
1 ALGORITHMS FOR MARKETING-MIX OPTIMIZATION

2 *Joachim Gudmundsson Pat Morin Michiel Smid*

3 ABSTRACT. Algorithms for determining quality/cost/price tradeoffs in saturated markets are consid-
4 ered. A product is modeled by d real-valued qualities whose sum determines the unit cost of producing
5 the product. This leads to the following optimization problem: given a set of n customers, each of
6 whom has certain minimum quality requirements and a maximum price they are willing to pay, design
7 a new product and select a price for that product in order to maximize the resulting profit.

8 An $O(n \log n)$ time algorithm is given for the case, $d = 1$, of products having a single quality,
9 and $O(n(\log n)^{d+1})$ time approximation algorithms are given for products with any constant number,
10 d , of qualities. To achieve the latter result, an $O(nk^{d-1})$ bound on the complexity of an arrangement
11 of homothetic simplices in \mathbb{R}^d is given, where k is the maximum number of simplices that all contain
12 a single points.

13 1 Introduction

14 Revealed preference theory [13] is a method of determining a course of business action through the
15 review of historical consumer behaviour. In particular, it is a method of inferring an individual's or a
16 group's preferences based on their past choices. The *marketing mix* [10] of a product consists of the 4
17 Ps: Product, price, place, and promotion. In the current paper, we present algorithms for optimizing
18 the first two of these by using data about consumers' preferences. That is, we show how, given data on
19 consumer preferences, to efficiently choose a product and a price for that product in order to maximize
20 profit.

21 Refer to Figure 1. A product $P = (p, q_1, \dots, q_d)$ is defined by a real-valued *price*, p , and
22 a number of real-valued orthogonal *qualities*, q_1, \dots, q_d . The *market* for a product is a collection
23 of customers $C = \{C_1, \dots, C_n\}$, where $C_i = (p_i, q_{i,1}, \dots, q_{i,d})$. A customer will purchase the least
24 expensive product that meets all her minimum quality requirements and whose price is below her
25 maximum price. That is, the customer C_i will *consider* the product $P = (p, q_1, \dots, q_d)$ if $p \leq p_i$ and
26 $q_j \geq q_{i,j}$ for all $j \in \{1, \dots, d\}$. The customer C_i will *purchase* the product if it has the minimum price
27 among all available products that C_i considers.

28 We consider markets that are *saturated*. That is, for every customer C_i there is an existing
29 product that satisfies C_i 's requirements and among all products that satisfy C_i 's requirements, C_i will
30 choose the least expensive product. From the point of view of a manufacturer introducing one or more
31 new products, this means that all customers are *Pareto optimal*, i.e., there are no two customers C_i
32 and C_j such that $q_{i,k} > q_{j,k}$ for all $k \in \{1, \dots, d\}$ and $p_i < p_j$. This is because no customer will
33 every purchase a product that is not Pareto optimal, since there is a lower-priced alternative that also




Figure 1: A sample market with $d = 1$ and $n = 3$. A customer will consider any product that is in their upper left quadrant.

34 satisfies all their minimum quality requirements. Therefore, every customer C_i can be replaced with
 35 the (Pareto optimal) product that they purchase.

36 As an example, consider a market for computers in which an example customer C_i may be
 37 looking for a computer with a minimum of 8 GB of RAM, a CPU benchmark score of at least 3000,
 38 a GPU benchmark score of at least 2000, and be willing to pay at most \$1500. In addition, there is
 39 already a computer on the market which meets these requirements and retails for \$1200. Thus, this
 40 customer would be described by the vector $(1200, 8, 3000, 2000)$. If a manufacturer introduces a new
 41 product $(1199, 8, 3500, 2000)$ (a computer with 8 GB of RAM, a CPU benchmark score of 3500 and a
 42 GPU benchmark score of 2000 retailing for \$1199) then this customer would select this new product
 43 over their current choice.

44 By appropriately reparameterizing the axes, we can assume that the cost, $\text{cost}(P)$, of manufac-
 45 turing a product $P = (p, q_1, \dots, q_d)$ is equal to the sum of its qualities

$$46 \quad \text{cost}(P) = \sum_{i=1}^d q_i .$$

47 The *profit per unit sold* of P is therefore

$$48 \quad \text{ppu}(P) = p - \text{cost}(P) .$$

49 In this paper we consider algorithms that a manufacturer can use when planning a new product to
 50 introduce into an existing saturated market with the goal being to maximize the profit obtained. More
 51 precisely, given a Pareto-optimal *market* of customers $M = \{C_1, \dots, C_n\}$, each having d qualities, the
 52 $\text{PRODUCTDESIGN}(d)$ problem is to find a product $P^* \in \mathbb{R}^{d+1}$ such that

$$53 \quad \text{profit}(P^*) = \text{ppu}(P^*) \times |\{i : C_i \text{ purchases } P^*\}|$$

54 is maximized.

55 The term *marketing mix* is probably the most famous phrase in marketing. Not surprisingly,
 56 economists and market researchers have considered methods of optimizing the marketing mix in various
 57 scenarios. As there are many different models of the problem, it is difficult to compare algorithms.

58 Most models of marketing-mix problems involve a constant number of real-valued input pa-
59 rameters. This sometimes leads to problems where the optimal solution is one of a constant number
60 of possible closed forms (see, e.g., Thomson and Teng’s optimal constant price model [12]). In other
61 cases, a closed form is not achievable, but a (sometimes approximate) solution can be obtained using
62 numerical optimization techniques (see, e.g., Balanchandran and Gensch [2], Thomson and Teng [12],
63 Naik *et al.* [11], Deal [6], Erickson [9]). In all cases, it is expected that the model parameters are derived
64 from real-world data, such as surveys or sales figures, and involves fitting of the model parameters to
65 the available data.

66 The work in the current paper is different from this previous work in several ways. For one, it
67 is one of the few works that deals primarily with the first two P’s, product and pricing. Most existing
68 literature focuses on the marketing P’s, namely place and promotion, and to a lesser extent, pricing
69 of an existing product. Secondly, it deals directly with data about individual consumers rather than
70 aggregating this data so that it fits a particular model of consumer behaviour.

71 We believe that this models very well what happens in online shopping for high cost products
72 such as computers, cameras, and televisions. In such markets, savvy consumers have good data available
73 about both the specifications and the cost of all available products so that marketing efforts are
74 (arguably) less important than the quality and prices of the products. On the other hand, online
75 sellers such as Amazon have large amounts of data about users’ past purchases and can use this data
76 as input to the problem. In particular, these sellers know the specifications $q_{i,1}, \dots, q_{i,d}$ and prices p_i
77 of huge quantities of items sold and can use this data to advise a manufacturer that is designing a new
78 product.

79 In the remainder of the paper we give an $O(n \log n)$ time algorithm for PRODUCTDESIGN(1)
80 (Section 2), and $O(n(\log n)^{d+1})$ time approximation schemes for PRODUCTDESIGN(d) (Section 3 and
81 Section 4). Section 5 summarizes our results and concludes with directions for future research.

82 2 One-dimensional products

83 In this section, we consider the simplest case, when a manufacturer wishes to introduce a new product
84 in which the quality of a product has only one dimension. Examples of such markets include, for
85 example, suppliers to the construction industry in which items (steel I-beams, finished lumber, logs)
86 must have a certain minimum length to be used for a particular application. An overly long piece can
87 be cut down to size, but using two short pieces instead of one long piece is not an option.

88 Throughout this section, since $d = 1$, we will use the shorthand $P = (p, q)$ for the product
89 being designed and q_i for $q_{i,1}$. Thus, we have a set of customers $M = \{(p_1, q_1), \dots, (p_n, q_n)\}$ and we
90 are searching for a point $P^* = (p^*, q^*)$ that maximizes

$$91 \text{profit}(p^*, q^*) = (p^* - q^*) |\{i : p^* \leq p_i \text{ and } q^* \geq q_i\}| .$$

92 Our algorithm is an implementation of the *plane-sweep* paradigm [4]. The correctness of the
93 algorithm relies on two lemmas about the structure of the solution space. The first lemma is quite
94 easy:

95 **Lemma 1.** *The value (p^*, q^*) that maximizes $\text{profit}(p^*, q^*)$ is obtained when $p^* = p_i$ and $q^* = q_j$ for*
96 *some $i, j \in \{1, \dots, n\}$.*




Figure 2: $\text{profit}(p, q) \leq \text{profit}(p, q')$ implies that $\text{profit}(p', q) \leq \text{profit}(p', q')$ for all $p' \leq p$.

97 *Proof.* First, observe the obvious bounds on p^* and q^* :

$$98 \quad \min\{p_i : i \in \{1, \dots, n\}\} \leq p^* \leq \max\{p_i : i \in \{1, \dots, n\}\}$$

99 and

$$100 \quad \min\{q_i : i \in \{1, \dots, n\}\} \leq q^* \leq \max\{q_i : i \in \{1, \dots, n\}\} .$$

101 Consider the arrangement of lines obtained by drawing a horizontal and vertical line through each
 102 customer (p_i, q_i) for $i \in \{1, \dots, n\}$. Within each cell of this arrangement, the function $\text{profit}(p, q)$ is
 103 a linear function of p and q and it is bounded. Therefore, within a particular cell, the function is
 104 maximized at a vertex. Since each vertex is the intersection of a horizontal and vertical line through a
 105 pair of customers, the lemma follows. \square

106 The following lemma, illustrated in Figure 2, is a little more subtle and illustrates a manufac-
 107 turer's preference for lower-quality products:

108 **Lemma 2.** *Let $q' \leq q$ and let p be such that $0 < \text{profit}(p, q) \leq \text{profit}(p, q')$. Then, for any $p' \leq p$,*
 109 *$\text{profit}(p', q) \leq \text{profit}(p', q')$.*

110 *Proof.* By definition, $\text{profit}(p, q) = a(p - q)$ and $\text{profit}(p, q') = a'(p - q')$, where a and a' are the number
 111 of customers who would consider (p, q) and (p, q') , respectively. These customers are all taken from the
 112 set $M_{\geq} = \{(p_i, q_i) \in M : p_i \geq p\}$.

113 Now, consider the customers in the set $M' = \{(p_i, q_i) \in M : p' \leq p_i < p\}$. By the assumption
 114 that customers are Pareto optimal, any customer (p_i, q_i) in M' has $q_i \leq q'$, so all of these customers
 115 will consider either (p', q') or (p', q) if either one is offered. Therefore,

$$\begin{aligned}
 \text{profit}(p', q') &= (a' + |M'|)(p' - q') \\
 &= a'(p' - q') + |M'|(p' - q') \\
 &\geq a'(p' - q') + |M'|(p' - q) && \text{since } q > q' \\
 &= a'(p - q') + a'(p' - p) + |M'|(p' - q) \\
 &\geq a'(p - q') + a(p' - p) + |M'|(p' - q) && \text{since } a \geq a' \text{ and } (p' - p) < 0 \\
 &\geq a(p - q) + a(p' - p) + |M'|(p' - q) && \text{by assumption} \\
 &= a(p' - q) + |M'|(p' - q) \\
 &= \text{profit}(p', q) ,
 \end{aligned}$$

117 as required. □

118 Lemma 2 allows us to apply the plane sweep paradigm with a sweep by decreasing price. It
119 tells us that, if a product (p, q') gives better profit than the higher-quality product (p, q) at the current
120 price p , then it will always give a better profit for the remainder of the sweep. In particular, there will
121 never be a reason to consider a product with quality q for the remainder of the algorithm's execution.

122 Let the customers be labelled $(p_1, q_1), \dots, (p_n, q_n)$ in decreasing order of p_i , so that $p_{i+1} \leq p_i$
123 for all $i \in \{1, \dots, n-1\}$. At any point in the sweep algorithm, there is a current price p , which
124 starts at $p = \infty$ and decreases during the execution of the algorithm. At the start of the algorithm the
125 algorithm's event queue Q , which is represented as a balanced binary search tree, is initialized to contain
126 the values p_n, \dots, p_1 . At all times, the algorithm maintains a list L of qualities $q_1^* > q_2^* > \dots > q_m^*$
127 such that $\text{profit}(p, q_1^*) > \text{profit}(p, q_2^*) > \dots > \text{profit}(p, q_m^*)$. The quality q_1^* is the optimal quality for
128 the current price, p . By the time the algorithm terminates, the quality of the globally-optimal solution
129 will have appeared as the first element in L . To complete the description of the algorithm, all that
130 remains is to show how L and Q are updated during the processing of events in the event queue.

131 There are two kinds of events in the event queue. *Insertions* occur at the values p_1, \dots, p_n .
132 *Deletions*, which we describe next, occur when the relative order of two adjacent items in L changes.
133 Consider a consecutive pair of the elements q_i^* and q_{i+1}^* in L . When q_i^* and q_{i+1}^* became adjacent in L ,
134 it was at some price $p = p_t$ such that $\text{profit}(p_t, q_i^*) > \text{profit}(p_t, q_{i+1}^*)$. Let a_i and a_{i+1} be the number of
135 customers who would consider (p_t, q_i^*) and (p_t, q_{i+1}^*) , respectively. Then,

$$136 \quad \text{profit}(p_t, q_i^*) = (p_t - q_i^*)a_i$$

137 and

$$138 \quad \text{profit}(p_t, q_{i+1}^*) = (p_t - q_{i+1}^*)a_{i+1}$$

139 Now, looking forward in time to a later step in the execution of the algorithm, when $p = p_{t'}$, with
140 $t' > t$, we find that

$$141 \quad \text{profit}(p_{t'}, q_i^*) = (p_{t'} - q_i^*)(a_i + t' - t)$$

142 and

$$143 \quad \text{profit}(p_{t'}, q_{i+1}^*) = (p_{t'} - q_{i+1}^*)(a_{i+1} + t' - t) .$$

144 We are interested in identifying the price $p_{t'}$ where the inequality $\text{profit}(p_{t'}, q_i^*) > \text{profit}(p_{t'}, q_{i+1}^*)$
145 changes to $\text{profit}(p_{t'}, q_i^*) \leq \text{profit}(p_{t'}, q_{i+1}^*)$. When this happens, q_i^* can be safely discarded from L
146 since, by Lemma 2, $\text{profit}(p, q_i^*)$ will never again exceed $\text{profit}(p, q_{i+1}^*)$ for the remainder of the sweep.
147 The value $p_{t'}$ is a deletion event.

148 Note that Lemma 2 also allows for binary search on the value $p_{t'}$. In particular, the interval
149 $[p_j, p_{j+1}]$ containing $p_{t'}$ can be found in $O(\log n)$ time, after which the value of $p_{t'}$ can be obtained by
150 solving the linear equation

$$151 \quad (p_{t'} - q_i^*)(a_i + t' - t) \leq (p_{t'} - q_{i+1}^*)(a_{i+1} + t' - t) ,$$

152 for t' . (The equation is linear because the values a_i and a_{i+1} are constant in the interval (p_j, p_{j+1}) .)
153 Thus, whenever two new elements become adjacent in L , we can add the appropriate deletion event to
154 Q in $O(\log n)$ time.

155 When processing an insertion event p_i , we remove from the tail of L all values q_j^* such that
156 $\text{profit}(p_i, q_j^*) \leq \text{profit}(p_i, q_i)$ and then append q_i onto L . While deleting the elements of L , we also

157 remove all the associated deletion events from Q . Appending q_i to L causes at most one new pair of
158 elements in L to become adjacent, so we add the appropriate deletion event to Q as described above.
159 The time to process the event p_i is there $O((k_i + 1) \log n)$, where k_i is the number of elements that are
160 deleted from L .

161 When processing a deletion event that deletes q_i^* from L , we simply delete q_i^* from L and its (at
162 most two) associated deletion events from Q . This may cause a new pair of elements, q_{i-1}^* and q_{i+1}^* ,
163 in L to become adjacent, so we add the appropriate deletion event to Q . In this way, a deletion event
164 can be processed in $O(\log n)$ time.

165 Note that, after all the processing associated with an event p_t is complete, the first element, q_1^* ,
166 in L is the value that maximizes $\text{profit}(p_t, q_1^*)$. Thus, the algorithm need only keep track, throughout
167 its execution, of the highest profit obtained from the first element of L , and output this value at the
168 end of its execution. This completes the description of the algorithm.

169 **Theorem 1.** *There exists an $O(n \log n)$ time algorithm for PRODUCTDESIGN(1).*

170 *Proof.* The correctness of the algorithm described above follows from 2 facts: Lemma 1 ensures that
171 the optimal solution is of the form (p_i, q_i^*) for some $i \in \{1, \dots, n\}$, and Lemma 2 ensures that the
172 optimal solution appears at some point as the first element of the list L .

173 The running time of the algorithm can be bounded as follows: The total number of deletion
174 events processed is at most n , since each such event removes some value q_i from L for some $i \in$
175 $\{1, \dots, n\}$. Thus, the cost of processing all deletion events is $O(n \log n)$. Similarly, $\sum_{i=1}^n k_i \leq n$ since
176 k_i is the number of elements deleted from L during the processing of p_i . Thus, the time required to
177 handle all insertion events, and therefore the running time of the entire algorithm, is $O(n \log n)$ \square

178 The following theorem shows that a running time of $\Omega(n \log n)$ is inherent in this problem, even
179 when considering approximation algorithms.

180 **Theorem 2.** *Let M be an instance of PRODUCTDESIGN(1) and (p^*, q^*) be a solution that maximizes*
181 *profit(p^*, q^*). In the algebraic decision tree model of computation, any algorithm that can find a solution*
182 *(p, q) such that $2 \cdot \text{profit}(p, q) > \text{profit}(p^*, q^*)$ has $\Omega(n \log n)$ running time in the worst-case.*

183 *Proof.* We reduce from the integer ELEMENT-UNIQUENESS problem, which has an $\Omega(n \log n)$ lower
184 bound in the algebraic decision tree model [14]: Given an array $A = [x_1, \dots, x_n]$ containing n integers,
185 are all the elements of A unique?

186 We convert A into an instance of PRODUCTDESIGN(1) in $O(n)$ time as follows (refer to Figure 3).
187 For each x_i , $i \in \{1, \dots, n\}$ we introduce a customer (p_i, q_i) with $p_i = q_i + 1/2$ and $q_i = x_i$. If
188 there exists a value x in A that occurs 2 or more times, then the product $(x + 1/2, x)$ gives a value
189 $\text{profit}(x + 1/2, x) \geq 1$. On the other hand, if there is no such x , then

- 190 1. any product (p, q) with $p - q > 1/2$ can not be sold to any customers and
- 191 2. any product (p, q) with $p - q > 0$ can be sold to at most 1 customer.




Figure 3: Reducing ELEMENT-UNIQUENESS to PRODUCTDESIGN(1,1).

192 Therefore, if all the elements of A are unique, then $\text{profit}(p^*, q^*) = 1/2$, otherwise $\text{profit}(p^*, q^*) \geq 1$.
 193 The result follows. \square

194 3 A near-linear approximation algorithm for bidimensional products

195 In this section, we consider algorithms for PRODUCTDESIGN(2), in which products have 2 qualities. As
 196 a baseline, we first observe that, if we fix the value of q_2 , then the optimal solution of the form (p, q_1, q_2)
 197 can be found using a single application of the algorithm in Theorem 1. Therefore, by successively solving
 198 the problem for each $q_2 \in \{q_{2,1}, \dots, q_{2,n}\}$ and taking the best overall solution we obtain an $O(n^2 \log n)$
 199 time algorithm for PRODUCTDESIGN(2).

200 More generally, PRODUCTDESIGN(d) can be solved using $O(n^{d-1})$ applications of Theorem 1
 201 resulting in an $O(n^d \log n)$ time algorithm. Unfortunately, these are the best results known for $d \geq 2$,
 202 and, as discussed in Section 5, we suspect that an algorithm with running time $o(n^d)$ will be difficult
 203 to achieve using existing techniques. Therefore, in this section we focus our efforts on obtaining a
 204 near-linear approximation algorithm.

205 Fix some constant $\epsilon > 0$. Given an instance M of PRODUCTDESIGN(d), a point $P \in \mathbb{R}^{d+1}$ is a
 206 $(1 - \epsilon)$ -approximate solution for M if $\text{profit}(P) \geq (1 - \epsilon) \text{profit}(P^*)$ for all $P^* \in \mathbb{R}^{d+1}$. An algorithm
 207 is a (high probability) *Monte-Carlo* $(1 - \epsilon)$ -approximation algorithm for PRODUCTDESIGN(d) if, given
 208 an instance M of size n , the algorithm outputs a $(1 - \epsilon)$ -approximate solution for M with probability
 209 at least $1 - n^{-c}$ for some constant $c > 0$.

210 Let $r = \max\{\text{ppu}(C_i) : i \in \{1, \dots, n\}\}$ and observe that r is the maximum profit per unit
 211 that can be achieved in this market. Let $E = 1/(1 - \epsilon)$ and let $\ell = \lceil \log_E n \rceil$ and observe that
 212 $\ell = O(\epsilon^{-1} \log n)$.¹ For each $i \in \{0, 1, 2, \dots, \ell\}$, define the plane $H_i = \{(p, q_1, q_2) : p - q_1 - q_2 = r(1 - \epsilon)^i\}$.
 213 The following lemma says that a search for an approximate solution can be restricted to be contained
 214 in one of the planes H_i .

215 **Lemma 3.** *For any product $P^* = (p^*, q_1^*, q_2^*)$, there exists a product $P = (p, q_1, q_2)$ such that $P \in H_i$
 216 for some $i \in \{0, \dots, \ell\}$ and $\text{profit}(P) \geq (1 - \epsilon) \text{profit}(P^*)$.*

¹This can be seen by taking the limit $\lim_{\epsilon \rightarrow 0^+} (\epsilon / \log(E))$ using one application of L'Hôpital's Rule.




Figure 4: The intersection of H_i with customers' quadrants is a set of homothetic equilateral triangles.

217 *Proof.* There are two cases to consider. If $\text{ppu}(P^*) \leq r/n$ then $\text{profit}(P^*) \leq r$, in which case we set
 218 $P = C_i$ where $\text{ppu}(C_i) = r$, so that $P \in H_0$ and $\text{profit}(P) = r \geq \text{profit}(P^*) \geq (1 - \epsilon) \text{profit}(P^*)$, as
 219 required.

220 Otherwise, $r/n < \text{ppu}(P^*) \leq r$. In this case, consider the plane H_i where $i = \lceil \log_E(r / \text{ppu}(P^*)) \rceil$.
 221 Notice, that for any point $P \in H_i$, $\text{ppu}(P) \geq (1 - \epsilon) \text{ppu}(P^*)$. More specifically, the orthogonal pro-
 222 jection $P = (p, q_1, q_2)$ of P^* onto H_i is a product with $p \leq p^*$, $q_1 \geq q_1^*$, and $q_2 \geq q_2^*$. Therefore,
 223 any customer who would consider P^* would also consider P , so $\text{profit}(P) \geq (1 - \epsilon) \text{profit}(P^*)$, as
 224 required. \square

225 Lemma 3 implies that the problem of finding an approximate solution to PRODUCTDESIGN(2)
 226 can be reduced to a sequence of problems on the planes H_0, \dots, H_ℓ . Refer to Figure 4. Each customer
 227 C_j considers all products in a quadrant whose corner is C_j . The intersection of this quadrant with
 228 H_i is a (possibly empty) equilateral triangle $\Delta_{i,j}$. The customer C_j will consider a product P in H_i
 229 if and only if P is in $\Delta_{i,j}$. Thus, the problem of solving PRODUCTDESIGN(2) restricted to the plane H_i
 230 is the problem of finding a point contained in the largest number of equilateral triangles from the set
 231 $\Delta_i = \{\Delta_{i,j} : j \in \{1, \dots, n\}\}$.

232 Note that the elements in Δ_i are *homothets* (translations and scalings) of an equilateral triangle,
 233 so they form a collection of *pseudodisks* and we wish to find the deepest point in this collection of
 234 pseudodisks. No algorithm with running time $o(n^2)$ is known for solving this problem exactly, but
 235 Aronov and Har-Peled [1] have recently given a Monte-Carlo $(1 - \epsilon)$ -approximation algorithm for this
 236 problem that runs in time $O(\epsilon^{-2} n \log n)$. By applying this algorithm to each of Δ_i for $i \in \{1, \dots, \ell\}$,
 237 we obtain the following result:

238 **Theorem 3.** *For any $\epsilon > 0$, there exists an $O(\epsilon^{-3} n (\log n)^2)$ time high-probability Monte-Carlo $(1 - \epsilon)$ -*
 239 *approximation algorithm for PRODUCTDESIGN(2).*

240 4 A near-linear approximation algorithm for constant d

241 In this section we extend the algorithm from the previous section to (approximately) solve $\text{PRODUCTDESIGN}(d)$
 242 for any constant value of d . The algorithm is more or less unchanged, except that the proof requires
 243 some new results on the combinatorics of arrangements of homothets.

244 As before, let $r = \max\{\text{ppu}(C_i) : i \in \{1, \dots, n\}\}$ and let $\ell = \lceil \log_E n \rceil$. For each $i \in$
 245 $\{0, 1, 2, \dots, \ell\}$, define the hyperplane $H_i = \{(p, q_1, \dots, q_d) : p - \sum_{i=1}^d q_i = r(1 - \epsilon)^i\}$. The follow-
 246 ing lemma has exactly the same proof as Lemma 3.

247 **Lemma 4.** *For any product $P^* = (p^*, q_1^*, \dots, q_d^*)$, there exists a product $P = (p, q_1, \dots, q_d)$ such that*
 248 *$P \in H_i$ for some $i \in \{0, \dots, \ell\}$ and $\text{profit}(P) \geq (1 - \epsilon) \text{profit}(P^*)$.*

249 Again, each customer C_j defines a regular simplex $\Delta_{i,j}$ in H_i such that C_j will consider $P \in H_i$
 250 if and only if $P \in \Delta_{i,j}$. In this way, we obtain a set $\Delta_i = \{\Delta_{i,1}, \dots, \Delta_{i,n}\}$ of homothets of a regular
 251 simplex in \mathbb{R}^d and we require an algorithm to find a $((1 - \epsilon)$ -approximation to) the point that is
 252 contained in the largest number of these simplices. The machinery of Aronov and Har-Peled [1] can be
 253 used to help solve this problem, but not before we prove some preliminary results, the first of which is
 254 a combinatorial geometry result.

255 **Lemma 5.** *Let Δ be a set of n homothets of a regular simplex in \mathbb{R}^d , for $d = O(1)$, and such that no*
 256 *point in \mathbb{R}^d is contained in more than k elements of Δ . Then, the total complexity of the arrangement,*
 257 *$A(\Delta)$, of the simplices in Δ is $O(nk^{d-1})$.*

258 *Proof.* We first consider the simpler case in which the elements of Δ are translates (without scaling) of
 259 a regular simplex. Suppose that the total complexity of $A(\Delta)$ is m . Then, by the pigeonhole principle,
 260 there is some element T in Δ whose surface is involved in m/n features of $A(\Delta)$. (Note that this implies
 261 that T intersects all the elements of a set $\Delta' \subseteq \Delta$ with $|\Delta'| = \Omega((m/n)^{1/(d-1)})$, since otherwise there
 262 are not enough elements in Δ' to generate m/n features on the surface of Δ .)

263 Observe that, since the elements of Δ' are all unit size and they all intersect T , that they are
 264 all contained in a ball of radius $O(1)$ centered at the center of T . Furthermore, since each element of
 265 Δ' has volume $\Omega(1)$ this implies that some point must be contained in $\Omega((m/n)^{1/(d-1)})$ elements of Δ' .
 266 Thus, we obtain the inequality $k \geq \Omega((m/n)^{1/(d-1)})$, or, equivalently, $m \leq O(nk^{d-1})$, as required.

267 Now, for the case where the elements of Δ are homothets (translations and scalings) of a regular
 268 simplex, we proceed as follows. Suppose that $|A(\Delta)| = rn$. Our goal is to show that $r = O(k^{d-1})$.
 269 Label the elements of Δ as T_1, \dots, T_n in increasing order of size and consider the smallest element T_i
 270 such that T_i contributes at least r features to $A(\{T_i, \dots, T_n\})$. Such a T_i is guaranteed to exist, since
 271 otherwise $|A(\Delta)| \leq rn$.

272 Now, T_i intersects all the elements in some set $\Delta' \subseteq \{T_{i+1}, \dots, T_n\}$ with $|\Delta'| = \Omega(r^{1/(d-1)})$.
 273 Shrink each element T' in Δ' so as to obtain an element T'' such that (a) the size of T'' is equal to
 274 the size of T_i , (b) $T'' \subseteq T'$, and (c) T'' intersects T_i . Call the resulting set of shrunken elements Δ'' .
 275 Condition (a) and the packing argument above imply that there is a point $p \in \mathbb{R}^d$ that is contained in
 276 $\Omega(r^{1/(d-1)})$ elements of Δ'' . Condition (b) implies that p is contained in $\Omega(r^{1/(d-1)})$ elements of Δ' and
 277 hence also Δ . Therefore, we conclude, as before, that $r = O(k^{d-1})$, which completes the proof. \square

278 **Remark.** Lemma 5 is somewhat surprising, since the union of n homothets of a regular tetrahedron in,
 279 for example, \mathbb{R}^3 can easily have complexity $\Omega(n^2)$. This fact makes it impossible to apply the “usual”
 280 Clarkson-Shor technique [5] that relates the complexity of the first k levels to that of the boundary of
 281 the union (the 0-level).

282 **Lemma 6.** *Let Δ be a set of n homothets of a regular simplex in \mathbb{R}^d such that no point of \mathbb{R}^d is*
 283 *contained in more than k simplices of Δ . Then the arrangement $A(\Delta)$ of Δ can be computed in*
 284 *$O(n(k^{d-1} + (\log n)^d))$ time.*

285 *Proof.* Computing the arrangement $A(\Delta)$ can be done in the following way. Sort the elements of Δ by
 286 decreasing size and construct $A(\Delta)$ incrementally by inserting the elements one by one. When inserting
 287 an element T , use a data structure (described below) to retrieve the elements of Δ that intersect T and
 288 discard the elements that are smaller than T . The proof of Lemma 5 implies that there will be at most
 289 $O(k)$ such elements. The intersection of the surfaces of these $O(k)$ elements with the surface of T has
 290 size $O(k^{d-1})$ and can be computed in $O(k^{d-1})$ time using $d + 1$ applications of the standard algorithm
 291 for computing an arrangement of hyperplanes in \mathbb{R}^{d-1} [7, 8]. Thus, ignoring the cost of finding the
 292 elements that intersect T , the overall running time of the algorithm is $O(nk^{d-1})$.

293 All that remains is to describe a data structure for retrieving the elements that intersect a given
 294 simplex $T \in \Delta$. In the following we describe a data structure that can be constructed in $O(n(\log n)^d)$
 295 time and can answer queries in $O(x + (\log n)^d)$ time, where x is the size of the output. This data
 296 structure will be constructed once and queried n times. The total size of the outputs over all n queries
 297 will be the $O(|A(\Delta)|) = O(nk^{d-1})$.

298 Refer to Figure 5. Suppose that every element $T \in \Delta$ is a homothet of the regular simplex V
 299 whose vertices are v_1, \dots, v_{d+1} and let n_1, \dots, n_{d+1} be the inner normals of the faces of V where n_i
 300 is the face opposite (not incident on) vertex v_i . For any $T \in \Delta$, let t_i be the image of v_i under the
 301 homothetic transformation that takes V onto T . Observe that, if h is a halfspace of \mathbb{R}^d with inner
 302 normal n_i , then h intersects T if and only if h contains t_i . Furthermore, every simplex $T \in \Delta$ is the
 303 intersection of d halfspaces h_1^T, \dots, h_{d+1}^T where the inner normal of h_i^T is n_i . Therefore, a simplex $T \in \Delta$
 304 intersects a simplex $T' \in \Delta$ if and only if h_i^T contains t'_i for all $i \in \{1, \dots, d + 1\}$.

305 This implies that the elements of Δ can be stored in a $(d + 1)$ -layer range tree [3]. The i th layer
 306 of this tree, for $i \in \{1, \dots, d + 1\}$, stores elements of Δ ordered by the projection of t_i onto n_i . In this
 307 way, the range tree can return the set of all simplices in Δ that intersect a query simplex $T \in \Delta$. The
 308 size of this range tree is $O(n(\log n)^d)$ and it can answer queries in time $O(x + n(\log n)^d)$ where x is the
 309 size of the query result. Since each simplex in Δ is passed as a query to this data structure exactly
 310 once, the total sizes of outputs over all n queries is equal to the number of pairs $T_1, \dots, T_2 \in \Delta$ such
 311 that $T_1 \cap T_2 \neq \emptyset$. But the number of such pairs is certainly a lower bound on $|A(\Delta)|$ so it must be at
 312 most $O(nk^{d-1})$. This completes the proof. \square

313 Lemma 6 can be used as a subroutine in the algorithm of Aronov and Har-Peled [1, Theorem 3.3],
 314 to obtain the following Corollary.

315 **Corollary 1.** *Let Δ be a set of n homothets of a regular simplex in \mathbb{R}^d such that some point $p \in \mathbb{R}^d$ is*
 316 *contained in δ elements of Δ . Then there exists an algorithm whose running time is $O(\epsilon^{-2d}n(\log n)^{d-1} +$
 317 $n(\log n)^d)$ and that, with high probability, returns a point $p' \in \mathbb{R}^d$ contained in at least $(1 - \epsilon)\delta$ elements
 318 of Δ .*




Figure 5: The simplex $T \in \Delta$ intersects $T' \in \Delta$ if and only if h^{T_i} contains t'_i for all $i \in \{1, \dots, d\}$.

319 As before, an approximate solution to $\text{PRODUCTDESIGN}(d)$ problem reduces to finding deepest
 320 point in each of the sets $\Delta_1, \dots, \Delta_\ell$ where Δ_i is a set of n d -simplices in H_i . By using the algorithm
 321 of Corollary 1 to do this we obtain the following result:

322 **Theorem 4.** *For any $\epsilon > 0$, there exists an $O(\epsilon^{-(2d+1)}n(\log n)^d + n(\log n)^{d+1})$ time high-probability*
 323 *Monte-Carlo $(1 - \epsilon)$ -approximation algorithm for $\text{PRODUCTDESIGN}(d)$.*

324 5 Conclusions

325 We have given an $O(n \log n)$ time exact algorithm for solving $\text{PRODUCTDESIGN}(1)$ and $O(n(\log n)^{d+1})$
 326 time approximation algorithms for solving $\text{PRODUCTDESIGN}(d)$. The running time of the exact
 327 $\text{PRODUCTDESIGN}(1)$ algorithm is optimal and no algorithm that produces a $(2 - \epsilon)$ -approximation,
 328 for any $\epsilon > 0$, can run in $o(n \log n)$ time.

329 In developing these algorithms, we gave a proof (the proof of Lemma 5) that shows that an
 330 arrangement of n fat convex objects in \mathbb{R}^d has complexity $O(nk^{d-1})$ where k is the maximum number
 331 of objects that contain any given point. We expect that this result, and the algorithm for approximate
 332 depth that arise from it [1], will find other applications.

333 An exact near-linear time algorithm for the case $d = 2$ seems to be out of reach. It appears as
 334 if this problem requires (at least) a solution to the problem of finding a point contained in the largest
 335 number of homothets of an equilateral triangle, a problem for which no subquadratic time algorithm
 336 is known. Is it possible to prove some kind of a lower bound? The related problem of finding the
 337 point contained in the largest number of unit disks is 3-SUM hard [1] providing some evidence that this
 338 problem will be difficult to solve in subquadratic time.

339 In this paper we considered the case where the problem is parameterized by the number, d ,
 340 of orthogonal qualities that a product may have. Another case to consider is the case in which a
 341 manufacturer wishes to introduce some number, k , $k > 1$, of new products into a market. Is this

342 problem NP-hard? Does it have a polynomial time approximation algorithm?

343 Acknowledgement

344 The authors would like to thank Gautam Das for bringing this class of problems to our attention and
345 Timothy Chan for helpful discussions on the subject of approximate depth.

346 References

- 347 [1] B. Aronov and S. Har-Peled. On approximating the depth and related problems. *SIAM Journal*
348 *on Computing*, 38(3):899–921, 2008.
- 349 [2] V. Balachandran and D. H. Gensch. Solving the “marketing mix” problem using geometric pro-
350 gramming. *Management Science*, 21(2):160–171, 1974.
- 351 [3] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications*
352 *of the ACM*, 18:509–517, 1975.
- 353 [4] J. L. Bentley and T. A. Ottman. Algorithms for reporting and counting geometric intersections.
354 *IEEE Transactions on Computing*, C-28:643–647, 1979.
- 355 [5] K. Clarkson and P.W. Shor. Applications of random sampling in computational geometry, II.
356 *Discrete & Computational Geometry*, 4:387–421, 1989.
- 357 [6] K. R. Deal. Optimizing advertising expenditures in a dynamic duopoly. *Operations Research*,
358 27(3):682–692, 1979.
- 359 [7] H. Edelsbrunner, J. O’Rourke, and R. Seidel. Constructing arrangements of lines and hyperplanes
360 with applications. *SIAM Journal on Computing*, 15:341–363, 1986.
- 361 [8] H. Edelsbrunner, R. Seidel, and M. Sharir. On the Zone Theorem for hyperplane arrangements.
362 *SIAM Journal on Computing*, 22(2):418–429, 1993.
- 363 [9] G. M. Erickson. *Dynamic Models of Advertising Competition: Open- and Closed-Loop Extensions*.
364 Kluwer Academic Publishers, Boston, MA, 1991.
- 365 [10] P. Kotler and K. Lane. *Marketing Management*. Prentice-Hall, 2005.
- 366 [11] P. A. Naik, K. Raman, and R. S. Winer. Planning marketing-mix strategies in the presence of
367 interaction effects. *Marketing Science*, 24(1):25–34, 2005.
- 368 [12] G. L. Thompson and J.-T. Teng. Optimal pricing and advertising policies for new product oligopoly
369 models. *Marketing Science*, 3(1):148–168, 1984.
- 370 [13] H. R. Varian. Revealed preference. In Michael Szenberg, Lall Ramrattan, and Aaron A. Gottesman,
371 editors, *Samuelsonian Economics in the 21st Century*, chapter 6, pages 99–115. Oxford University
372 Press, New York, 2006.
- 373 [14] A. C. Yao. Lower bounds for algebraic computation trees with integer inputs. *SIAM Journal on*
374 *Computing*, 20(4):655–668, 1991.