Negotiating over Bundles and Prices Using Aggregate Knowledge: Extended Abstract

D.J.A. Somefun[1], T.B. Klos[1], and J.A. La Poutré[1,2]

[1]*CWI, Kruislaan 413, NL-1098 SJ Amsterdam, The Netherlands*
[2]*Faculty of Technology Management, University of Eindhoven, The Netherlands,*
{koye,tomas,hlp}@cwi.nl

The full version of this paper appeared in:
Proceedings 5^{th} *Int'l. Conf. on E-Commerce and Web Technologies* (EC-WEB)
Lecture Notes in Computer Science, Berlin: Springer-Verlag
Preprint: http://db.cwi.nl/rapporten/abstract.php?abstractnr=1514

Combining two or more items and selling them as one good, a practice called bundling, can be a very effective strategy for reducing the costs of producing, marketing, and selling products. In addition, and maybe more importantly, bundling can stimulate demand for (other) goods or services, although this requires knowledge of customer preferences. Traditionally, firms first acquire such aggregate knowledge, for example through market research or sales data, and then use it for determining which bundle-price combinations they should offer. Especially for online shops, an appealing alternative approach would be to *negotiate* bundle-price combinations with customers: in that case, aggregate knowledge can be used to facilitate an *interactive search* for the desired bundle and price.

In this paper, we present an approach that allows a shop to use aggregate knowledge about *many* customers, in negotiations with *individual* customers. These negotiations concern selecting a subset from a collection of goods or services, viz. the bundle, together with a price for that bundle. Our method finds promising alternatives to the bundle currently under negotiation, so that the shop can recommend an alternative bundle when the current negotiation stalls. Specifically, when the shop deems the progress in the customer's consecutive offers too slow to expect that a deal about current bundle b is imminent, he selects from the neighborhood of b the most promising alternative bundle b', given his assessment of which bundle the customer is most interested in (the customer's 'interest bundle'), and recommends negotiating about bundle b'. Depending on the customer's response to this recommendation, the shop either (a) updates his assessment of the customer's interest bundle when the customer's response exceeds a certain threshold τ, or (b) recommends another alternative bundle from the neighborhood of b when the customer's response implies that b' is even worse than b was, or (c) simply continues negotiating about bundle b' in intermediate cases.

In our approach, aggregate knowledge is used for estimating customers' preferences for bundles, for the purpose of assessing which bundles are promising—in the current paper, aggregate knowledge means that the shop has access to the distribution underlying his customers' preferences.[1] In order to determine the performance of our proposed mechanism, we tested it against many simulated customers using a variety of different negotiation heuritics, and equipped with preferences drawn from various distinct random distributions. Results are presented in Figure 1, which distinguishes our system ('s') from a benchmark ('b') that recommends *random* rather than 'promising' bundles from the current bundle b's neighborhood. The three columns represent the customers using different negotiation strategies.

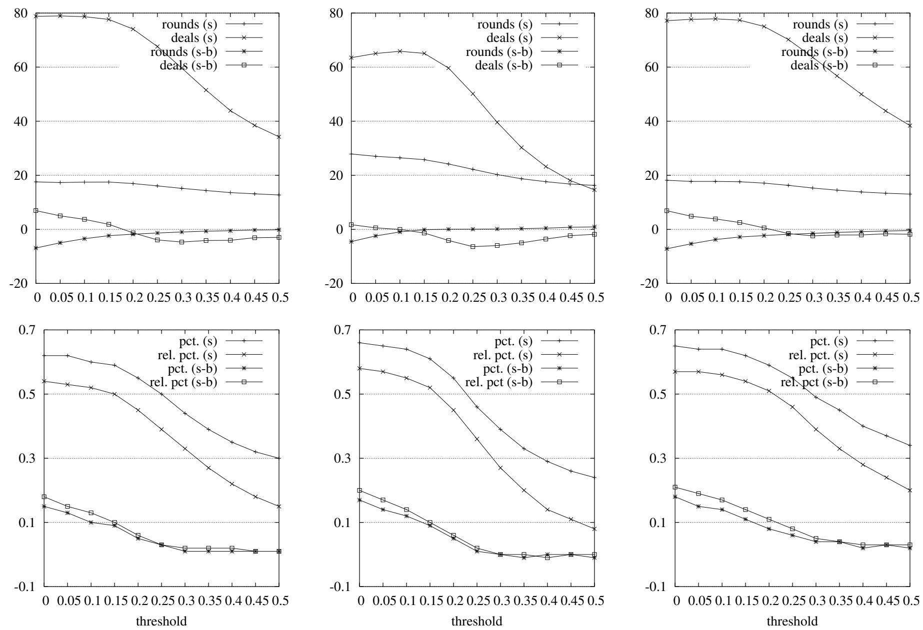

Figure 1: Various results when the customers use different negotiation strategies. The difference with the benchmark is indicated by the '(s-b)'-graphs.

The top row of graphs in Figure 1 shows how, irrespective of the negotiation strategy used by the customers, and for a range of values for the shop's threshold τ, our system 's' reaches more deals than the benchmark 'b,' and uses less rounds to reach those deals. The bottom row of graphs presents the quality of the deals reached as a percentage of the difference between the maximum and the minimum quality attainable ('pct.'), and as a percentage of the difference between the maximum quality and the quality of the first bundle negotiated about ('rel. pct.'). These graphs also witness our system's superior performance as compared to the benchmark, over the whole range of threshold values; the threshold keeps the system from adjusting its estimation of the customer's 'interest bundle' too easily, and from unchecked exploration without exploitation.

[1] In a different paper (Proceedings 6^{th} International Conference on Electronic Commerce, ACM Press), we let the shop learn the required aggregate knowledge online, continuously testing and adapting his estimations over the course of interactions with many customers.