Optimizing Merchant Revenue with Rebates

Rakesh Agrawal
Search Labs
Microsoft Research
rakesha@microsoft.com

Samuel Ieong
Search Labs
Microsoft Research
saieong@microsoft.com

Raja Velu*
School of Management
Syracuse University
rpvelu@syr.edu

ABSTRACT

We study an online advertising model in which the merchant reimburses a portion of the transacted amount to the customer in a form of rebate. The customer referral and the rebate transfer might be mediated by a search engine. We investigate how the merchants can set rebate rates across different products to maximize their revenue. We consider two widely used demand models in economics—linear and log-linear—and explain how the effects of rebates can be incorporated in these models. Treating the parameters estimated as inputs to a revenue maximization problem, we develop convex optimization formulations of the problem and combinatorial algorithms for solving them. We validate our modeling assumptions using real transaction data. We conduct an extensive simulation study to evaluate the performance of our approach on maximizing revenue, and found that it generates significantly higher revenues for merchants compared to other rebate strategies. The rebate rates selected are extremely close to the optimal rates selected in hindsight.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics; F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General

General Terms

Algorithms, Economics

Keywords

Internet advertising, Rebates

1. INTRODUCTION

Sponsored search, where merchants pay search engines for displaying their advertisements or for redirecting traffic to their websites, has played a key role in making internet search ubiquitous on the web. Under the *pay per click* model of sponsored search, merchants pay for clicks that bring users to their website. An increasingly popular variant is the *pay per action* model, where merchants pay the search engine only when users take some action such as placing an order at the website. See [10, 13, 16, 20] for concept, history, and analysis of sponsored search.

Cashback, an experimental program introduced by Microsoft in May 2008, added a novel twist to the pay per action model. Merchants first select products that they would like to advertise. They also select the referral commission they are willing to pay per transaction as a fraction of the sales price. As in pay per action, merchants only pay the search engine when a transaction occurs. But the search engine then returns the commission to the consumer in the form of a cash rebate. This program can thus be viewed as one where merchants offer rebates on their products, facilitated by the search engine. In effect, when a user searches for a product and buys it through a participating merchant, the user receives a percentage of the amount she pays as rebate.[1]

There has been little work on analyzing Cashback as an advertising program. An exception is [6], where Chen *et al.* proposed several different revenue sharing mechanisms that are reminiscent of Cashback, and compare their revenue properties from the search engine's perspective.

In this paper, we adopt the *merchant's* perspective, and examine the question of how merchants can best use rebates to maximize their revenue. Assuming that the search engine passes on the whole of commission received from a merchant for a transaction to the customer as rebate, the problem of determining how much commission should a merchant pay for a transaction becomes the same as computing optimum rebate.[2] We thus address the problem of selecting rebate rates to maximize revenue subject to a budget constraint, taking into account that a merchant may carry more than one product line, and that the effects of rebates on different product lines may differ.

Our approach is to tackle this problem in two steps. First,

*Work done while author was visiting Search Labs.

[1]The program has been discontinued as of July 2010. The study of the question of how to select rebate rates remains interesting nonetheless from a research viewpoint. There are also other websites where merchants can offer rebates to online shoppers, such as `ebates.com` and `fatwallet.com`.

[2]It is straightforward to extend our work to the case where the search engine charges a fixed amount of administrative fee for the rebate processing service or charges a fee which is some fraction of the amount of rebate handled.

WSDM'11, February 9–12, 2011, Hong Kong, China.
Copyright 2011 ACM 978-1-4503-0493-1/11/02 ...$10.00.

for different products, we estimate the sensitivity of demand to changes in prices and rebates using historical data. Second, treating the values estimated in the previous step as parameters, we formulate an optimization program to find the best rebate rate for every product. We consider two important economic demand models [18]—linear demand and log-linear demand—and show how rebates can be incorporated into these models. We also address the algorithmic question of revenue maximization under both demand models. While the two demand models give rise to different optimization problems, we provide efficient algorithms to solve both problems.

We evaluate our approach using both real and synthetic transaction data. We first validate our modeling assumptions with transaction data obtained from Microsoft Cashback operations. For evaluation, as we do not have direct access to merchants, we conduct an extensive simulation study using synthetic transaction data, and compare our proposed solution to a number of rebate selection heuristics on the amount of revenue each method generates. Our approach outperforms all tested heuristics in almost all of the cases.

The rest of the paper is organized as follows. We discuss related work from the economic and marketing literature in Section 2. We state our problem formulation in Section 3. In Section 4, we explain how the linear and the log-linear demand model can be extended to estimate demand sensitivity to rebate rates. In Section 5, we show how to solve the rebate optimization problem efficiently under both demand models. In Section 6, we validate our modeling assumptions using real transaction data collected over a year. In Section 7, we conduct an extensive simulation study using a synthetic data generation process designed to mimic real data, and evaluate the end-to-end performance of our approach to maximizing revenue. We conclude with the key findings and directions for future work in Section 8.

2. RELATED WORK

Merchants often choose between price cuts and rebates for stimulating sales [9, 17, 19]. The key difference between these two tactics is that price cuts offer discounts up front whereas rebates offer discounts after the product is purchased at the regular price. This difference leads to a phenomena known as “slippage”, where some rebates are not redeemed [11]. In the context of Cashback, the rebate credit to the customer collection process is automatic, hence slippage in its traditional sense is a non-issue. Nonetheless, users wait to collect the rebates creating a float for the search engine.

There has been work on finding rebate policies that maximize profits by Arcelus *et al.* [1, 2, 3] and Khouja *et al.* [14, 15]. Their work focuses on settings where merchants can choose to set the sales price, the rebate value, and the order size (for managing inventory). Our work is different in several ways. First, in our setting, the merchant carries multiple products, and the optimal rebates may differ for each product. Second, the size of the rebate program is governed by a budget that plays a central role in determining the size of the advertising program. On the other hand, we do not model inventory holding costs and treat prices as given, and focus on the selection of rebate rates.

Past work on modeling the relationship of demand and rebates is discussed further in Section 4.

3. PROBLEM SETUP

Let the products a merchant sells be $P = \{1, 2, \ldots, n\}$, and the budget for rebates be b. The relationship among demand, price, and rebates is product-dependent. For product i, let its price be p_i, its rebate rate be r_i, and its demand, which depends on p_i and r_i, be $q_i(p_i, r_i)$.

We define revenue as the net proceeds the merchant receives, i.e., the gross revenue from the products sold less the rebates paid. Our objective is to find rebate rates, one for each product, that maximizes revenue without exceeding the budget. Prices are treated as input parameter to the problem. The rebate optimization problem can be stated as follows.[3]

$$\begin{aligned} \max_{\boldsymbol{r}} \quad & \sum_{i \in P} q_i(p_i, r_i) p_i (1 - r_i) \\ \text{subject to} \quad & \sum_{i \in P} q_i(p_i, r_i) p_i r_i \le b \\ & 0 \le r_i \le 1 \qquad\qquad \forall i \in P \end{aligned} \tag{1}$$

Budgets are central to the problem formulation, as they are instrumental in controlling exposure risk in sponsored search [16]. Even though in the long run, a budget may be effectively unlimited as long as a merchant is making a profit, in the short run, it is needed to balance the allocation of capital across different operations. Hence, we treat the budget as given to the optimization. The presence of a budget presents a trade-off between offering more rebates on one product versus another. It also requires careful planning of rebates so as not to run out of budget before the end of an advertising campaign. In the (unlikely) scenario where a merchant is not constrained by a budget, selecting rebate rates can be simplified to selecting the optimal rate for each product individually.

The above formulation treats prices as input parameters, and tacitly assumes that prices stay unchanged during the rebate program. Given that rebate programs often run for short duration, this assumption is realistic. In situations where prices may change, one can rerun the optimization and adjust the rebate rates accordingly.

The formulation also assumes that there is no constraint on supply for each product. This assumption holds for digital goods [12], or when supply is large compared to demand. Limits on supply can be modeled by adding constraints of the form $q_i(p_i, r_i) \le S_i$ to the optimization problem (1). Our results generalize to this setting.

This work assumes that the demand of a product can be estimated independently of prices and rebates offered by other products, and leaves the general case where prices and rebates of different products may interact to future work.

To instantiate the optimization problem, the demand function $q_i(p_i, r_i)$ is estimated using historical data. This introduces uncertainty in the underlying problem. Our results directly apply if one is interested in maximizing expected revenue while satisfying the constraint in expectation. When the constraint has to be satisfied with high probability, one will need to extend proposed techniques employing ideas from *stochastic programming* [22].

[3] Our formulation and optimization technique allows one to specify limits on the minimum and maximum rebate rate for each product; this is useful when a merchant is running a targeted campaign. We keep the limits to be 0 and 1 for ease of exposition.

4. ESTIMATING DEMAND SENSITIVITY

To determine the rebates to offer, we start with estimating a relationship among prices, rebates, and quantity of goods sold (demand). We have chosen to treat prices (p) and rebates (r) as two separate variables rather than treating them as a net price variable ($p(1-r)$) for two reasons. First, consumers often consider the value of a dollar rebate to be different from a dollar discount. Second, in our experiments in Section 6, models that treat the two variables separately fit the data much better.

A simple yet widely-used demand model is the *linear demand model* [18]. It has been extended to treat price and rebate separately in [3, 14, 15].

$$q = \beta_0 + \beta_1 p + \beta_2 pr + \epsilon \tag{2}$$

where β_0, β_1, β_2 are parameters of the model to be estimated and ϵ is random noise. Given historical sales data of the form (q, p, r), we can estimate the parameters of the model using linear regression. We expect the coefficient β_1 to be negative and β_2 to be positive, consistent with the expected properties of a demand curve.

Another model we consider is the *log-linear demand model* [18]. It is an important and well-studied model in economics due to its interpretability, though it has been examined less in the context of rebates, the exception being [3]. To treat rebates separate from prices, we extend the model as follows.

$$\log q = \gamma_0 + \gamma_1 \log p + \gamma_2 \log(1-r) + \epsilon \tag{3}$$

where γ_0, γ_1, γ_2 are parameters of the model to be estimated and ϵ is random noise. As in the case of the linear demand model, we can learn the parameters of the model from historical data using linear regression.

In the economic literature, the coefficient γ_1 has been interpreted as the *price elasticity* of demand, and the coefficient γ_2 as the *rebate elasticity* of demand. For our model, both of these coefficients are expected to be negative, corresponding to the expectation that demand increases when price decreases or rebate increases.

The description above makes the simplifying assumption that demand relationship remains unchanged for the duration of the analysis. As rebate programs are often short, this assumption is reasonable. If demand relationship may shift over time, one may use ideas from time series analysis to model the time-dependent effects [4]. The central idea is to augment the basic model with time, and express demand as a function of price, rebate rates, and past demand. Our techniques also apply to such time-variant models; the limiting factor is whether one has sufficient data.

5. OPTIMIZING REBATE RATES

Treating the coefficients of demand sensitivity as parameters, we now discuss the question of how to select the optimal rebate rates. Depending on the demand model used, the solution to the optimization problem requires different techniques. In the following, we are going to consider the linear demand model and the log-linear demand model separately.

5.1 Linear Demand Model

Under the linear demand model (Equation (2)), the optimization program (1) becomes:

$$\begin{aligned}
\max_{\boldsymbol{r}} \quad & \sum_{i \in P} \Big(\beta_{0,i} + \beta_{1,i} p_i + \beta_{2,i}(p_i r_i)\Big) p_i (1 - r_i) \\
\text{subject to} \quad & \sum_{i \in P} \Big(\beta_{0,i} + \beta_{1,i} p_i + \beta_{2,i}(p_i r_i)\Big) p_i r_i \le b \\
& 0 \le r_i \le 1 \qquad \forall i \in P
\end{aligned}$$

The decision variables in this optimization are the rebate rates r_i. The prices p_i, the demand sensitivities ($\beta_{0,i}$, $\beta_{1,i}$, $\beta_{2,i}$), and the budget b, are all inputs to the problem.

To solve this optimization problem, we start with some preprocessing. Note that for any product i for which $\beta_{2,i} \le 0$, one should not offer any rebates. We thus set $r_i = 0$ for these products and remove them from further consideration. After this step, remaining products have $\beta_{2,i} > 0$.

Our problem has both a quadratic objective and a quadratic constraint. We can formulate the problem as a (convex) *quadratically constrained quadratic program* (QCQP), which can be solved in polynomial time using interior point methods [5]. There are also efficient off-the-shelf solvers for the problem [7].

Theorem 1. *Selecting rebate rates that maximize revenue under the linear demand model can be solved in polynomial time via a convex QCQP.*

Proof. The objective is quadratic in $\boldsymbol{r}$ and the budget constraint is quadratic in $\boldsymbol{r}$, hence we need to verify that the optimization is a convex one, i.e., the objective is concave (since it is maximized) and the constraints are convex [5].

Recall all products with negative $\beta_{2,i}$ have been removed during preprocessing. The objective can be expressed as a sum of quadratic and linear terms in r_i's. The coefficient to the quadratic terms are $-\beta_{2,i} p_i^2$, hence the individual quadratic terms are concave, and the overall function is concave. Similarly, the budget constraint can be expressed as a sum of quadratic and linear terms in r_i's. The coefficients to the quadratic terms are $\beta_{2,i} p_i^2$, hence they are convex and the constraint is convex. The other constraints are linear and hence also convex. $\square$

Next, we will develop a combinatorial algorithm for solving the optimization in the case of log-linear demand model. The technique can be adapted to linear demand model as well. However, given the speed of existing solvers for convex QCQP, a combinatorial algorithm for the linear demand case might not be needed.

5.2 Log-linear Demand Model

We start by re-writing the log-linear demand model (Equation (3)) as follows.

$$\begin{aligned}
\log q &= \gamma_0 + \gamma_1 \log p + \gamma_2 \log(1-r) \\
q &= \exp(\gamma_0 + \gamma_1 \log p) \exp(\log((1-r)^{\gamma_2})) = c(1-r)^{\gamma_2}
\end{aligned}$$

where c is a positive number independent of rebate rate r.

As the coefficients of the demand function (γ_0, γ_1, γ_2) and prices p are input to our optimization, we can compute c and treat it as part of the input as well. Note that each product will have its own constant c. Rewriting our optimization program (1), we have:

$$\begin{aligned} \max_{\boldsymbol{r}} \quad & \sum_{i \in P} c_i (1-r_i)^{\gamma_{2,i}} p_i (1-r_i) \\ \text{subject to} \quad & \sum_{i \in P} c_i (1-r_i)^{\gamma_{2,i}} p_i r_i \le b \\ & 0 \le r_i \le 1 \qquad \forall i \in P \end{aligned} \tag{4}$$

To solve this optimization problem, similar to the case for linear demand model, we start by removing products for which offering rebates do not improve revenue. In this case, for any product i, if $\gamma_{2,i} \ge -1$, we should set $r_i = 0$. We can verify this as follows. Denote the objective function of Eq. (4) by $f(\boldsymbol{r})$. Taking the partial derivative of f with respect to r_i,

$$\frac{\partial f}{\partial r_i} = -c_i p_i (\gamma_{2,i} + 1)(1-r_i)^{\gamma_{2,i}}, \tag{5}$$

which is non-positive when $\gamma_{2,i} \ge -1$. One arrives at the same conclusion by interpreting $\gamma_{2,i}$ as the rebate elasticity: when it is at least -1, a unit increase in rebates will generate at most a unit return in gross revenue, hence offering rebates do not improve net revenue. After preprocessing, all remaining products have $\gamma_{2,i} < -1$.

Unlike the case for linear demand model, however, after preprocessing, the resulting problem is a *non-convex* optimization problem. This is because we are maximizing an objective that is not necessarily concave. Indeed, evaluating the partial derivative in Eq. (5), when $\gamma_{2,i} < -1$, $\frac{\partial f}{\partial r_i} > 0$; the objective is actually convex rather than concave!

However, by a careful change of variables, one can find an equivalent optimization problem that is convex.

Theorem 2. *Selecting rebate rates that maximize revenue under the log-linear demand model can be (re)formulated as a convex optimization problem.*

Proof. For all $i \in P$, let $x_i = (1-r_i)^{\gamma_{2,i}+1}$. Rewriting rebate rate r_i in terms of x_i,

$$r_i = 1 - x_i^{1/(\gamma_{2,i}+1)}.$$

Changing the optimization variables in the problem from $\boldsymbol{r}$ to $\boldsymbol{x}$, and making the substitution to Eq. (4), we have

$$\begin{aligned} \max_{\boldsymbol{x}} \quad & \sum_{i \in P} p_i c_i x_i \\ \text{subject to} \quad & \sum_{i \in P} p_i c_i \left(x_i^{\gamma_{2,i}/(\gamma_{2,i}+1)} - x_i \right) \le b \\ & x_i \ge 1 \qquad \forall i \in P \end{aligned} \tag{6}$$

To verify that this problem is convex, first, note that the objective function is linear in $\boldsymbol{x}$, hence it is concave. Denote the budget constraint by $g(\boldsymbol{x})$. Taking the partial derivative of g with respect to x_i,

$$\frac{\partial g}{\partial x_i} = p_i c_i \left(\frac{\gamma_{2,i}}{\gamma_{2,i}+1} x_i^{\frac{-1}{\gamma_{2,i}+1}} - 1 \right)$$

When $\gamma_{2,i} < 1$, $\frac{\partial g}{\partial x_i} > 0$, hence g is a convex function, and the problem is a convex optimization problem. $\square$

Since the optimization problem can be formulated as a convex program, in theory it can be solved in polynomial time using interior point method. However, unlike convex QCQP, off-the-shelf solvers are slow for general convex programs. Hence, we develop a primal-dual-based combinatorial algorithm for it. The details are described in Algorithm 1.[4]

Algorithm 1: Primal-dual algorithm for (6)

```
forall the i ∈ P do x_i ← 1
S ← {}
Sort(h_i(x_i))    // Ensures h_1(x_1) ≥ ··· ≥ h_|P|(x_|P|)
for i ← 1 to |P| do
    Add i to S
    forall the j ∈ S do
        Increase all x_j while maintaining the invariant
        h_1(x_1) = ··· = h_i(x_i) = λ until
        (C1) λ = h_{i+1}(x_{i+1}), or (C2) ∑_{i∈S} g_i(x_i) = b.
        if (C1) then Break
        if (C2) then Terminate
    end
end
```

The correctness of the algorithm is based on two technical lemmas. We state the two lemmas here. Let

$$\begin{aligned} f_i(x_i) &= p_i c_i x_i, \\ g_i(x_i) &= p_i c_i \left(x_i^{\gamma_{2,i}/(\gamma_{2,i}+1)} - x_i \right), \\ h_i(x_i) &= \frac{\partial f_i(y)}{\partial x_i} / \frac{\partial g_i(y)}{\partial x_i} = \frac{1}{\frac{\gamma_{2,i}}{\gamma_{2,i}+1} x_i^{-1/(\gamma_{2,i}+1)} - 1} \ . \end{aligned}$$

Lemma 1. *A feasible solution $\boldsymbol{x}^*$ is optimal for (6) if and only if there exists a constant c such that*

- *For $i \in P$ where $x_i^* > 1$, $h_i(x_i^*) = c$;*
- *For $j \in P$ where $x_j^* = 1$, $h_j(x_j^*) \le c$.*

Lemma 2. *The function $h_i(x_i)$ is non-negative and monotonically decreasing in x_i.*

Theorem 3. *Algorithm 1 finds the optimal x to optimization problem (6) in $O(n^2) + T$ time, where T is the time needed to solve for the root of a univariate polynomial equation, and is bounded by the parameters of the problem.*

Proof. One can verify that $\frac{\partial g_i}{\partial x_i} > 0$. Hence, $\sum_{i \in P} g_i(x_i)$ increases as the x_i's increase. Condition (C2) will eventually be satisfied and the algorithm always terminates.

Because of Lemma 2, for each j, as x_j increase, the value $h_j(x_j)$ decreases, hence it is possible to maintain the invariant. At termination, the conditions to Lemma 1 will be satisfied, since for all $i, j \in S$, $h_i(x_i) = h_j(x_j) = \lambda$, and for $k \notin S$, $x_k = 1$, and $h_i(x_i) \ge h_k(x_k)$ due to sorting. Hence the algorithm finds the optimal solution.

In actual implementation, one does not increase x_j continuously; the description is only for intuition. Instead, one checks for conditions (C1) and (C2) discretely. Let $\lambda_i = h_i(x_i)$. In the outer loop, while the budget is not exceeded, we add i to S, and solve for $h_1(x_1) = \cdots = h_i(x_i) = \lambda_{i+1}$.

[4] An alternative approach, suggested by an anonymous reviewer, is to perform the substitution $x_i = p_i c_i r_i (1-r_i)^{\gamma_{2,i}}/b$ to Eq. (4), which gives rise to optimizing an convex objective over a simplex constraint with respect to x_i. Off-the-shelf solvers run more efficiently for this type of problems.

The moment the budget is exceeded, we know that we cannot add another product to S. Note that x_i can be expressed as a function of λ, i.e., $x_i = h_i^{-1}(\lambda)$. This inverse is well-defined since h_i is monotonically non-decreasing in x_i. Hence we solve a polynomial equation in one variable, λ, $\sum_{i \in S} g_i(h_i^{-1}(\lambda)) = b$. The LHS is monotonic in λ, hence its root can be found by binary search. The number of steps is bounded by parameters to the optimization problem. Note that only one equation needs to be solved.

The rest of the algorithm is bounded by $O(n^2)$, as there is a maximum of n loops, each of which takes $O(n)$ time. $\square$

6. MODEL VALIDATION

A central assumption in our demand model is that potential customers do not value price and rebates equally, and hence the effects of these variables on demand should be treated separately. We validate this assumption using real transaction data in this section.

6.1 Transaction Data

We obtained transaction data from Microsoft Cashback operations over a year. We grouped the transactions first by merchant, and then by product. We randomly sampled 40 thousand such groups of merchant-product pairs, constraining to pairs for which the product was sold at least 5 times, and for which the units sold were not identical for all days. We retrieved all transactions for the selected pairs.

This sampling process results in about 3 million transactions for evaluation. Each row of data describes one transaction, and includes information such as the merchant, the product, the date of sales, the price, the rebate rate, and the number of units sold. This provides the input to our demand estimation experiments discussed next.

6.2 Demand Estimation Evaluation

For each merchant-product pair, we compute for each day the average price, the average rebate, and the number of units sold for each group of transactions, and run regression to estimate the parameters to four models. The first two models are the linear and the log-linear models presented in Section 4. They model price and rebates separately. The other two models, the linear net-price (Linear-NP) and the log-linear net-price (Log-linear-NP), are the control models that only use the net prices (computed as prices times one minus rebate rates) to model demand. If our assumption is valid, we should see an increase in the explanatory powers of the first two models over the latter two.

To measure the explanatory powers of the models, we use the *coefficient of determination* (commonly known as R^2), which measures the proportion of the variability in the observations accounted for by the statistical model [21]. The value lies between 0 and 1, and a higher value suggests a better fit. Since the models have different number of independent variables, we also compute the *adjusted* R^2 of the models.[5] Adjusted R^2 takes into account the difference in the number of independent variables of the models, and a higher value suggests that the improvement of explanatory power due to the additional variable(s) cannot be explained by chance alone. The results are reported in Table 1.

[5] Adjusted $R^2 = 1 - (1 - R^2)\frac{(n-1)}{(n-m-1)}$, where n is the number of samples and m is the number of independent variables [21].

Model	R^2	Adjusted R^2
Linear	0.4379	0.2292
Log-linear	0.4420	0.2348
Linear-NP	0.2827	0.1730
Log-linear-NP	0.2898	0.1810

Table 1: R^2 and adjusted R^2 on transaction data.

Both the R^2 and the adjusted R^2 values of the models are significantly higher when prices and rebates are modeled as two separate variables (Linear and Log-linear) than when they are modeled as net prices (Linear-NP and Log-linear-NP). This indicates that potential customers under real transaction conditions do not value price and rebates equally. It validates our demand modeling assumptions, and reinforces the importance of treating rebate selection as a different problem from price selection.

7. SIMULATION STUDY

Our objective is to develop an approach that the merchants can use to maximize their revenue. We conducted two sets of experiments. In the first set of experiments, we consider the case of a single product, and evaluate if our approach finds good rebate rates that both satisfy the budget and maximize revenue, and whether it is better than the alternatives considered. We also study the robustness of our approach under various parameter settings. In the second set of experiments, we consider the case of multiple products. We investigate whether our approach can discover the product that is more sensitive to rebates, and whether it can find the right trade-off between rebates offered on one and the other.

Ideally we would like to evaluate our approach with real operations, but we do not have direct access to merchants. Historical data cannot help with evaluating the efficacy of rebate selection as such evaluation requires counterfactual changes to the rebates offered. To circumvent this difficulty, we design a synthetic data generator that aims to mimic real transactions, and evaluate our approach using simulation.

7.1 Synthetic Transaction Generator

To better understand the characteristics of real transactions, we examined the sample transactions used for model validation in Section 6. We observed that in most transactions, only a single unit of product is sold. This suggests that most potential customers face a *discrete choice*—given the price and rebate, whether to purchase the product or not. Hence, we adopt the following process for transaction generation.

For each day, the number of potential customers of a merchant for a given product, referred to as *traffic* henceforth, is drawn according to a Poisson distribution, parameterized by μ, the average traffic per day. We note that the traffic to a website has also been modeled as a Poisson distribution in [23]. Each potential customer faces a binary choice of whether to purchase the product. Following the discrete choice literature [24], the decision is modeled using a binary logit function; for price p and rebate rate r, the probability t that the potential customer is going to purchase the product is given by

$$t = \frac{1}{1 + \exp(-(\alpha_0 + \alpha_1 p + f\alpha_1 pr))}, \tag{7}$$

where α_0, α_1, and f are parameters specific to the product.

The probability t can be interpreted as the conversion rate of the merchant. The parameter α_1 captures how sensitive the potential customers are to price changes. The parameter f captures the relative value of rebate to price. The parameter α_0 can be viewed as an offset that helps determine the conversion rate; in our experiments, we vary the minimum conversion rate of the merchant and compute the corresponding value for α_0.

To determine if a potential customer makes a purchase, a number is drawn uniformly at random between 0 and 1. If its value is less than t, the product is bought. We refer this sample to be the potential customer's *deal-seeking attitude*, as a higher value means the person is seeking for a "better deal" (lower prices or higher rebates). Finally, the number of units sold per day is obtained by aggregating over all traffic.

This synthetic transaction generator is designed to mimic how transactions take place for merchants. The key assumption is based on the discrete choice process, supported in the economic literature [18, 24]. We believe it does not create a bias that favors our proposed approach; our approach uses only transaction data, and it is unaware of how traffic is generated, as well as the discrete choice process underlying the decisions of the potential customers.

7.2 Experimental Setup

We want to conduct an end-to-end evaluation of our approach, starting from demand estimation and ending with measuring the revenues generated based on rebate optimization. Therefore, in each trial, we first generate transaction data over some pre-determined price and rebate ranges, corresponding to a period during which a merchant is learning the demand relationship. The data is then fitted to the linear and the log-linear models and the parameters are estimated. Then, over a of 12-week evaluation period, given a budget parameter b, the merchant fixes the price of the product at price p and selects a rebate rate either based on our optimization routines or some other heuristics.

The merchant offers the said rebate until budget runs out, after which zero rebates are offered for the remaining period. We measure the revenue generated both over the duration of the rebate program, and over the entire 12 weeks, averaged over 500 trials for each experiment. The former "while rebate lasts" (WRL) scenario is appropriate when a merchant is required to offer rebates to participate in the program, whereas the latter "entire duration" (ED) scenario is appropriate when that is not the case. The revenue under ED is at least as high as the revenue under WRL, and strictly higher when budget is exhausted due to potential customers that purchase at zero rebates.

Revenue is sensitive to both traffic and the deal-seeking attitudes of the potential customers. To control for this variability, instead of running the simulation independently for each rebate program, we *couple* the simulations together: for each trial, we sample one set of potential customers along with their attitudes, and evaluate all rebate programs with respect to them. Any difference in revenue is therefore due only to the choice of rebates, but not due to differences in traffic or the attitudes of the potential customers.

Given a set of potential customers, we can compute the optimal rebate rate (in hindsight) that would maximize revenue for this specific instance. Of course, this revenue cannot be achieved in reality as it requires foreknowledge of the number of potential customers and their attitudes, but it can serve as an instance upper bound for each trial. We refer this upper bound as the optimal revenue.

7.2.1 Alternative Approaches to Selecting Rebates

We compare our approach to the following heuristics. The first two heuristics are selected due to their popularity in Cashback data. The last heuristic, motivated by feedback control, tries to adapt to demand patterns, and constitute a competitive baseline for comparison.

1. **Fixed rates (Fix-r)**. Fix rebate rate at r for the entire period, until budget runs out. In our experiment, we try three popular rates—5%, 10%, and 15%. Note that our approach also selects a fixed rate, although the rate is determined algorithmically to optimize revenue.
2. **Hi-Lo**. A merchant alternates between offering high rebate rates and low ones. In our experiment, the high rate is set at 15% and the low at 5%, and the merchant changes the rebate rate every week.
3. **Adaptive**. A merchant changes the rebate rates depending on the remaining budget. When the remaining budget is higher than expected, the rates are increased; if lower, they are decreased. In our experiment, we start with a rate of 5% and adjust the rates multiplicatively by a factor of 1.5 when the budget fails to track by more than 10%. These parameters were chosen after some basic tuning and appear to do well in simulation.

7.2.2 Performance Metric

Our metric for evaluating the different approaches to selecting rebates is the **% of optimal revenue** achieved. The revenues achieved under both WRL and ED scenarios are measured. Naturally, the higher the value for this metric, the better the approach.

When we evaluate our proposed approach, we will also examine the rebate rates selected according to the linear and the log-linear models for each trial, and compare them to the optimal rebate rates. This helps to measure how close we are to the optimal choice.

7.2.3 Simulation Parameters

There are altogether five parameters that we vary in our simulation study. Four of these parameters, μ, α_0, α_1, and f, govern the synthetic transaction generation process. As mentioned, we do not explicitly select α_0, but rather determine α_0 based on the minimum conversion rate, t_{min}, which we vary in our experiments. The parameter t_{min} corresponds to the expected conversion rate when a merchant selects the default price and offers zero rebate for the product. The fifth parameter is the budget b, which determines how much rebates are available during the evaluation period. These parameters, along with their default values and the ranges with which we have experimented, are summarized in Table 2.

Throughout the study, price p is set at 100. This is without loss of generality, as an increase in price can be mapped to a corresponding decrease in α_1 and in budget b. Hence, variation in price is implicitly tested when we vary the parameters α_1 and b.

Parameter	Default	Min	Max
μ (traffic)	100	10	400
α_1 (price sensitivity)	-0.08	-0.06	-0.10
f (relative value of rebate)	0.8	0.4	1.2
t_{min} (min conversion rate)	0.04	0.005	0.08
b (budget)	$5,000$	$2,500$	$10,000$

Table 2: Summary of simulation parameters.

Program	% of optimal revenue (standard deviation)	
	while rebate lasts (WRL)	entire duration (ED)
Fix-5%	83.1 (2.0)	82.7 (2.0)[6]
Fix-10%	86.4 (2.5)	98.0 (1.0)
Fix-15%	54.4 (1.6)	92.5 (1.8)
Hi-Lo	74.7 (2.7)	94.4 (1.7)
Adaptive	92.6 (2.4)	92.4 (2.4)
Linear	**97.8 (1.5)**	97.6 (1.6)
Log-linear	93.3 (4.2)	**98.8 (0.9)**

Table 3: % of optimal revenue achieved by different rebate programs under WRL and ED.

The default values for the parameters α_1, f, and t_{min} are chosen to give rise to realistic conversion rates. Under the range of parameter values considered, depending on the rebates offered, the conversion rates for the product fall in the range of 0.5% to 20%. These values are within ranges observed in online marketing [8]. The default value for traffic is set to 100. This selection is based on an educated guess, since this number could not be reliably estimated from the Cashback data. Consequently, we vary its value in a wide range as we test for sensitivity. When we vary traffic, we also vary budget at the same time to keep the average amount spent per potential customer constant.

The default value for budget is set to be $5,000$, corresponding to spending about 50 cents per potential customer, or \$5 per conversion, assuming a 10% conversion rate (for a product sold at \$100). We believe this value is a realistic estimate of the amount merchants are willing to pay per transaction.

7.3 Experiment 1: Single-Product Case

Under the default parameter values to the simulation, the average fractions of optimal revenue achieved (and their standard deviations) for the different rebate approaches are presented in Table 3.

Optimization under the linear and the log-linear models are respectively the best methods under WRL and ED, and achieves close to 100% of the optimal revenue. These are very strong performance numbers, especially considering that the optimal revenues are determined with the benefit of hindsight. The differences compared to other rebate programs are statistically significant in both cases under paired t-tests (with p-values < 0.0001).

Examining the results closer, we find two trends. On the one hand, rebate programs that use up more of the budget generally achieve higher revenue. On the other hand, if budget is exhausted too early, less revenue is achieved. Indeed, the strong performance of our approach can be attributed to the ability to balance the two factors. The importance of balancing these factors can be illustrated using the three fixed-rate programs, Fix-5%, Fix-10%, and Fix-15%. At a rebate rate of 5%, there is leftover budget, and one can improve revenue by setting a higher rate. At rebate rates of 10% and 15%, budget is fully exhausted. Comparing the two, we see that a rebate rate of 10% does better as it exhausts the budget later than a rebate rate of 15%.

At the default parameter values, the average optimal rebate rate is 8.7%. The average rebate rates computed using the linear model and the log-linear models are 8.4% and 9.3% respectively. They are both close to the optimal rates. The linear model performs better than the log-linear model under WRL as it does not exhaust the budget before the end. On the other hand, under ED, as revenue still counts after budget is exhausted and the merchant starts offering zero rebates, the log-linear model performs better as it puts the entire budget to work.

7.3.1 Sensitivity Analysis

We next vary each of simulation parameters (holding others to their default values) to understand their influence on the performance of our approach. The results are presented in Figure 1. In this figure, the panels on left right show the rebate rates computed by our approach using the linear and the log-linear demand models, along with the optimal rates, as a function of the parameter value. The panels on the left show the % of optimal revenue achieved. We have shown revenue plots only for the linear model under the WRL scenario and for the log-linear model under the ED scenario. It is because the linear (resp. log-linear) model consistently outperformed the log-linear (resp. linear) model under the WRL (resp. ED) scenario. Similarly, we only show plots for our approach since it outperformed the alternatives in more than 90% of cases.

We observe the following with respect to the rebate rates:

1. Overall, the rebate rates selected by our approach are very close to the optimal rates, with an average difference of less than 0.5%. Comparing the rebate rates chosen, the linear model tends to pick rates slightly below optimal, whereas the log-linear model picks ones slightly above.

2. When the traffic is very small (μ =10), the probability that there is no transaction for an entire day is 67%. This "missing" data problem manifests itself in adversely affecting demand estimation. The resultant error in the regression coefficients lead to the suboptimal values of rebate rate. Computed rebate rates start tracking the optimal with moderate increase in traffic (Figure 1(a)).

3. As expected, the computed rebate rates decrease as the sensitivity of the demand to the price of the product (parameter α_1) increases (Figure 1(c)). Similarly, rebate rates decrease as the relative value of rebate (parameter f) increases (Figure 1(e)).

4. As minimum conversion rate (parameter t_{min}) increases, the average number of transactions increases, and rebate rate has to decrease to match the budget (Figure 1(g)). The decrease is more significant for smaller values of t_{min} due to the inverse proportional relationship between conversion rate and rebate rate, which in turn is due to the budget constraint (number of transactions times rebate per transaction must be less than a fixed budget).

[6]Note that the optimal revenue depends on the scenario, and is higher under ED. Hence, we see that the fraction of optimal revenue achieved for Fixed-5% is lower under ED than under WRL even though its revenue stays the same.

Rebate rate sensitivity

Revenue sensitivity

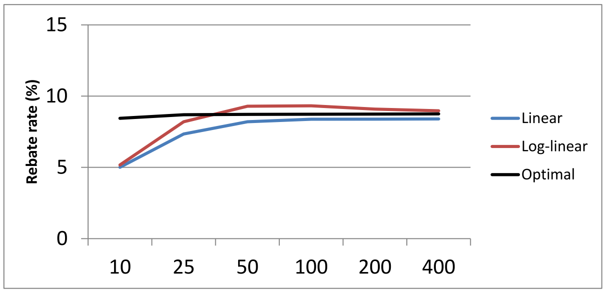

(a) Sensitivity to μ

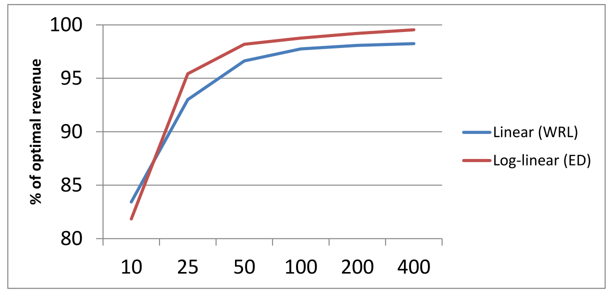

(b) Sensitivity to μ

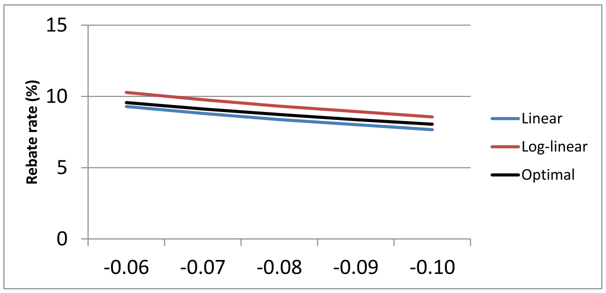

(c) Sensitivity to α_1

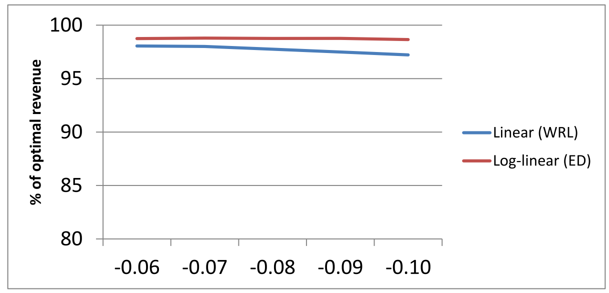

(d) Sensitivity to α_1

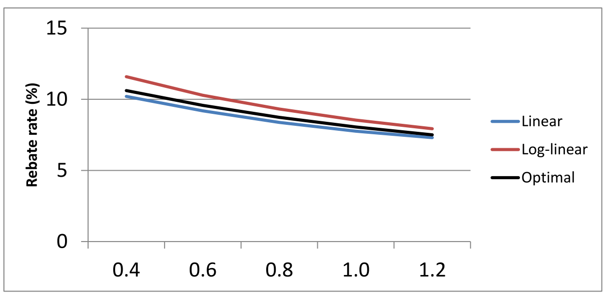

(e) Sensitivity to f

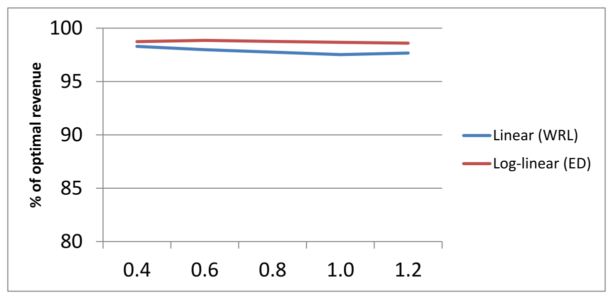

(f) Sensitivity to f

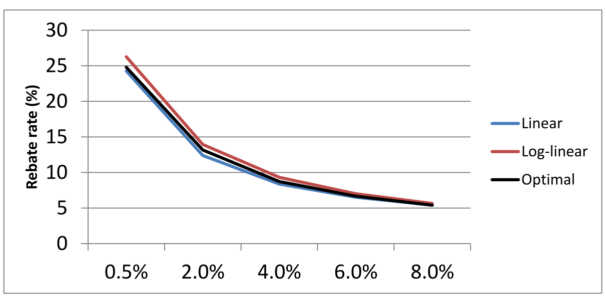

(g) Sensitivity to t_{min}

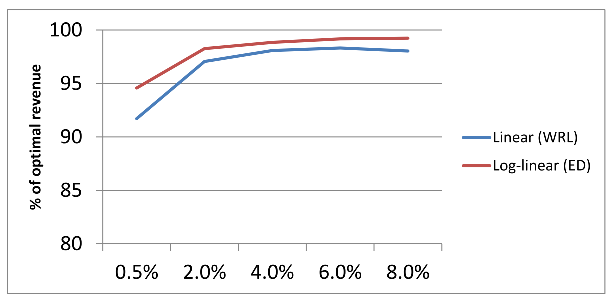

(h) Sensitivity to t_{min}

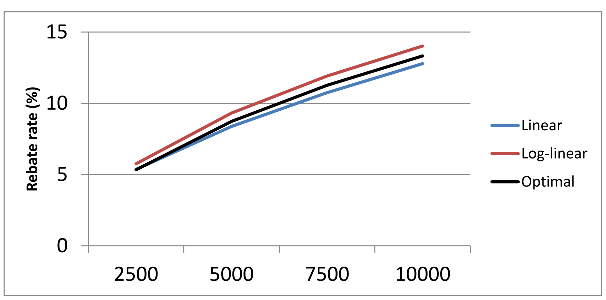

(i) Sensitivity to b

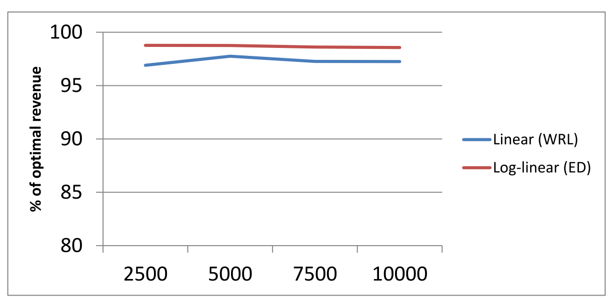

(j) Sensitivity to b

Figure 1: Sensitivity analysis.

5. As budget increases (parameter b) increases, the model picks up larger rebate rates as there is no incentive to leave the budget unspent (Figure 1(i)).

We observe the following with respect to the revenues:

1. Following the strategy of using the linear model when optimizing revenue under WRL, and adopting the log-linear model under ED, one does extremely well, achieving average % of optimal revenue of at least 95% in most cases. Only when the traffic is very small ($\mu = 10$) or when transactions are very rare ($t_{min} = 0.5\%$), our approach does not achieve 95% of maximum possible revenue (Figures 1(b) and (h)). The reasons are due to estimation errors because of "missing values", and that when transactions are rare, the relative value of each transaction increases, and so missing out a few transactions becomes more costly. But once we have sufficient data, and transactions are not rare, our approach starts performing at a very high level.

2. For the parameters α_1 and f, the performance of our approach is very consistent, with little variation in % of optimal revenue achieved across all choices of parameters (Figures 1(d) and (f)) . This is due to the data-driven nature of our approach. As these parameters vary, the rebate sensitivity of the product changes. By leveraging transaction data, our approach identifies these changes during demand estimation and optimize accordingly.

3. The performance of our approach is also very consistent for the budget parameter b. As can be seen in Figure 1(i), despite large variations in the optimal rates, our approach found rates that are close to the optimal ones for different budget values. Here, the reason is due to the algorithmic nature of our approach. As our approach takes budget as in input parameter and selects rebate rates through optimization, it can adapt well to different budgets well.

7.3.2 Summary

Our approaches for selecting rebate rates achieves close to the best possible revenue, and their performances are consistent across almost all choices of simulation parameters. Based on the experimental results, the linear model works better under WRL, and the log-linear model works better under ED. Both models manage to select rebate rates that are very close to the optimal ones.

7.4 Experiment 2: Multi-Product Case

In this experiment, we investigate whether our approach can identify products that are more sensitive to rebates and find the right trade-offs amongst the rebates offered for different products. We simulate a situation where a merchant carries two products. The first product, P1, has the same default parameter values as the ones used in the single product case. The parameters for the second product, P2, is identical in all aspects except for f, its rebate sensitivity. Holding f for the first product constant, we vary f of the second product.

For the base case, f of P2 is set at 0.4. This is smaller than f of P1, which is 0.8, and hence P2 is less responsive to rebates than P1. The revenues achieved under different approaches are shown in Table 4.

[7]When there are two or more products, we do not know how to compute the optimal revenue under ED efficiently. Based on our experiments on single products, the optimal revenue is < 1% higher than under WRL. Hence we use the optimal revenue under WRL as the approximate benchmark.

Program	% of optimal revenue (standard deviation)	
	while rebate lasts (WRL)	entire duration (ED)[7]
Fix-5%	81.4 (1.5)	81.4 (1.5)
Fix-10%	91.0 (2.2)	95.5 (1.0)
Fix-15%	57.3 (1.4)	91.7 (1.3)
Hi-Lo	80.9 (2.4)	93.1 (1.2)
Adaptive	88.9 (3.7)	93.4 (1.5)
Linear	**97.1 (1.3)**	97.2 (1.4)
Log-linear	92.8 (3.7)	**98.6 (0.9)**

Table 4: % of upper bound on revenue achieved by different rebate programs with two products.

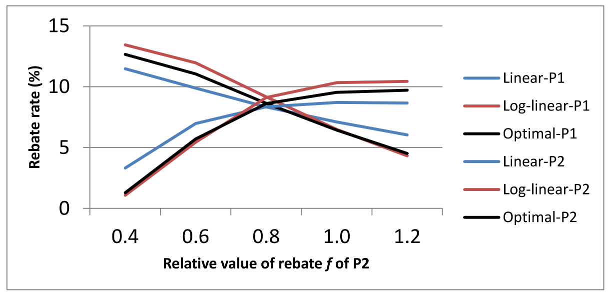

Figure 2: Rebate rates selected under the two demand models compared to optimal, as the rebate sensitivity f of P2 varies.

Like the case of a single product, our approach using the linear and the log-linear models achieves the highest % of optimal revenue respectively under WRL and ED. The differences in performance are statistically significant under a paired t-test (with p-value < 0.0001). The differences are larger in this experiment compared to the case with only a single product. This is due to the importance of offering different rebate rates on products with different rebate sensitivities. For the current setting, the optimal rebate rate is about 12.7% for P1 and 1.2% for P2. A single rebate rate will fail to take into account these differences.

To complete the experiment, we vary the rebate sensitivity f of P2 and compare the computed rebate rates using coefficients produced by the linear and the log-linear models to the optimal rates. The result is presented in Figure 2. As P2 becomes more sensitive to rebates, both approaches correctly increase the rebates rates for P2, and they closely mirror the optimal rates for *both* products. As a sanity check, when the rebate sensitivity of both products are equal, both the optimal rebate rate and the rates selected by our approach are roughly equal as well.

7.4.1 Summary

Our approach performs relatively even better when there are multiple products. The demand estimation step correctly identifies the product more sensitive to rebates, and the optimization selects correspondingly higher rebate rate for the more sensitive product. The approach is robust to changes in rebate sensitivities, and strikes a good balance among the rates selected for different products.

8. CONCLUDING REMARKS

We studied the problem of how online merchants can best use rebates to maximize their revenue. Our solution consists of two steps—an estimation step and an optimization step. Our estimation routine builds on classical demand models in economics, and extends them to model the effect of rebates separately from that of prices. We develop efficient solutions to the optimization problem under both the linear and the log-linear demand model, drawing upon ideas from convex optimization. We validated our modeling assumptions using transaction data obtained from Microsoft Cashback operations, and conducted an extensive simulation study to evaluate the performance of our proposed approach. We found that across a wide range of parameters, our approach consistently generates higher revenue than other approaches, and achieves close to the maximum possible revenue.

Through these simulation studies, we found that selecting good rebate rates requires carefully balancing two factors—putting the entire budget to use and spreading the budget over the entire period. Our approach does well in balancing these factors, and hence performs better than other approaches. Between the linear and the log-linear model, the former is more suited for the scenario where revenues are measured while rebate lasts (WRL), whereas the latter is more suited for the scenario where revenues are measured over the entire duration (ED). The rebate rates selected are often within 1% of the optimal rates. We also note that the performances of the linear and the log-linear models are very close both in real and synthetic data. This is an interesting and somewhat surprising finding, as economists have often favored the log-linear model over the linear model for demand, and merits further investigation.

The optimization approach we presented in this paper can be extended in several ways. For example, it can be used to maximize profits instead of revenue by taking into account the cost of production. It can also be used to solve more sophisticated problems that include additional constraints such as minimum and maximum rebate rate per product or limits on the supply of each product.

There are important future directions to explore. One direction is techniques for optimizing rebates when demand of one product may be affected by prices and rebates of others. From the estimation standpoint, this presents a challenge due to its requirement for large volume of data. The optimization problem can no longer be formulated as a convex program and new techniques will be needed.

In this work, we do not consider the revenue generated during the period when a merchant is learning a demand model for the product. If this period is considered as part of the evaluation, merchants face a new problem that may require interleaving exploration for additional data and exploiting the demand model. Solving this problem optimally (or approximately optimally) will require new techniques.

9. ACKNOWLEDGMENTS

We thank Arun Sacheti and Larry Colagiovanni for helpful discussions and provision of data for this research, and the anonymous reviewers for their helpful feedback.

10. REFERENCES

[1] F. J. Arcelus, S. Kumar, and G. Srinivasan. Retailer's response to alternate manufacturer's incentives under a single period, price-dependent, stochastic-demand framework. *Decision Sciences*, 36:599–626, 2005.

[2] F. J. Arcelus, S. Kumar, and G. Srinivasan. Pricing and rebate policies for the newsvendor problem in the presence of a stochastic redemption rate. *Int. J. Production Economics*, 107:467–482, 2007.

[3] F. J. Arcelus, S. Kumar, and G. Srinivasan. Price and rebate policies in the two-echelon supply chain with asymmetric information under price-dependent, stochastic demand. *Int. J. Prod. Econ.*, 113, 2008.

[4] G. Box, G. M. Jenkins, and G. Reinsel. *Time Series Analysis: Forecasting and Control*. Prentice Hall, 1994.

[5] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[6] Y. Chen, A. Ghosh, P. McAfee, and D. Pennock. Sharing online advertising revenue with consumers. In *WINE*, pages 556–565, 2008.

[7] J. Dahl and L. Vandenberghe. `CVXOPT`: A python package for convex optimization, 2008.

[8] B. Eisenberg. The average conversion rate: Is it a myth? `http://www.clickz.com/3628276`, seen on August 1, 2010.

[9] V. Folkes and R. D. Wheat. Consumers' price perceptions of promoted products. *J. Retailing*, 1995.

[10] A. Ghose and S. Yang. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Sci.*, 55(10):1605–1622, 2009.

[11] S. M. Gilpatric. Slippage in rebate programs and present-biased preferences. *Marketing Science*, 2008.

[12] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *SODA*, pages 735–744, 2001.

[13] B. J. Jansen and T. Mullen. Sponsored search: an overview of the concept, history, and technology. *Int. J. Electronic Business*, 6(2):114–131, 2008.

[14] M. Khouja. A joint optimal pricing, rebate value and lot sizing model. *European Journal of Operational Research*, 174:706–723, 2006.

[15] M. Khouja, M. Hadzikadic, and M. A. Zafffar. An agent based modeling approach for determining optimal price-rebate schemes. *Simulation Modeling: Practice and Theory*, 16:111–126, 2008.

[16] S. Lahaie, D. M. Pennock, A. Saberi, and R. V. Vohra. Sponsored search auctions. In *Algorithmic Game Theory*, chapter 28. Cambridge University Press, 2007.

[17] Q. Lu and S. Moorthy. Coupons versus rebates. *Marketing Science*, 26(1):67–82, 2007.

[18] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[19] T. T. Nagle and J. E. Hogan. *The Strategy and Tactics of Pricing*. Prentice Hall, 4th edition, 2006.

[20] H. Nazerzadeh, A. Saberi, and R. Vohra. Dynamic cost-per-action mechanisms and applications to online advertising. In *WWW*, pages 179–187, 2008.

[21] S. Ross. *A First Course in Probability*. Prentice Hall, 6 edition, 2001.

[22] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming*. SIAM, 2009.

[23] H. L. R. Tierney and B. Pan. A Poisson regression examination of the relationship between website traffic and search engine queries. MPRA Paper 19895, 2010.

[24] K. Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2003.