Negotiation with Price-dependent Probability Models ⋆

Antonio Bella, Cèsar Ferri, José Hernández-Orallo, and María José Ramírez-Quintana

Universidad Politécnica de Valencia, DSIC, Valencia, Spain

Abstract. Negotiation and agreement generally require models of the peers who are involved in the negotiation. One typical area where negotiation takes place is in selling and retailing, which is also known as Customer Relationship Management (CRM). Customers and products are usually modelled using previous retailing experiences with similar or dissimilar customers and products. Machine learning is typically used to construct these models, which can be used to design mailing campaigns, to recommend new products, to do cross-selling, etc. Many CRM problems can already be solved through rankers, recommender systems, etc., provided that there are good models of customer and product behaviours available. A related but more general problem is when models are used to negotiate with one or more features of the product (or, less frequently, the customer) such as prices, bonuses, warranties, etc. Additionally, if it is possible to make several bids until an agreement is reached, methods must be devised so that the maximum profit is obtained by the seller. In this work, we present a taxonomy of CRM problems, from which we distinguish those that have already been solved and those whose solutions are still pending. Then, we extend classical purchase probability rankings to the notion of profit probability curves (price-dependent distributions), and we propose a simple negotiation solution for these cases.

Keywords: negotiation, agreement, bargaining, CRM, ranking, probability estimation, negotiable features, machine learning.

1 Introduction

The evolution of small-sized retailing to mass Customer Relationship Management (CRM) has usually implied greater automation of the selling process, the selection of products, customer prescriptions, cross-selling, up-selling, market segmentation, etc. Techniques such as mailing campaign design, recommender systems, and others (e.g. data-mining and machine learning [5] from the area of business intelligence) are nowadays common in these applications. However, this evolution has not usually included traditional techniques in selling such as

⋆ This work has been partially supported by the EU (FEDER) and the Spanish MEC/MICINN, under grant TIN 2007-68093-C02 and the Spanish project "Agreement Technologies" (Consolider Ingenio CSD2007-00022).

bargaining and other kinds of negotiation because of the difficulty of automating them and also because more general models about customer behaviour are needed. With the diversification and specialisation of services, there is a need for new techniques that deal with selling scenarios where price (or any other negotiable feature) can be adjusted to the customer (price discrimination using custom discounts and offers) in order to achieve the goal of mass customisation.

Consider, for instance, a classical mailing campaign design or market targeting problem. Given n equal products and m customers, the subset of customers to whom a product must be offered (customer prescription) must be devised. In order to do this, machine learning is usually employed to learn probabilisitc models, which assess the probability of buying for each customer, from which a ranking of customers is performed. If the campaign has both fixed and variable costs, there are simple techniques (e.g. profit lifts) to calculate the subset of customers that maximises the expected profit [4]. Similar techniques can be used if there are many kinds of products where several rankings and constraints (stocks) must be taken into account [2].

A very different type of problem, however, is when product or customer buying expectations can be influenced by changing one or more features of the product or the customer, such as prices, bonuses, warranties, etc. For instance, the probability of a customer buying a product changes dramatically if the price is modified. And it is frequent to adjust prices (or, in a subtler way, to offer discounts) to those customers who are less eager to buy a product. Setting different prices for various customers can allow us to maximise profit.

In this scenario, we need to move from probabilistic models (where, given a product and a customer, we have a probability of buying) to profit probability curves or price-dependent distributions (where, given a product and customer, we see the evolution of the probability of buying according to a certain feature (e.g. price)). This more complex scenario has not been addressed to date.

Additionally, negotiating with price or other negotiable features normally implies that counter-offers from both buyer and seller are allowed up to a maximum number of bids or until an agreement is reached. In other words an offer can be made to a customer at a given price, and if the customer does not buy the product, a lower price (a special discount offer for that particular customer) can be offered. This is the beginning of a negotiation process between the seller and the buyer that we have already studied in [3]. In that work, we developed several methods to derive the profit probability curves, and we also introduced new negotiation strategies.

However, in general, the case when a negotiation can take place, in parallel, for several products and customers must also be solved. In this case, not only must the profit probability curves for each pair of product and customer be derived, but is also necessary to analyse how to combine these curves to develop new optimal negotiation strategies in these much more complex scenarios.

This work attempts to address this problem. We place ourselves on the side of the seller. Since the main overall objective is to sell as much as possible with the highest possible net price. Therefore, agreements are considered to be optimal when they are optimal for the seller. Nonetheless, the techniques that we discuss in this paper can also be used when both buyer and seller model each other, or when a buyer wants to negotiate with several sellers.

The paper is organised as follows. In Section 2, we devise a taxonomy of CRM prescription problems, from which we distinguish those that have been solved in the past and those which we propose a solution. In Section 3, we present a detailed description of a specific scenario with one kind of product, a negotiable price, and M customers, and explain how it is possible to solve other kinds of CRM prescription problems using the same strategy. In Section 4, we present our conclusions and future work.

2 A Taxonomy of CRM Prescription Problems

There are a great variety of different CRM prescription problems that can be defined only by taking into account the cardinality of the different kinds of products and customers $(1{-}1, N{-}1, 1{-}M, N{-}M)$ and the presence or absence of negotiation. As we have stated in the introduction, if a feature is negotiable, then we can introduce some kind of negotiation into the CRM process; however, if it is non-negotiable (fixed), then we are dealing with a traditional CRM prescription problem. Focusing on the seller, in this paper, when we refer to the price of a product we mean the net price that the seller obtains for the product, e.g., if a customer buys a product for 3 euros, the net price is 3 euros since fixed and variable costs are not included (except if mailing is involved).

In any of these situations, data-mining techniques can help the seller by modelling customer behaviour in order to make good decisions. In this context, that means obtaining as much profit as possible.

Table 1. Different CRM prescription problems that consider the number of different kinds of products to sell, whether the net price for the product is fixed or negotiable, and the number of customers.

Case	Kinds of products	Net price	Number of customers	Approach
1	1	fixed	1	Trivial
2	1	fixed	M	Customer ranking [4]
3	N	fixed	1	Product ranking [4]
4	N	fixed	M	Joint Cut-off [2]
5	1	negotiable	1	Negotiable Features [3]
6	1	negotiable	M	This work
7	N	negotiable	1	This work
8	N	negotiable	M	Future work

Table 1 shows eight different CRM prescription problems that are defined by considering the number of products and customers involved as well as the fixed or negotiable nature of the net price for each product. The last column shows several approaches, that have already been proposed in the literature for solving some of these problems. It also shows others that are proposed in this paper for solving the remainder. We discuss each of them in more detail below.

2.1 CRM Prescription Problems without Negotiation

Case 1 in Table 1 (one kind of product, fixed net price, and one customer) is trivial. In this scenario, the seller offers the product to the customer at a fixed price and the customer may or not buy the product. The seller cannot do anything more because s/he does not have more products to sell. S/he cannot negotiate the price of the product with the customer, and s/he does not have any more customers for the product.

Case 2 in Table 1 (one kind of product, fixed net price, and M customers) is the typical case of a mailing campaign design. The objective is to obtain a customer ranking to determine the set of customers to whom the mailing campaign should be directed in order to obtain the maximum profit. Data-mining can help in this situation by learning a probabilistic estimation model from previous customer data that includes information about similar products that have been sold to them. This model will obtain the buying probability for each customer, so by putting them in order of decreasing buying probability, the most desirable customers will be at the top of the ranking. Using a simple formula for marketing costs, we can establish a threshold/cut-off in this ranking. The customers above the threshold will be offered the product. This is usually plotted using the so-called lift charts.

Case 3 in Table 1 (N kind of products, fixed net price, and one customer) is symmetric to case 2. Instead of N customers and one product, in this case, there are N different products and only one customer. The objective is to obtain a product ranking for the customer. Similarly, data-mining can help to learn a probabilistic estimation model from previous product data that have been sold to similar customers. This model will predict the buying probability for each product, so by putting them in order of decreasing buying probability, the most desirable products for the customer will be at the top of the ranking. This case overlaps to a great extent with recommender systems.

Case 4 in Table 1 (N kinds of products, fixed net price, and M customers) is studied in [2]. This situation is more complex than the cases 2 and 3, since there is a data-mining model for each product. In other words, there are N rankings of customers (one for each product) and the objective is to obtain the set of customers that gives the maximum overall profit. Note that, normally, the best local cut-off of each model (the set of customers that gives the maximum profit for one product) does not give the best global result. Moreover, several constraints would have to be fulfilled (limited stock of products, the customers can only buy one product), which usually happens in real situations. Two different methods are proposed in [2] to obtain the global cut-off: one is based on merging the prospective customer lists and using the local cut-offs, and the other is based on simulation. The study in [2] shows that use simulation to adjust model cut-off obtain better results than more classical analytical methods.

2.2 CRM Prescription Problems with Negotiation

In the above four cases, we have assumed that the net price of the product is fixed. When this is not the case, the objective of the sellers changes. They do not have to sell the product, but they have to sell it at the maximum price, which the seller and the buyer can negotiate. Figure 1 (Left) shows how the behaviour of a customer can be approximated. The customer buys the product if the price is less than or equal to the maximum price that s/he could pay for it. Figure 1 (Right) shows the expected profit for each price (the expected profit is equal to the buying probability multiplied by the price). In this case, it is clear that the maximum expected profit is obtained when the seller sells the product at the maximum price that the customer can pay.

We will show later, that the seller does not know this real model and tries to learn the most accurate model from previous data by using data-mining techniques.

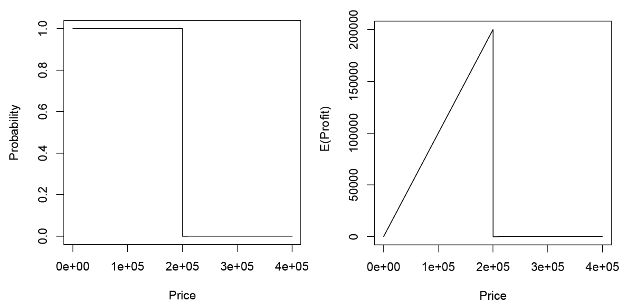

Fig. 1. Real behaviour of the buying model of a customer. **Left**: Probability distribution function. **Right**: Associated expected profit.

In [3], a formal definition of a negotiable feature is given and a case with one kind of product, a negotiable price, and one customer is studied (Table 1, case 5). In this case, there is one probabilistic data-mining model for the customer, which was learned from previous product data. Several data-mining algorithms and several negotiation strategies were studied experimentally. As can be observed in [3] the best results were obtained by approaching the buying model curve by the cumulative distribution function of a normal distribution. The mean μ is equal to the value obtained by a regression model and the standard deviation σ is equal to the standard error of the model (mean absolute error, mae). Figure 2 shows an example of a probabilistic buying model of a customer who is interested in buying a flat that is approximated by a normal distribution with $\mu = 200,000$ and $\sigma = 20,000$. As can be observed in this figure, the mean price has a probability of 0.5, so the maximum expected profit is located just to the left of that point (the seller will decrease the price to increase the probability of buying the product if s/he can only make one offer to the customer). If the seller can make several offers, then the offers can be distributed along the curve to maximise the profit of the product, according to several negotiation strategies.

There are some details that should be taken into account from our work in [3]. Note that the normal distribution is limited on the left, when working with the expected profit curve (i.e., the probability density function) because the expected profit is zero when the price of the product is zero. Also note that, in [3], only symmetric normal distributions were considered, but it would also be interesting

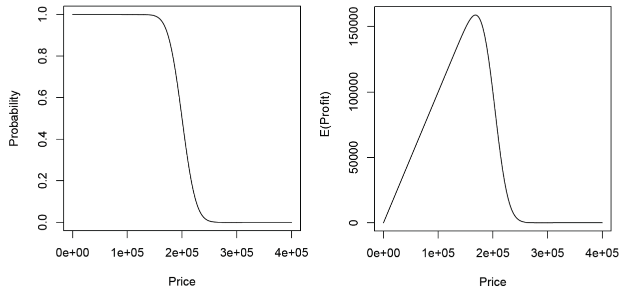

Fig. 2. Probabilistic buying model of a customer approximated by a normal distribution with $\mu = 200,000$ and $\sigma = 20,000$. **Left**: Probability distribution function. **Right**: Associated expected profit.

to study non-symmetric normal distributions (i.e., to consider different standard deviations on the left and on the right of the mean, as is common in real life. For example, it is easy to buy a product simply because it is cheap (although not essential), therefore, in this situation the standard deviation on the left will be greater than the one on the right. An example of a skew-normal distribution [1] can be seen in Figure 3.

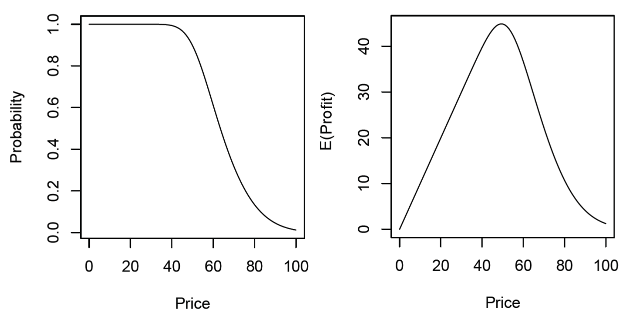

Fig. 3. Probabilistic buying model of a customer approximated by a skew-normal distribution with location $\xi = 50$, scale $\omega = 2$, and shape $\alpha = 3$. **Left**: Probability distribution function. **Right**: Associated expected profit.

The cases 6, 7 and 8 in Table 1 have not yet been studied. However, they can be understood as an extension of case 5 combined with the rankings of customers and products that are used in the approaches proposed in Table 1 for the cases 2 and 3.

In case 6 in Table 1 (one kind of product, a negotiable price, and M customers), there is a curve for each customer, that is similar to the curve in case 5 (Figure 4, Left). If the seller can only make one offer to the customers, the seller will offer the product at the price that gives the maximum expected profit (in relation to all the expected profit curves) to the customer whose curve achieves the maximum. However, if the seller can make several offers, the seller will distribute the offers along the curves following a negotiation strategy. In this case, the seller not only changes the price of the product, but the seller can also change the customer that s/he is negotiating with, depending on the price of the product (that is, by selecting the customer in each bid who gives the greatest expected profit at this price). Therefore, these curves can be seen as an evolution of the customer ranking for each price.

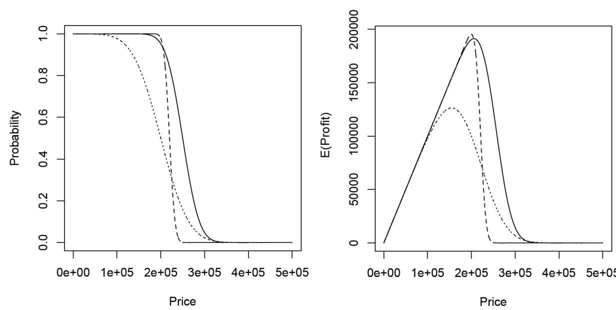

Fig. 4. Probabilistic buying models of 3 different customers approximated by 3 normal distributions with $\mu_1 = 250,000$ and $\sigma_1 = 30,000$, $\mu_2 = 220,000$ and $\sigma_2 = 10,000$, and $\mu_3 = 200,000$ and $\sigma_3 = 50,000$. **Left**: Probability distribution function. **Right**: Associated expected profit.

Case 7 in Table 1 (N kind of products, a negotiable price, and one customer) is symmetric to case 6. Instead of one curve for each customer, there is one curve for each product that was learned from the previous customer data. In this case, the curves represent a ranking of products for that customer. The learned data-mining models will help the seller to make the best decision about which product the seller offers to the customer and at what price. Figure 4 is an example of this case since the curves would represent three different products to be offered to one customer.

Case 8 in Table 1 (N kind of products, a negotiable price, and M customers) is the most complex of all. There is one data-mining model for each product and customer (i.e., $N \times M$ curves). The objective is to offer the products to the customer at the best price in order to obtain the maximum profit. Multiple scenarios can be proposed for this situation: each customer can buy only one product; each customer can buy several products; if the customer buys something, it will be more difficult to buy another product; there is limited stock; etc.

To solve cases 6, 7 and 8, we propose extending the classical concept of ranking customers or products to profit probability curves in order to obtain a ranking of customers or products for each price (similar to cases 2 and 3). For example, Figure 4 shows that, for a price of 300,000 euros the most desirable customer is the one represented by the solid line, the second most desirable one is the customer represented by the dotted line, and the least desirable customer is the one represented by the dashed line. The situation changes for a price of 200,000 euros; at that point the most desirable customer is the one represented by the dashed line, the second most desirable one is the customer represented by the solid line, and the least desirable customer is the one represented by the dotted line. Therefore, an important property of these probabilistic buying models is that there is a change in the ranking at the point where two curves cross.

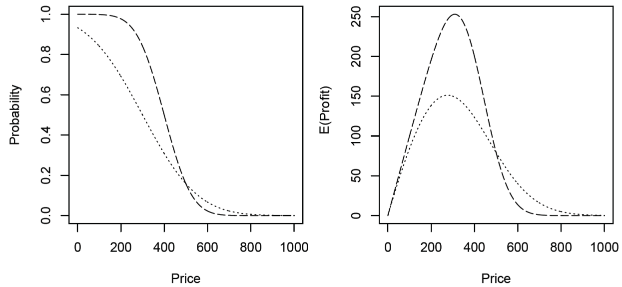

Fig. 5. Probabilistic buying models of 2 different customers approximated by 2 normal distributions with $\mu_1 = 400$ and $\sigma_1 = 100$, and $\mu_2 = 300$ and $\sigma_2 = 200$. **Left**: Probability distribution function. **Right**: Associated expected profit.

3 Scenario with Negotiable Price and several Customers

To study case 6 in more depth, we start with the simplest situation with two customers, and we explain the negotiation strategy that the seller will follow by means of an example.

In Figure 5, there are two curves representing the buying model of two different customers. The buying model of the first customer follows a normal distribution with $\mu_1 = 400$ and $\sigma_1 = 100$, and it is represented by a dashed line. The buying model of the second customer follows a normal distribution with $\mu_2 = 300$ and $\sigma_2 = 200$, and it is represented by a dotted line. These are the models; however, the real situation is that the maximum buying price for customer 1 is 100 euros and 150 euros for customer 2.

We assume a simple negotiation process for this example. The negotiation strategy that we describe is similar to the Best Local Expected Profit (BLEP) strategy explained in [3], but in that case the number of offers was limited to n.

Table 2. Left: Trace of the negotiation process. **Right:**Trace of the negotiation process with the ordering pre-process.

Offer	Price	Customer	Accepted
1	309	1	No
2	214	1	No
3	276	2	No
4	149	1	No
5	101	1	No
6	150	2	Yes

Offer	Price	Customer	Accepted
1	309	1	No
2	276	2	No
3	214	1	No
4	150	2	Yes

The strategy consists of offering the product at the price that obtains the maximum expected profit for the customer whose curve reaches the maximum. If the customer accepts the offer, the process is finished. If the customer does not accept the offer, his/her curve is normalised taking into account the following: the probabilities of buying that are less than or equal to the probability of buying at this price will be set to 0; and the probabilities greater than the probability of buying at this price will be normalised between 0 and 1. This process is

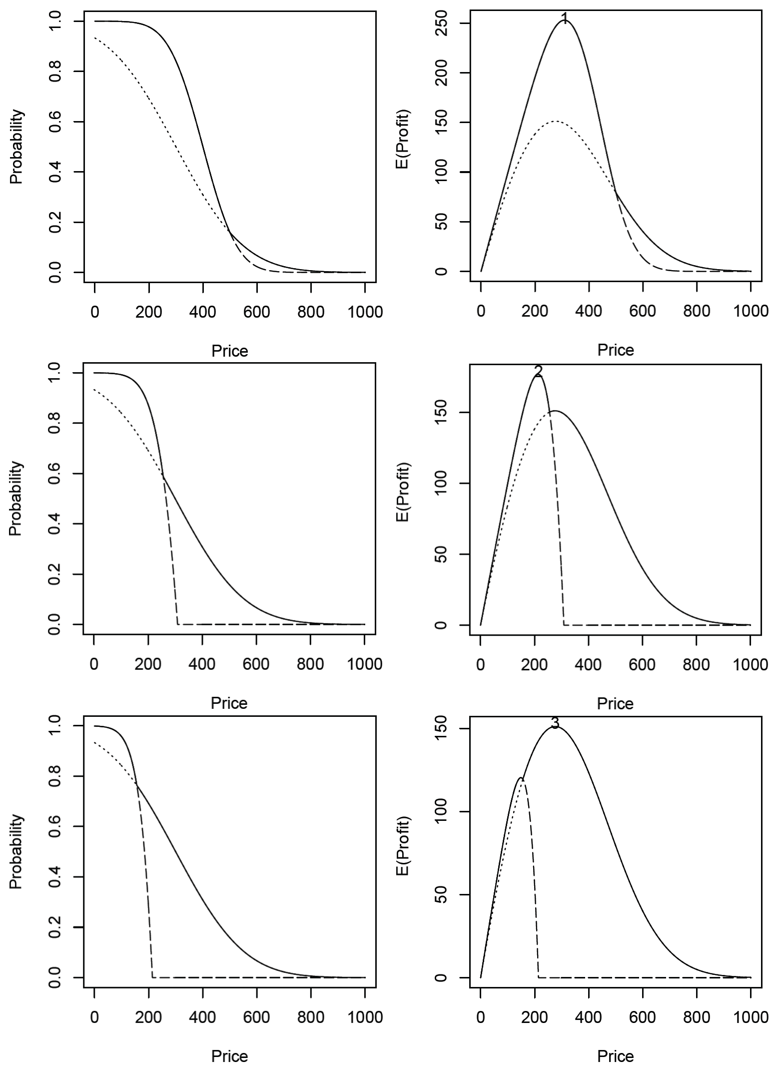

Fig. 6. Points 1, 2 and 3 in the negotiation process. **Left**: Probability distribution function. **Right**: Associated expected profit.

repeated with each customer until one of them accepts an offer[1]. The trace of the negotiation process is described in Table 2 (Left) and shown graphically in Figures 6 and 7. In each iteration, the maximum of the functions is calculated (the envelope curve). The envelope curve is represented by a solid line in Figures 6 and 7.

Note that as Table 2 (Left) shows, the third offer is greater than the second one. This is because there is more that one customer in the negotiation process and the offer is made at the price that maximises the expected profit at each iteration. Therefore, it is easy to propose an improvement for this negotiation strategy with a limited number of offers, which is similar to BLEP with n bids. This improvement is a pre-process that consists of calculating the n points and ordering them by the price before starting the negotiation. Following the example shown in Table 2 (Left), if there are only 4 bids no one will buy the product. However, with our improvement (the pre-process) customer 2 will buy the product at a price of 150 euros as shown in Table 2 (Right).

This negotiation scenario suggests that other negotiation strategies can be proposed for application to problems of this type in order to obtain the maximum profit. One problem with the BLEP strategy is that it is very conservative. It might be interesting to implement more aggressive strategies that make offers at higher prices (graphically, more to the right). A negotiation strategy that attempts to do this is one of the strategies proposed in [3], the Maximum Global Optimisation (MGO) strategy (with n bids). The objective of this strategy is to obtain the n offers that maximise the expected profit by generalising an optimisation formula that was presented in [3].

In case 6, we have presented an example with two customers and one product, but it would be the same for more than two customers. In the end, there would be one curve for each customer, and the same negotiation strategies could be applied.

Case 7 (N kind of products, a negotiable price, and one customer) is the same as case 6, but the curves represent the buying model of each product for each customer, and a ranking of products will be obtained for each price.

Case 8 (N kind of products, a negotiable price, and M customers) can be studied using the same concept of expected profit curves, but there will be $N \times M$ curves. The use of simulation or some kind of evolutionary computation will be necessary to obtain a good solution because case 8 is similar to case 4 where the best point in each curve does not give the best global solution. For each of the N kind of products, there will be M curves that belong to the buying model of each customer. Figure 8 presents a simple example with two products and two customers, where a customer can only buy a maximum of one product. In this example, customer 1 (dashed line) has the maximum expected profit for both products, corresponding to 80 euros for product 1 and 112 euros for product 2. The best local decision would be to offer product 2 to customer 1; however, customer 2 (dotted line) has the maximum expected profit of 27 euros

[1] More stop conditions are possible (e.g. limited number of offers (BLEP with n bids), minimum selling price, etc.)

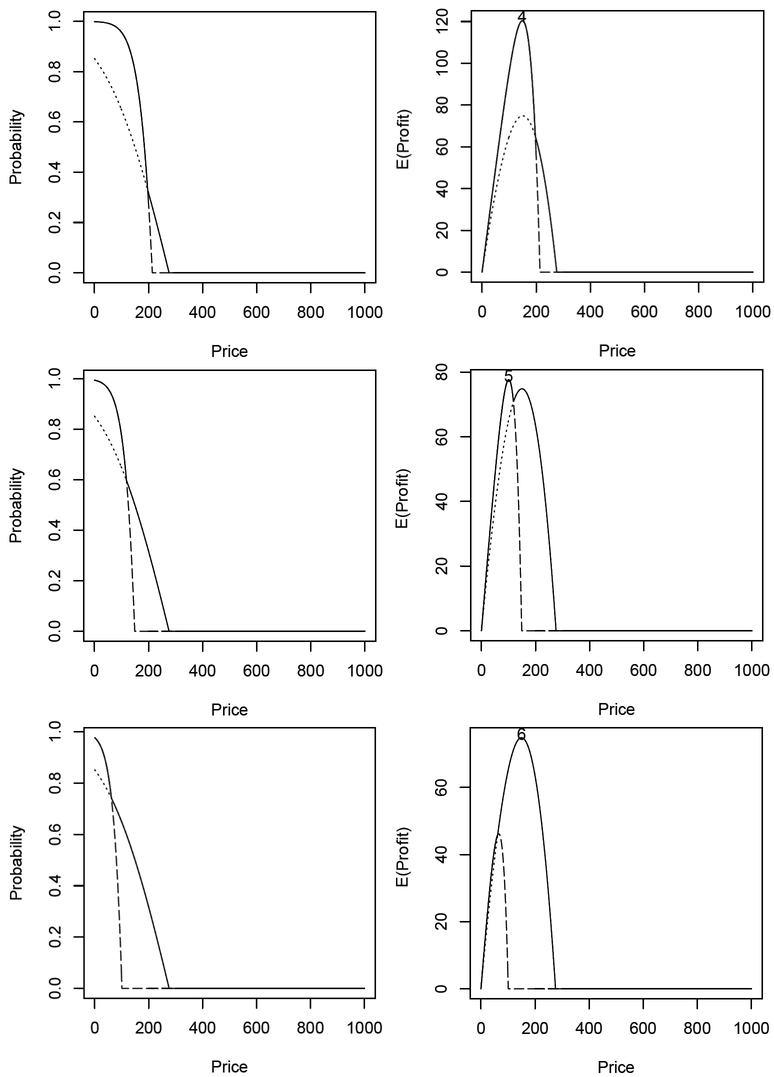

Fig. 7. Points 4, 5 and 6 in the negotiation process. **Left**: Probability distribution function. **Right**: Associated expected profit.

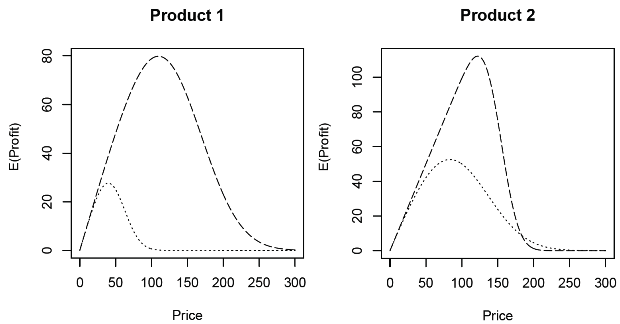

Fig. 8. Example of probabilistic models of two different customers. **Left**: Product 1. **Right**: Product 2.

for product 1 and 52 euros for product 2. Thus, for a better global solution, it would be better to offer product 2 to customer 2 and product 1 to the customer 1 since the overall profit would be greater.

4 Conclusions

In this paper, we have devised a taxonomy of CRM prescription problems, where automated learning can help a seller to make a good decision about which product should be offered to which customer and at what price in order to obtain as much overall profit as possible.

Some of these problems have already been studied, and we have explained the approaches proposed to solve them. In the cases that have not yet been studied (negotiable price, several products and/or several customers), we have proposed a solution based on the extension of rankings to expected profit curves, in which there is a ranking of customers and/or products for each price of the product.

As future work, we plan to study the performance of the proposed methods with experiments applying the negotiation strategies described in Section 3 and other suitable negotiation strategies to cases 6, 7 and 8 shown in Table 1.

References

1. A. Azzalini. A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*, 12:171–178, 1985.
2. A. Bella, C. Ferri, J. Hernández-Orallo, and M.J. Ramírez-Quintana. Joint cut-off probabilistic estimation using simulation: A mailing campaign application. In *IDEAL*, volume 4881 of *LNCS*, pages 609–619. Springer, 2007.
3. A. Bella, C. Ferri, J. Hernández-Orallo, and M.J. Ramírez-Quintana. Feature Dependent Models. Technical Report http://users.dsic.upv.es/∼abella/papers/FDM.pdf, Universidad Politécnica de Valencia, 2009.
4. M.J.A. Berry and G.S. Linoff. *Mastering Data Mining: The Art and Science of Customer Relationship Management*. Wiley, 1999.
5. T.M. Mitchell. *Machine Learning*. McGraw-Hill, 1997.