p', v')\} + (1 - p)\gamma(p', v')) - \pi(p', v'). \end{aligned}$$

F2: Individual rationality (IR): for any type (p, v)

$$\alpha(p, v)(p \max\{v, \gamma(p, v)\} + (1 - p)\gamma(p, v)) - \pi(p, v) \geq 0.$$

These two conditions say that (1) the buyer cannot gain by misreporting to the mechanism, and (2) the truthful report earns the buyer a payoff at least equal to his outside option, which we assume is zero. We will next introduce a key condition to rule out partial MBGs.

F3: For any type (p, v), $\gamma(p, v) > 0$ implies $\gamma(p, v) = \pi(p, v)$.

F3 requires any positive MBG to be a full reimbursement of the price of the good. Note, importantly, that F3 does not impose that a MBG be offered by the seller. Our imposition of F3 is inspired by the following guideline of the Federal Trade Commission:

> A seller... should use the term "Money Back Guarantee"... only if the seller refunds the *full* purchase price ... at the purchaser's request. (Our italics; citation needed.)

Despite having a strong foundation in actual business practice, F3 is an important restriction in mechanism design. In principle it may be desirable for the seller to offer only a partial MBG. The possibility of partial MBGs brings about a host of interesting potential policies for the seller and makes mechanism design significantly more challenging.

For any mechanism (α, γ, π) which satisfies F1, F2 and F3, and any type (p, v) let $R^{(\alpha,\gamma,\pi)}(p, v)$ be the seller's ex post payoff when the buyer type is (p, v), i.e.,

$$R^{(\alpha,\gamma,\pi)}(p, v) = \begin{cases} \pi(p, v) - (1 - p)\alpha(p, v)\gamma(p, v) & \text{if } \gamma(p, v) \leq v, \\ \pi(p, v) - \alpha(p, v)\gamma(p, v) & \text{if } \gamma(p, v) > v. \end{cases}$$

We are interested in the following *optimal mechanism design problem*:

$$\max_{(\alpha,\gamma,\pi):T \to C} \int_0^1 \int_0^1 R^{(\alpha,\gamma,\pi)}(p, v) f(p, v) dp dv$$
$$\text{s.t. F1, F2 and F3.} \tag{P1}$$

Problem P1 is that of maximizing the expected payoff of the seller by choosing a mechanism which satisfies the three feasibility constraints above. We will call a mechanism *optimal* if it solves P1.

In principle, the seller could offer a high enough MBG which will induce the buyer to return the good regardless of its fit. If this is true for some type (p, v), under F3, the buyer's payoff is $\gamma(p, v) - \pi(p, v) = 0$. Correspondingly, the seller's ex post payoff from interacting with this type is $\pi(p, v) - \gamma(p, v) = 0$ as well. Our first observation is that the seller will never find offering such a high MBG profitable.

Lemma 1 *Suppose (α, γ, π) satisfies F1, F2 and F3. If $\gamma(p', v') > v'$ for some (p', v'), then there exists a mechanism $(\tilde{\alpha}, \tilde{\gamma}, \tilde{\pi})$ which satisfies F1, F2, F3 such that for every (p, v), $\tilde{\gamma}(p, v) \leq v$ and $R^{(\alpha,\gamma,\pi)}(p, v) = R^{(\tilde{\alpha},\tilde{\gamma},\tilde{\pi})}(p, v)$.*

Proof. Fix (α, γ, π) in satisfaction of F1, F2 and F3. Define $(\tilde{\alpha}, \tilde{\gamma}, \tilde{\pi})$ as follows:

$$(\tilde{\alpha}(p, v), \tilde{\gamma}(p, v), \tilde{\pi}(p, v)) = \begin{cases} (\alpha(p, v), \gamma(p, v), \pi(p, v)) & \text{if } \gamma(p, v) \leq v, \\ (0, 0, 0) & \text{if } \gamma(p, v) > v. \end{cases}$$

Clearly $\tilde{\gamma}(p, v) \leq v$ for all (p, v) and $(\tilde{\alpha}, \tilde{\gamma}, \tilde{\pi})$ satisfies F3. Furthermore $R^{(\tilde{\alpha},\tilde{\gamma},\tilde{\pi})}(p, v) = R^{(\alpha,\gamma,\pi)}(p, v)$ for all (p, v) as whenever $\gamma(p, v) > v$, $R^{(\alpha,\gamma,\pi)}(p, v) = 0$ by F3.

Note that for every (p, v)

$$\begin{aligned} &\tilde{\alpha}(p, v)(p \max\{v, \tilde{\gamma}(p, v)\} + (1 - p)\tilde{\gamma}(p, v)) - \pi(p, v) \\ &\qquad = \alpha(p, v)(p \max\{v, \gamma(p, v)\} + (1 - p)\gamma(p, v)) - \pi(p, v). \end{aligned}$$

This is trivially true if $\gamma(p, v) \leq v$ as $(\tilde{\alpha}(p, v), \tilde{\gamma}(p, v), \tilde{\pi}(p, v)) = (\alpha(p, v), \gamma(p, v), \pi(p, v))$. $\gamma(p, v) > v$, on the other hand both sides are zero. F2 directly follows from this observation.

Also note that if $(p', v') \neq (p, v)$

$$\begin{aligned} &\alpha(p', v')(p \max\{v, \gamma(p', v')\} + (1 - p)\gamma(p', v')) - \pi(p', v') \\ &\qquad \geq \tilde{\alpha}(p', v')(p \max\{v, \tilde{\gamma}(p', v')\} + (1 - p)\tilde{\gamma}(p', v')) - \pi(p', v'). \end{aligned}$$

As before this is trivially true if $\gamma(p', v') \leq v'$. If $\gamma(p', v') > v'$, on the other hand, the left-hand side is $p(\max\{v, \gamma(p', v')\} - \gamma(p', v')) \geq 0$ while the right-hand side is 0. Combining the inequalities in the last two displays and using the hypothesis that (α, γ, π) satisfies F1, we conclude that $(\tilde{\alpha}, \tilde{\gamma}, \tilde{\pi})$ satisfies F1 as well. ■

Let us formulate the condition that the good be returned to the seller for a MBG only if it is not a fit, as a feasibility condition.

F4: For any type (p, v), $\gamma(p, v) \leq v$.

Restricting attention to mechanisms satisfying this condition simplifies the seller's payoff and, by Lemma 1, is without loss of generality in solving for the optimal mechanism. In other words, in order to find the optimal mechanism which solves P1 the seller need only solve

$$\max_{(\alpha,\gamma,\pi):T\to C} \int_0^1 \int_0^1 [\pi(p, v) - (1-p)\alpha(p, v)\gamma(p, v)] f(p, v) dp dv$$
$$\text{s.t. F1, F2, F3 and F4.} \tag{P2}$$

Next we introduce a class of mechanisms which are feasible in P2.

Definition 1 *Let $0 \leq k \leq m \leq 1$. A mechanism (α, γ, π) is a (k, m) mechanism if*

$$(\alpha(p, v), \gamma(p, v), \pi(p, v)) = \begin{cases} (1, m, m) & \text{if } v \geq m \text{ and } p \leq \frac{k}{m}, \\ (1, 0, k) & \text{if } p > \frac{k}{m} \text{ and } pv \geq k, \\ (0, 0, 0) & \text{otherwise.} \end{cases}$$

The (k, m) class contains three kinds of mechanisms as illustrated below, depending on whether MBGs are offered at all, or offered to some types and not to others. If $0 < k < m = 1$, as in the first diagram, the seller does not offer MBGs to any type. The mechanism allocates the good to the agent if his expected value $pv \geq k$ at the price $\pi = k$. This is exactly the case where multidimensionality of the consumer's type (p, v) is inconsequential and the mechanism divides different types with respect to the product

pv.

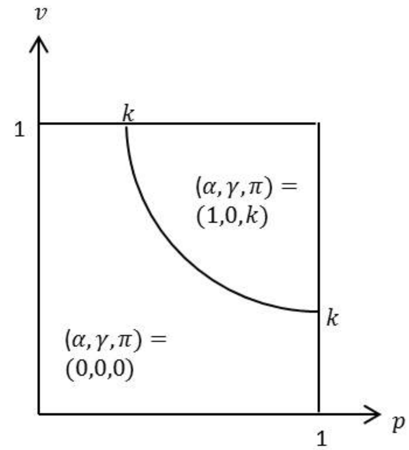

Figure 1
$0 < k < m = 1$

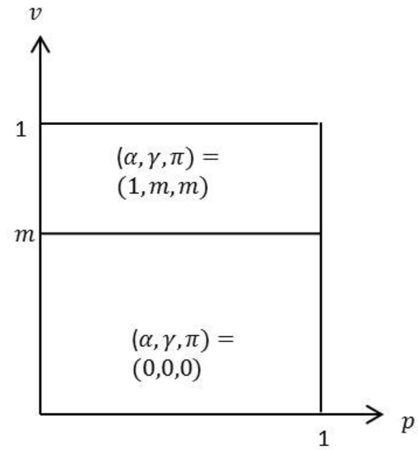

Figure 2
$0 < k = m < 1$

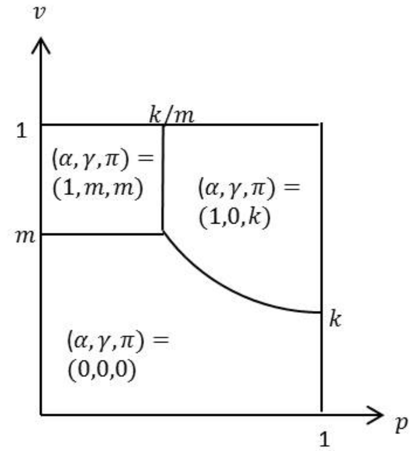

Figure 3
$0 < k < m < 1$

At the other extreme are mechanisms which give MBGs to all types who receive the good. These mechanisms have $m = k$. A typical such mechanism is given in the second diagram. Mixing of these two policies is also feasible, as in the third diagram, by choosing $0 < k < m < 1$. In this case, the mechanism allocates the good to some types at a discount price k with no MBG and to other types at a higher price m with a full MBG.

It is straightforward to check that any (k, m) mechanism satisfies conditions F1-F4 and is there fore feasible in problem P2. Next we will show that if a mechanism is feasible in problem P2 then it is "almost" a (k, m) mechanism. First we record a useful consequence of conditions F1-F4.

Lemma 2 *If (α, γ, π) satisfies F1-F4, then there exists $m \in (0, 1)$ such that if $\gamma(p, v) > 0$, then $\gamma(p, v) = m$.*

Proof. Suppose, towards a contradiction, that for two distinct types (p, v) and (p', v'), $0 < \gamma(p', v') < \gamma(p, v)$. Then $\alpha(p', v') = \alpha(p, v) = 1$ and the payoff to the (p, v) type from a truthful report is

$$pv + (1 - p)\gamma(p, v) - \pi(p, v) = p(v - \gamma(p, v))$$

where we use F3 in substituting $\gamma(p, v)$ for $\pi(p, v)$. If, instead, the (p, v) type reports (p', v'), then his payoff would be

$$pv + (1 - p)\gamma(p', v') - \pi(p', v') = p(v - \gamma(p', v')).$$

F1 requires $p(v-\gamma(p,v)) \geq p(v-\gamma(p',v'))$ which implies, since $p>0$, $\gamma(p',v') \geq \gamma(p,v)$, a contradiction. ■

For any mechanism (α,γ,π) satisfying F1-F4, let $(T_1^{(\alpha,\gamma,\pi)}, T_2^{(\alpha,\gamma,\pi)}, T_3^{(\alpha,\gamma,\pi)})$ be the partition of the type space defined by

$$
\begin{aligned}
(p,v) &\in T_1^{(\alpha,\gamma,\pi)} \Leftrightarrow \alpha(p,v)=1 \text{ and } \gamma(p,v)>0, \\
(p,v) &\in T_2^{(\alpha,\gamma,\pi)} \Leftrightarrow \alpha(p,v)=1 \text{ and } \gamma(p,v)=0 \\
(p,v) &\in T_3^{(\alpha,\gamma,\pi)} \Leftrightarrow \alpha(p,v)=0.
\end{aligned}
$$

Note that all types in $T_1^{(\alpha,\gamma,\pi)}$ receive the same MBG by Lemma 2. We will say that two mechanisms are *almost identical* if they generate the same partitions, except perhaps at the boundaries. Given our assumption that types have a strictly positive density, almost identical mechanisms earn the seller the same revenue as they differ only on a set of zero measure.

Proposition 1 *Any mechanism which satisfies F1-F4 is almost identical to some* (k,m) *mechanism.*

Proof. Let (α,γ,π) satisfy F1-F4. The proof relies on Lemma 2 as well as the following three observations.

Claim 1: If $\alpha(p,v)=1$, then $\alpha(p',v')=1$ for all (p',v') such that $p'>p$ and $v'>v$.

Proof of Claim 1: Suppose that $\alpha(p,v)=1$ but $\alpha(p',v')=0$ for some (p',v') such that $p'>p$ and $v'>v$. If $\gamma(p,v)=0$, then

$$0 \geq p'v' - \pi(p,v) > pv - \pi(p,v) \geq 0$$

where the weak inequalities follow from incentive compatibility. This is clearly impossible. If $\gamma(p,v)=m>0$, on the other hand, the impossibility follows similarly from incentive compatibility:

$$0 \geq p'(v'-m) > p(v-m) \geq 0.$$

Claim 2: If $\alpha(p,v)=\alpha(p',v')=1$ and $p'<p$, then $\gamma(p',v') \geq \gamma(p,v)$.

Proof of Claim 2: Suppose that $\alpha(p,v)=\alpha(p',v')=1$, $p'<p$ but $\gamma(p',v')<\gamma(p,v)$. Then, by Lemma 2, for some $m>0$ $\gamma(p,v)=m$ and $\gamma(p',v')=0$. Incentive compatibility gives

$$p(v-m) \geq pv - \pi(p',v') \text{ and } p'v' - \pi(p',v') \geq p'(v'-m).$$

Rearranging we get $pm \le \pi(p', v') \le p'm$, which is an impossibility since $m > 0$.

Claim 3: If $\gamma(p, v) > 0$, then $\gamma(p', v') = \gamma(p, v)$ for all $(p', v') \in (0, p) \times (\gamma(p, v), 1)$.

Proof of Claim 3: Suppose $\gamma(p, v) = m > 0$, $p' < p$ and $v' > m$. We will first show that $\alpha(p', v') = 1$. If not, incentive compatibility implies $0 \ge p'(v' - m)$, an impossibility. Now suppose $\gamma(p', v') = 0$. The incentive compatibility conditions are exactly as in the proof of Claim 2 and the same contradiction follows.

We can now go back to the proof of Proposition 1. If one of the sets in the partition $(T_1^{(\alpha,\gamma,\pi)}, T_2^{(\alpha,\gamma,\pi)}, T_3^{(\alpha,\gamma,\pi)})$ is empty, then the result follows straightforwardly. We will deal here with the case in which all three sets are nonempty. By Lemma 2, there exists $m^* > 0$ such that for any $(p, v) \in T_3^{(\alpha,\gamma,\pi)}$, $\gamma(p, v) = m^*$. By Claim 3 above, $\sup\{p : \gamma(p, v) = 1\} = \sup\{p : \gamma(p, v') = 1\}$ for any $v, v' > m^*$. Let p^* be this supremum and $k^* = m^* p^*$. By Claim 1 above, $\alpha(p, v) = 1$ if $p > p^*$ and $v > m$. It follows that $\gamma(p, v) = 0$ for such (p, v) since $p > p^*$. Clearly $k^* = \inf\{p'v' : p > p^* \text{ and } v > m\}$ and if for some (p, v) such that $v \le m$ and $pv > k^*$, $\alpha(p, v) = 0$, incentive compatibility fails. Hence $(\alpha(p, v), \gamma(p, v)) = (1, 0)$ for all (p, v) such that $p > p^*$ and $pv > k^*$. By incentive compatibility and individual rationality $\pi(p, v) = k^*$ for any such type. Hence (α, γ, π) is a (k, m) mechanism with $(k, m) = (k^*, m^*)$. ■

Hence, the seller need only find the optimal one among the (k, m) mechanisms, as we record in the following corollary.

Corollary 1 *If the pair* (k^*, m^*) *solves*

$$\max_{k,m} \int_m^1 \int_0^{k/m} mpf(p, v)dpdv + \int_m^1 \int_{k/m}^1 kf(p, v)dpdv + \int_k^m \int_{k/v}^1 kf(p, v)dpdv$$
$$s.t.\ 0 \le k \le m \le 1 \tag{P3}$$

then the (k^*, m^*) *mechanism solves the optimal mechanism design problem P1.*

The objective in the reformulated problem P3 in Corollary 1 is precisely the expected payoff of the seller at a (k, m) mechanism. The first double-integral is over all types which receive the good at the price m and together with the option of returning the good for the MBG m. The seller's revenue at any such type is $m - (1 - p)m = pm$. The second and third double-integrals give the seller's expected payoff over all types which receive the good at a discount k but without the MBG. Note that if $m = 1$, the objective becomes $\int_k^1 \int_{k/v}^1 kf(p, v)dpdv$, the revenue in the mechanism which involves no MBGs (Figure 1 above), and if $m = k$, the objective becomes $\int_m^1 \int_0^1 mpf(p, v)dpdv$, the revenue in the mechanism which gives MBGs to all types who receive the good (Figure 2 above).

Before we exhibit the use of this result in computing an optimal mechanism in a specific example, we will show that under general conditions, the optimal mechanism contains MBGs. In other words, the solution to the reformulated problem P3 has $m < 1$.

Proposition 2 *The optimal mechanism offers money-back guarantee to some types if one of the following two conditions holds:*

1. *p and v are independently distributed.*

2. *The density f is continuously differentiable.*

Proof. For any $k \in (0, 1)$, consider the $(k, 1)$ mechanism. This mechanism offers no MBGs. We will show that if $m \in (k, 1)$ is sufficiently close to 1, then the seller's expected payoff is larger in the (k, m) mechanism than it is in the $(k, 1)$ mechanism.

Switching from the $(k, 1)$ mechanism to a (k, m) mechanism entails a loss of expected revenue for all types (p, v) such that $p \in (\frac{v}{m}, \frac{k}{m})$ and $v \in (m, 1)$. At any such type the $(k, 1)$ mechanism earns the seller k, the price of the good, whereas the $(k, 1)$ mechanism brings the expected revenue mp, difference between price m and the expected MBG payment $(1-p)m$ back to the buyer. Note $mp < k$ for this range of p. The benefit from said switch occurs at types (p, v) such that $p \in (0, \frac{v}{m})$, $v \in (m, 1)$. The $(k, 1)$ mechanism does not serve these types. The (k, m) mechanism serves these types with the MBG m and earns the seller mp in expectation. Hence the switch is profitable if

$$E[mp|p \in (0, \frac{k}{m}), \, v \in (m, 1)] > E[k|p \in (\frac{v}{m}, \frac{k}{m}), \, v \in (m, 1)].$$

To establish that this is the case, we will show that for some $m \in (k, 1)$

$$\frac{\Pr\{p \in (\frac{v}{m}, \frac{k}{m}) \text{ and } v \in (m, 1)\}}{E[p|p \in (0, \frac{k}{m}), \, v \in (m, 1)]} < \frac{m}{k}.$$

Since the numerator in the left-hand side converges to 0 as m goes to 1, and since the right-hand side is larger than 1, and it suffices if the denominator of the left-hand side has positive limit, i.e.,

$$\lim_{m\to 1} E[p|p\in(0,\frac{k}{m}),\ v\in(m,1)] > 0.$$

Taking the conditional expectation

$$E[p|p\in(0,\frac{k}{m}),\ v\in(m,1)] = \frac{\int_m^1\int_0^{\frac{k}{m}} pf(p,v)dpdv}{\int_m^1\int_0^{\frac{k}{m}} f(p,v)dpdv}.$$

Suppose that p and v are independently distributed with strictly positive densities f_p and f_v respectively. The conditional expectation becomes

$$\frac{\int_m^1\int_0^{\frac{k}{m}} pf(p,v)dpdv}{\int_m^1\int_0^{\frac{k}{m}} f(p,v)dpdv} = \frac{\int_m^1 f_v(v)dv\int_0^{\frac{k}{m}} pf_p(p)dp}{\int_m^1 f_v(v)dv\int_0^{\frac{k}{m}} f_p(p)dp} = \frac{\int_0^{\frac{k}{m}} pf_p(p)dp}{\int_0^{\frac{k}{m}} f_p(p)dp}$$

and

$$\lim_{m\to 1} E[p|p\in(0,\frac{k}{m}),\ v\in(m,1)] = \frac{\int_0^k pf_p(p)dp}{\int_0^k f_p(p)dp} > 0$$

by the positivity of the densities and the number k, as we wanted to show.

If f and p are not independent, we apply L'Hopital's rule:

$$\lim_{m\to 1}\frac{\int_m^1\int_0^{\frac{k}{m}} pf(p,v)dpdv}{\int_m^1\int_0^{\frac{k}{m}} f(p,v)dpdv} = \lim_{m\to 1}\frac{\frac{d}{dm}\left[\int_m^1\int_0^{\frac{k}{m}} pf(p,v)dpdv\right]}{\frac{d}{dm}\left[\int_m^1\int_0^{\frac{k}{m}} f(p,v)dpdv\right]}.$$

Define

$$\begin{aligned} H(m,v) &= \int_0^{\frac{k}{m}} pf(p,v)dp, \text{ and} \\ G(m,v) &= \int_0^{\frac{k}{m}} f(p,v)dp. \end{aligned}$$

Since f is continuously differentiable, so are the integrands in these expressions and we can use Leibnitz Theorem as follows:

$$\begin{aligned}
\lim_{m\to 1}\frac{\frac{d}{dm}\left[\int_m^1 H(m,v)dv\right]}{\frac{d}{dm}\left[\int_m^1 G(m,v)dv\right]} &= \lim_{m\to 1}\frac{-H(m,m)+\int_m^1 \frac{\partial}{\partial m}H(m,v)dv}{-G(m,m)+\int_m^1 \frac{\partial}{\partial m}G(m,v)dv} \\
&= \lim_{m\to 1}\frac{-H(m,m)+\int_m^1 \frac{k}{m}f(\frac{k}{m},v)dv}{-G(m,m)+\int_m^1 f(\frac{k}{m},v)dv} \\
&= \lim_{m\to 1}\frac{H(m,m)}{G(m,m)}.
\end{aligned}$$

Now $m \mapsto H(m,m)$ and $m \mapsto G(m,m)$ are continuous because f is so. Hence

$$\lim_{m\to 1}\frac{H(m,m)}{G(m,m)} = \frac{H(1,1)}{G(1,1)} = \frac{\int_0^k xf(x,1)dx}{\int_0^k f(x,1)dx} = \Pr\{p|p<k \text{ and } v=1\} > 0,$$

which is what we needed to show. ■

In the simple scenario studied in the following example, all types who receive the good are also offered the MBG, in other words $k = m < 1$.

Example 1 Suppose that p and v are independently and uniformly distributed. The expected payoff of the seller in a (k,m) mechanism is

$$\int_m^1\int_0^{k/m} mpdpdv + \int_m^1\int_{k/m}^1 kdpdv + \int_k^m\int_{k/v}^1 kdpdv.$$

Maximizing the expression with respect to (k,m) we find that the optimal mechanism is the (k,m) mechanism with $k = m = \frac{1}{2}$. Hence the optimal mechanism offers MBGs to all types who receive the good at the price of $\frac{1}{2}$. The corresponding expected revenue of the seller is $\frac{1}{8}$.

3 Limit cases: one-dimensional private information

We next analyze two special cases in which the seller knows one of the two dimensions of the buyer's type (p,v). In these limit cases, without resorting to an analog of F4, we can more directly analyze the optimal mechanism design problems. Surprisingly, we find that F3 is trivially satisfied by the optimal mechanisms in both cases, but the nature of optimal mechanisms are quite different.

Scenario 1: v is private information and p is common knowledge. Suppose that v is buyer's private information, i.e., him type, taking values in $(0, 1)$, while $p \in (0, 1)$ is common knowledge. The following definitions are direct adaptations of those given for the multidimensional model earlier and they will apply, with the appropriate domain modification, in the second scenario below as well. A *mechanism* is a map (α, γ, π) : $(0, 1) \to C$ associating an outcome $(\alpha(v), \gamma(v), \pi(v))$ with every type v. A mechanism (α, γ, π) is *incentive compatible* if for every $v, v' \in (0, 1)$, $\alpha(v)[pv + (1 - p)\gamma(v)] - \pi(v) \geq \alpha(v')[pv+(1-p)\gamma(v')]-\pi(v')$. A mechanism (α, γ, π) is *individually rational* if for every $v \in (0, 1)$, $\alpha(v)[pv+(1-p)\gamma(v)]-\pi(v) \geq 0$. If a mechanism (α, γ, π) is incentive compatible and individually rational, it earns the seller the expected profit $\int_0^1 [\pi(v)-(1-p)\alpha(v)\gamma(v)]f(v)dv$ where f is the distribution of v. The revelation principle (Myerson, 1981) that the profit maximizing selling strategy is given by the mechanism which maximizes this expected payoff within the class of incentive compatible and individually rational mechanisms. Hence the optimal mechanism design in this scenario is :

$$\begin{aligned}&\max_{(\alpha,\gamma,\pi)} \int_0^1 [\pi(v) - (1 - p)\alpha(v)\gamma(v)]f(v)dv \\ &\text{s.t. IC and IR.}\end{aligned} \tag{P4}$$

The following result indicates that in this scenario, the seller has no incentive to use MBGs.

Proposition 3 *There is a solution to problem P4 which involves no money-back guarantees.*

Proof. Let (α, γ, π) be incentive compatible and individually rational with $\gamma \neq 0$, i.e., some type receives MBG. Consider the alternate mechanism (α, γ', π') where

$$\begin{aligned}\gamma'(v) &= 0, \text{ and} \\ \pi'(v) &= \pi(v) - (1 - p)\alpha(v)\gamma(v)\end{aligned}$$

for all v. Suppose that the buyer's type is v and he reports v' to the mechanism (α, γ', π'). His payoff is

$$\begin{aligned}\alpha(v')pv - \pi'(v') &= \alpha(v')pv - \pi(v') + (1 - p)\alpha(v')\gamma(v') \\ &= \alpha(v')[pv + (1 - p)\gamma(v')] - \pi(v')\end{aligned}$$

which is exactly his payoff from reporting v' to the mechanism (α, γ, π) when his type is v. Hence (α, γ', π') is incentive compatible and individually rational because (α, γ, π) is so. Furthermore the two mechanisms generate the same revenue for the seller ex post at every type v. Hence in order to solve P4 it suffices to maximize seller's expected profit by selecting an incentive compatible and individually rational mechanism from among those that involve no MBGs. ■

The intuition behind Proposition 3 is that the part of buyer's payoff from a nonzero MBG, $\alpha(1-p)\gamma$, is type-independent, just like his utility for money in this quasilinear framework. Hence the seller can substitute a strictly positive γ with a lower π, without changing incentive properties of the mechanism, while keeping his ex post revenue constant.

Scenario 2: p is private information and v is common knowledge Suppose now that the buyer's type is p which takes values in $(0,1)$, while $v \in (0,1)$ is common knowledge between the buyer and the seller. The definition of a mechanism remains the same as above, except that the argument of α, γ and π is p rather than v. The optimal mechanism design problem is

$$\begin{aligned} &\max_{(\alpha,\gamma,\pi)} \int_0^1 [\pi(p) - (1-p)\alpha(p)\gamma(p)] dF(p) \\ &\text{s.t. IC and IR.} \end{aligned} \tag{P5}$$

Note now, as opposed to Scenario 1, that the buyer's payoff from a MBG depends on his type p, which opens the possibility that MBGs can be used to increase seller revenue.

Proposition 4 *There is a solution to problem P5 which offers full money back guarantees and leaves the buyer with zero payoff regardless of his type.*

Proof. Consider the mechanism $(\alpha^*(p), \gamma^*(p), \pi^*(p)) = (1, v, v)$ for all p. First note that this mechanism is incentive compatible, as it is constant in the type p. Next note that it is individually rational as the payoff to truthful reporting is $\alpha^*(p)(pv + (1-p)\gamma^*(p)) - \pi^*(p) = 0$ for all p. To show that it is optimal for the seller, take any other incentive compatible and individually rational mechanism (α, γ, π). For any type p, the payoff of the seller from $(\alpha^*, \gamma^*, \pi^*)$ is

$$\begin{aligned}\pi^*(p) - (1-p)\alpha^*(p)\gamma^*(p) &= pv \\ &\geq \alpha(p)pv \\ &\geq \pi(p) - (1-p)\alpha(p)\gamma(p)\end{aligned}$$

where the first inequality is by $\alpha(p) \in \{0, 1\}$ and the second is by the individual rationality of (α, γ, π). Note that the last expression is the payoff of the seller from the mechanism (α, γ, π) when the buyer's type is p. Hence $(\alpha^*, \gamma^*, \pi^*)$ earns a weakly larger payoff to the seller compared to (α, γ, π) at every type. This completes the proof. ■

We would like to point out that if MBGs are not admissible, then Scenarios 1 and 2 are identical. With the possibility of MBGs, however, they lead to two very distinct outcomes. Whereas in Scenario 1 MBGs do not improve the seller's profit, in Scenario 2 it becomes feasible through MBGs to extract the buyer's full surplus. In other words the optimal mechanism identified in Proposition 4 serves all buyer types and leave each buyer type zero information rent. This is noteworthy, especially because full surplus extraction occurs in a single-agent framework. In contrast, Cremer and McLean (1988) show that the seller can extract full surplus in a multiagent problem with interdependent values, using a mechanism which is not ex post individually rational. We would also like to emphasize that the optimal mechanism of Proposition 2 does not rely on any distributional assumptions regarding the buyer's type.

4 Conclusion

A seller can screen buyers who are privately informed about their valuation and fit risk through MBGs. When fit risk is common knowledge, MBGs are not useful in optimally screening buyers. This is not to say that MBGs are of no use. MBGs can still be used to signal quality or enhance competition amongst sellers, issues that we do not consider in this paper.

We would like to conclude by highlighting two possible extensions of our results.

We consider full MBGs only in this paper, an assumption that has support in actual business practice. This restriction significantly simplifies our analysis since, by Proposition 1, we can identify the class of mechanisms that offer only full MBGs fairly easily. If partial

MBGs are possible, then the class of feasible mechanisms for the seller enlarges. However if one were to come up with a tractable description of this class, one could analyze under what conditions full MBGs dominate partial MBGs, and vice versa.

A second possible extension pertains to the nature of private information. We assume that if a good does not fit then the buyer's value for it is zero, which is common knowledge. It would also be interesting to study the role of MBGs when non-fit value is also a part of the buyer's private information.

References

[1] Anderson, E. T, K. Hansen and D Simester (2009), "The option value of returns: Theory and empirical evidence," *Marketing Science*, 28 (3), 405-423

[2] Armstrong, M. (1996), "Multiproduct nonlinear pricing," *Econometrica*, 64(1), 51-75

[3] Courty, P. and H. Li (2000), "Sequential screening," *Review of Economic Studies*, 67, 697-717.

[4] Cremer and McLean (1988), "Full extraction of the surplus in Bayesian and dominant strategy auctions," *Econometrica*, 56, 1247–1257.

[5] Davis, S., Gerstner and M. Hagerty (1995), "Money-back guarantees in retailing: Matching products to consumer tastes," *Journal of Retailing*, 71(1), 7-22.

[6] Dowling, M. (1999), "Getting less in return," *Catalog Age*, 16(4) 1, 18.

[7] Guide, D., Jr., G. Souza, L. Van Wassenhove, and J. Blackburn (2006), "Time value of commercial product returns," *Management Science*, 52(8), 1200-1214.

[8] Heal, Geoffrey (1977), "Guarantees and risk-sharing," *Review of Economic Studies*, 44(3), 549-560.

[9] Heiman, A., B. McWilliams and D. Zilberman (2001), "Reducing purchasing risk with demonstrations and money-back guarantees," *Marketing Management Journal*, 11, 1-15.

[10] Mann, D. and J. Wissink, "Money-back warranties vs. replacement warranties: A simple comparison," *American Economic Review*, 80(2), 432-436.

[11] Manelli, A. and D. Vincent (2007), "Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly," *Journal of Economic Theory*, 137, 153–185.

[12] Matthews, S. and J. Moore (1987), "Monopoly provision of quality and warranties: An exploration in the theory of multidimensional screening," *Econometrica*, 55 (2), 441-467.

[13] Matthews, S. and N. Persico (2005), "Information acquisition and the excess refund puzzle," *Working paper*, University of Pennsylvania

[14] McWilliams, B. (2012), "Money-back guarantees: Helping the low quality retailer," *Management Science*, 58(8), 1521-1524.

[15] Moorthy, S., and K. Srinivasan (1995), "Signaling quality with a money-back guarantee: The role of transaction costs," *Marketing Science*, 14(4), 442-466.

[16] Myerson, R. (1981), "Optimal auction design," *Mathematics of Operations Research.*

[17] Padmanabhan, V. and I. P. L. Png (1997), "Manufacturer's returns policies and retail competition," *Marketing Science*, 16(1), 81-94

[18] Rochet, J-C. and P. Chone (1998), "Ironing, sweeping and multidimensional screening," *Econometrica*, 66(4), 783-826.

[19] Yeon-Koo Che (1996), "Customer return policies for experience goods," *The Journal of Industrial Economics*, XLIV(1), 17-24.