A Decision Procedure for Bundle Purchasing with Incomplete Information on Future Prices

Scott Buffett and Bruce Spencer

Institute for Information Technology - e-Business, National Research Council

46 Dineen Drive, Fredericton, New Brunswick, Canada, E3B 9W4

{scott.buffett, bruce.spencer}@nrc.gc.ca

Biographical Sketches

Scott Buffett is a research officer in the Internet Logic group with the National Research Council in the Institute for Information Technology - e-Business, and a Ph.D candidate in the Faculty of Computer Science at the University of New Brunswick. His interests involve the application of artificial intelligence and economics research to e-commerce related topics such as e-procurement, auctions, privacy and automated negotiation.

Bruce Spencer is the group leader of the Internet Logic group with the National Research Council in the Institute for Information Technology - e-Business, and an adjunct professor of Computer Science with the Faculty of Computer Science at the University of New Brunswick, with eleven years experience in that Faculty. He received his Ph.D from the University of Waterloo in 1990. His research interests center around automated reasoning, applying reasoning systems to the Web, artificial intelligence, logic programming and software engineering.

A Decision Procedure for Bundle Purchasing with Incomplete Information on Future Prices

Abstract. This paper introduces the online bundle purchasing problem (OBPP) as a new computational challenge induced by e-commerce technology. The task of the OBPP is to decide which of possibly many satisfactory combinations (bundles) of items should be purchased, from whom, and when, to maximize the buyer's overall satisfaction. Satisfaction, formalized as multiattribute utility, includes attitudes toward quality, reputation and risk. The Prequote-Quote-Rescind (PQR) protocol communicates probabilistic and temporal information on items' future prices and availabilities. A comparison set, defined as a set of bundles in which all items will be available for a fixed interval and their prices known, determines future intervals when purchase decisions will be fully informed. A decision procedure is provided that makes effective use of comparison sets, and provably improves a buyer's expected utility compared with a naïve decision procedure.

Keywords: e-commerce, utility theory, bundle procurement, protocols, comparison shopping, purchasing decision procedure

1 Introduction

With the high volume of purchasing options available on the Internet today, strategic tools that can cleverly ascertain the true value of several different options are becoming extremely important. More and more businesses are turning to Web-based pricing tools that sift through large volumes of data on product revenues, inventory levels and consumer activity to determine how much to charge for certain items during certain periods of time [10]. Still others are resorting to dynamic pricing [9], where prices can change over time, as well as across consumers or across packages of goods and/or services. Dynamic pricing is utilized not only to maximize profit by responding to changes in supply/demand, but also to discourage the use of price comparing "shopbots" by rendering them unreliable. These pricing strategies translate into higher profits for business, mostly at the expense of the consumer. To combat this trend, buyers need the decision analysis technology that can properly assess not only the current purchasing options, but also the positive or negative potential of future opportunities. Jacobson and Obermiller [8] support this by demonstrating the importance of considering expected future price in consumer decision-making. When making a purchase decision, consumers choose from various alternatives, where one such alternative is delaying the purchase altogether. The utility of this alternative is therefore dependent on the expectation of future prices. The problem of assessing this utility becomes more difficult when we consider expected future prices over several alternatives. In this paper, we focus on the task of computing utility over several options, while considering the buyer's preference over the options as well as the incomplete information on prices.

Consider a potential buyer in need of one of possibly many acceptable *bundles* of items. In our context, we consider a bundle to be a set of items as determined by the buyer. Acquiring all items in any particular bundle will be considered a success by the buyer (with varying degrees). The buyer must assess his/her preferences for the attributes (e.g. total cost, item preference, item suppliers, compatibility of the items, etc.) of all bundles for which items are currently available, as well as those for bundles containing items that will be available in the future. For these future bundles, the buyer might have only a probability measure on the outcome of the cost. At certain times, the buyer must decide whether to buy a bundle of currently available items, or to take the risk of letting these pass and waiting for future opportunities. The goal is to buy the single bundle that best meets the buyer's preferences. We refer to this as the *online bundle purchasing problem* (OBPP). A common approach to deciding whether or not to purchase a given bundle of items is to compare the utility of purchasing the bundle with the expected utility of all other bundle purchases. The bundle in question is then purchased if and only if the utility of purchasing it meets or exceeds the expected utility of any other possible bundle purchase. If the purchasing model requires the buyer to choose a bundle and commit to it immediately, then this is an optimal strategy. However, if the task is to simply decide whether or not to commit to the current bundle, then the expected utility of not committing is almost certainly higher (and never lower) than the expected utility of any future bundle purchase.

In this paper, we propose an approach that can be used to more accurately assess the expected utility of not purchasing a bundle. This assessment is based on the expected utility of future *time intervals* during which the buyer will have the option of choosing amongst possibly several bundle purchases for which all prices are known. We consider the setting where at any given time, the purchaser has price quotes for some items that are available for some fixed period of time, as well as the knowledge of incoming quotes, availabilities, sales, price fixing, etc., for other items available during definite future periods. The purchaser might have only a probability measure on these future prices. Purchasing decisions therefore have to be made online with incomplete information. The paper is organized as follows: After providing a brief review of the literature describing background and related work in section 2, section 3 formally defines the OBPP and proposes the Prequote-Quote-Rescind (PQR) protocol for information exchange between a supplier and a purchaser for probabilistic and temporal information. Section 4 then formally describes a naïve decision procedure that always pursues the best bundle in terms of expected utility, while section 5 proposes a decision procedure that pursues the best *time interval.* This is done by grouping bundles with common purchase intervals into *comparison sets*. Section 6 shows that using this improved decision procedure always yields expected utility at least high as the naïve procedure, and derives an expression measuring the improvement for a simple case. Finally section 7 offers conclusions and plans for future work.

2 Literature Review

We make use of *expected utility theory* (EUT) [11, 16] to facilitate the decision process. Basing purchasing decisions on expected utility maximization (as opposed to expected cost minimization, for example) seems appropriate since one may wish to be sensitive to the decision-maker's attitude toward risk, as well as preferences for attributes such as item quality, compatibility with other items, and supplier reliability.

The online market clearing [1] problem parallels our online bundle purchasing problem in that transactions are made in a continuous setting where transaction possibilities may pass by and information about future possibilities may be incomplete or unknown. The difference is, in online market clearing buy bids and sell bids are matched to maximize some value such as seller profit or market liquidity. Online bundle purchasing is concerned with a different problem. Here the task is to match a buyer's needs with items offered by suppliers so that a single, complete bundle is purchased. We are not concerned with excess supply or market liquidity, but only with the satisfaction of the buyer.

Recent work has focused on decision procedures for bundle purchasing where there are multiple auctions in which to bid. Boutilier *et al.* [2, 3] consider the model where a bundle of items[1] must be purchased by participating in a subset of several *sequential* auctions. These auctions are first-price sealed-bid, have known start/end times, and do not overlap. At each decision point (auction start

[1] Items are referred to as "resources".

time), the optimal bidding strategy is computed and the amount to bid (if any) in the current auction is determined. Our work differs from this in both the auction mechanism used as well as in the timings, as we use the request-for-quote mechanism and allow quotes to be open in parallel. Byde [4] considers multiple simultaneous auctions, but the purchaser's goal in this case is to buy only one single item. The problem where there are multiple simultaneous auctions has been examined by Byde *et al.* [5]. In their model, the purchaser attempts to buy possibly multiple units of only a single good. Finally, Preist *et al.* [14] discuss bundle purchasing[2] in the setting where there are multiple simultaneous auctions. While their problem is more daunting than ours since they consider English, Dutch and sealed-bid auctions, their decision-making method only pursues the set of auctions that maximizes expected utility. That is, the expected utility of not participating in the current auctions is judged to be the expected utility of the optimal future set. Since the algorithm does not truly commit to this set, but instead re-evaluates its options at each decision point, this expected utility is not an accurate account of the true expected utility of the choice. The main idea of our paper is to predict how the algorithm will behave in the future in order to more accurately estimate the true expected utility of a choice.

[2] Bundle purchasing is referred to as "service composition"

3 The Online Bundle Purchasing Problem

3.1 Problem Formalization

Let I be a set of items and $\mathcal{B} \subseteq 2^I$ be a set of *bundles*, where each $b \in \mathcal{B}$ is a combination of items that meets the needs of the buyer. At any given time, let I contain only those items that are known to be available either currently or during some definite future time period. I and $\mathcal{B}$ may therefore change over time as new availabilities arise and others pass by. Each $i \in I$ has a quoted cost $c(i)$ and each $b \in \mathcal{B}$ has cost $c(b)$ equal to the sum of its item costs. Note that when two instances of the same item are offered by two different suppliers, or by the same supplier but as part of two different offers, they are treated as two different items. If an item i is currently available, then assume the buyer knows $c(i)$. Otherwise, the buyer has a probability measure $p : Z \rightarrow \Re$ on the outcome of the cost of i, where Z is the set of monetary units. This could be any discrete or continuous distribution obtained from market history, from the supplier directly, a third party, or even subjectively decided upon by the buyer. The goal in the OBPP is to make decisions that maximize expected utility, ultimately giving the buyer the greatest chance of purchasing the $b \in \mathcal{B}$ that is most preferable in terms of b and $c(b)$. In our context, we denote the utility of a bundle purchase by the function $u : \mathcal{B} \times Z \rightarrow \Re$, and the expected utility of a bundle purchase by $E[\tilde{u}(b)]$ where for each $b \in \mathcal{B}$, $\tilde{u}(b)$ is a random variable for the unknown outcome of the utility of purchasing b, based on the probability measure for the outcome of the cost of b.

3.2 The PQR Protocol

The Prequote-Quote-Rescind (PQR) protocol is a message-passing protocol for information exchange between a supplier and a purchaser for probabilistic and temporal information. It defines when information will become known by the purchaser about items such as cost, the distribution of possible outcomes on cost, the time a quote will be offered, and the time a quote will be terminated. This information can then be used when planning purchases. Let $[t_0, t_n] \subseteq \Re$ be the period of time during which a buyer needs to purchase some bundle b of items I, and let $t_p : I \rightarrow \Re$, $t_q : I \rightarrow \Re$ and $t_r : I \rightarrow \Re$ assign time points to items $i \in I$, where $t_p(i)$ is the *prequote* time, $t_q(i)$ is the *quote* time and $t_r(i)$ is the *rescind* time for i, and $t_p(i) \leq t_q(i) < t_r(i)$. The intervals $[t_p(i), t_q(i)]$ and $[t_q(i), t_r(i)]$ are known as the *prequote interval* and the *quote interval* for i, respectively. The quote time is the time at which the quote will be offered, the rescind time is the time at which the quote expires, and the prequote time is the time at which the buyer learns the quote and rescind times. It is assumed that at the prequote time the buyer also learns or determines the probability measure on the cost outcome of the item. Note that at time t, $\forall i \in I$, $t_p(i) \leq t < t_r(i)$. So an item is added to I when the prequote is received, and is removed at the rescind time. Table 1 summarizes the time periods during which the buyer will have information on the cost, potential cost, and availability of an item.

Table 1. Summary of time periods during which the buyer will have certain information about i

4 A Naïve Decision Procedure

This section formalizes a naïve decision procedure for the OBPP. The strategy presented is simple: Each time an item in a bundle b is about to expire, $u(b, c(b))$ is computed, as well as the expected utility of all other valid bundle purchases. If $u(b, c(b))$ is higher than all expected utilities, then buy b. Else, let it expire. Before the formal definition of the decision procedure is given, the following terms are defined:

Definition 1 For a bundle b, $pi(b) = [t_q(b),\ t_r(b)] = \bigcap_{i \in b}[t_q(i), t_r(i)]$ is known as the *purchase interval* of b. All items in a bundle can be purchased at any time during its purchase interval.

Definition 2 A bundle b is *valid* iff $pi(b) \neq \phi$.

Let t be the current time, let I and $\mathcal{B}$ be as defined in section 3.1, and let $t_p(i)$, $t_q(i)$ and $t_r(i)$ assign time points to each item $i \in I$ as defined in section 3.2. Also, let $\mathcal{B}_v \subseteq \mathcal{B}$ be the set of valid bundles and let $\{E[\tilde{u}(b')] \mid b' \in \mathcal{B}_v\}$ be the set of expected utilities of valid bundles. If $\exists b \in \mathcal{B}_v$ such that $t = t_r(b) - \varepsilon$, then the purchaser decides whether to buy b using the following rule:

> If $u(b, c(b)) \geq \max\{E[\tilde{u}(b')] \mid b' \in \mathcal{B}_v\}$, then purchase b.
>
> Else, allow b to expire.

This defines the set of decision points to be $\{t_r(b) - \varepsilon \mid b \in \mathcal{B}_v\}$. In theory, any bundle is available for purchase at any time during its purchase interval, but it would be unwise to commit to purchasing it much before $t_r(b)$. First of all, since the cost of the bundle is fixed until $t_r(b)$, and $t_r(b)$ is known by the purchaser, there is no need to commit any earlier. Secondly, since new information on other bundles may arise, it would be best to wait until the last moment (perhaps leaving a minimal amount of time ϵ before $t_r(b)$ to perform the transactions or inform the suppliers of the buyer's intentions). Therefore, decisions only need to be made at (or just before) these $t_r(b)$ time points. At such a time, the utility of purchasing b is compared with the utilities and expected utilities of other available and future prospects, and a decision on whether or not to buy b is made.

This method is referred to as the naïve decision procedure since it merely pursues the bundle with the greatest expected utility, without using any strategy or taking any other factors into account. A more intelligent method that takes into account the impact of possible future options is now given.

5 An Improved Decision Procedure

In the naïve decision procedure, each time a valid bundle purchase is about to expire, it must be determined whether or not there is another bundle purchase that is likely to be better. Instead of determining whether or not *there is a future purchase that is likely to be better*, it should be determined whether or not *it is likely that a future purchase will be better.* These are two different questions, as explained with a simple example:

Example 1. Consider playing a game of chance with a typical six-sided fair die. You roll it once and get a 4. You are then given a decision to make: either you end the game and take \$4, or you can choose to give up the \$4 and roll the die twice more, winning the equivalent dollar amount of the higher of your two rolls. While the expected value of each roll is just 3.5, the *expected higher value* is

$$E(\tilde{x}_h) = \sum_{k=1}^{6} kP(\tilde{x}_h = k) = 4.47 \tag{1}$$

where $\tilde{x}_h$ is the uncertain higher outcome of the two rolls. So the decision-maker would expect on average to make \$4.47 if he/she chooses to continue.

The same idea comes up when making decisions about purchases. If one is trying to choose between making a purchase now and waiting until later, and it is known that there is a future time period where two or more bundles will be offered, the buyer needs to compare the utility of the current bundle with the *expected highest utility* of those future bundles, since the buyer will have the luxury of comparing them at that time and choosing the one with the highest utility. The notion of a *comparison set*, which is a set of bundles for which there is a period of time that the buyer will have complete information, is now introduced.

5.1 Comparison Sets

Recall that the purchase interval $pi(b)$ for a bundle b is the period of time during which the prices of all items in b are known, and all items are available for purchase.

Definition 3 Let $\mathcal{B}_v$ be a set of valid bundles and let $CS \subseteq \mathcal{B}_v$. CS is a *comparison set* of $\mathcal{B}_v$ iff it is maximal such that $ci(CS) = \bigcap_{b \in CS} pi(b)$ is non-empty. The interval $ci(CS)$ is called the *comparison interval* of CS. The *comparison set cover* $csc(\mathcal{B}_v)$ of $\mathcal{B}_v$ is the set of all comparison sets of $\mathcal{B}$.

Note that $ci(CS)$ is the period of time during which the prices of all items in all bundles in CS are known, and all items are available for purchase. So the buyer has complete information on all bundles in CS. Note that any bundle in $\mathcal{B}_v$ will appear in at least one comparison set even if by itself, and may also appear in more than one. Thus $csc(\mathcal{B}_v)$ is a covering of $\mathcal{B}_v$.

Algorithm 1 (Construction) The comparison set cover $csc(\mathcal{B}_v)$ for $\mathcal{B}_v$ is constructed by first finding the comparison intervals, and then determining the comparison sets from those. This is done as follows. Let T be a sorted list of the time points in $\{t_q(b) \mid b \in \mathcal{B}_v\} \cup \{t_r(b) \mid b \in \mathcal{B}_v\}$ from earliest to latest. Ties between a t_q and a t_r time are broken be placing the t_r time first, and all other ties are broken arbitrarily. For each pair of consecutive elements t_k and t_{k+1} in T, if t_k is a t_q time and t_{k+1} is a t_r time, then $[t_k, t_{k+1}]$ is a comparison interval, and $CS = \{b \in \mathcal{B}_v \mid [t_k, t_{k+1}] \subseteq pi(b)\}$ is therefore a comparison set. The comparison set cover $csc(\mathcal{B}_v)$ is then the set of all of these comparison sets.

Example 2. Let $\mathcal{B}_v = \{b_1, b_2, b_3, b_4, b_5\}$ where each bundle has a purchase interval as depicted by horizontal lines in Figure 1 (e.g. the purchase interval for b_1 is $[0, 3]$). The comparison intervals are indicated by dotted vertical lines.

The comparison set cover for $\mathcal{B}_v$ is then $csc(\mathcal{B}_v) = \{CS_1, CS_2, CS_3\}$, where $CS_1 = \{b_1, b_2\}$, $CS_2 = \{b_2, b_3, b_4\}$, and $CS_3 = \{b_5\}$.

Fig. 1. Comparison set cover of $\mathcal{B}_v$ in Example 2.

Since all items in all bundles in a given comparison set CS are available during a common interval and all prices are known, if the buyer chooses to buy during this period, he will choose the bundle in CS with the highest purchase utility. The utility one would expect to achieve during this period is therefore equal to the expected highest utility of the bundles in CS, referred to hereafter simply as the expected utility of CS and denoted by $E[\tilde{u}(CS)]$.

5.2 The Proposed Decision Procedure

Let b be the bundle currently available at cost $c(b)$ that is about to expire. Also let $\mathcal{B}_v$ be the set of valid bundles and let $csc(\mathcal{B}_v)$ be the comparison set cover for $\mathcal{B}_v$, each $CS \in csc(\mathcal{B}_v)$ with expected utility $E[\tilde{u}(CS)]$.

> If $u(b, c(b)) \geq \max\{E[\tilde{u}(CS)] \mid CS \in csc(\mathcal{B}_v)\}$, then purchase b.
>
> Else, allow b to expire.

Using this decision procedure is proven to provide the buyer with a higher expected utility than using the naïve decision procedure (see Theorem 1 in section 6.1).

5.3 Calculating the Expected Utility of a Comparison Set

Unfortunately, computing expected utilities of comparison sets can be a quite complex. When using continuous random variables to represent item prices, in order to calculate the exact expected utility of a comparison set, one would have to solve the multiple integral

$$E[\tilde{u}(CS)] = \int_0^1 \dots \int_0^1 \max\{x_1, \dots, x_n\} \prod_{i=1}^{n} p_i(x_i)\; dx_1 \dots dx_n \tag{2}$$

where $x_1, \dots, x_n$ are the utilities of the bundles in CS and $p_1, \dots, p_n$ are their respective probability density functions. Since no closed-form expression exists for even the single integral of a normal probability density function [12], if some or all of the p_i are normal (or some other complex form) then it is unlikely that the above can be expressed in closed form. Making this computation even more difficult is that there may be a strong interdependence between many of the p_i, since many bundles have common items. While it is possible to compute the value of this integral exactly if it has certain properties (e.g. if the p_i are uniformly distributed, or have only a few discrete outcomes), in general we suggest using a *Monte Carlo* method for estimating this value.

Monte Carlo methods [6, 13] involve simulation to approximately solve a mathematical problem. Such a method can be used to estimate the expected highest utility of a set of bundles in a comparison set. This is done by first properly modeling the system of items residing in the bundles in question, which includes the probability distributions for costs of the items as well as possible interdependencies between the item costs, if they exist. The results of several independent simulations of the random elements involved in the system are then obtained. For each simulation, the outcomes of the item prices are used to determine the utility of purchasing each bundle, and the highest is noted. The average of these results is then taken as the unbiased estimate of the expectation. Simulations are run until the standard error $\sigma/\sqrt{n}$ is small enough to achieve the desired confidence in the estimate. To help shrink $\sigma/\sqrt{n}$, a variance reduction technique referred to as the *antithetic variate* sampling method [7] is used. With this method, price outcomes are selected in pairs that mutually compensate for each other's variations. See Hammersley and Handscombe [6] for a more detailed description of the approach.

6 Analysis

This section presents a theorem that proves that using the improved decision procedure yields a higher expected utility than using the naïve decision procedure. A theorem is also given for determining the increase in utility to be expected when using the improved decision procedure over the naïve procedure, for a simple case. We then provide an example in order to shed some light on what this increase means.

6.1 The Improved Decision Procedure Yields Higher Expected Utility

For a comparison set cover $csc(\mathcal{B}_v)$ of valid bundles $\mathcal{B}_v$, let $j = \max\{E[\tilde{u}(b)] \mid b \in \mathcal{B}_v\}$ and $k = \max\{E[\tilde{u}(CS)] \mid CS \in csc(\mathcal{B}_v)\}$.

Lemma 1. $j \leq k$.

Proof. Let CS_j be a comparison set containing b such that $E[\tilde{u}(b)] = j$. Since the expected utility of choosing from a number of bundles must be at least as high as the expected utility of any one of the bundles, $E[\tilde{u}(CS_j)] \geq j$. So $k = \max\{E[\tilde{u}(CS)] \mid CS \in csc(\mathcal{B}_v)\} \geq E[\tilde{u}(CS_j)] \geq j$. □

Lemma 2. The buyer's expected utility Eu_{im} when using the improved decision procedure is greater than or equal to k.

Proof. By induction on the size of $csc(\mathcal{B}_v)$.

Base Case. Let $|csc(\mathcal{B}_v)| = 1$. Then $Eu_{im} = k$.

Induction. Let $csc(\mathcal{B}_v)$ be of arbitrary size, let CS_i be the current comparison set, and let b be the bundle in CS_i available at cost z such that $u(b, z)$ is the maximum in CS_i. Assume that the expected utility when using the informed decision procedure for $csc(\mathcal{B}_v) \setminus CS_i$ is at least $\max\{E[\tilde{u}(CS)] \mid CS \in \ csc(\mathcal{B}_v) \setminus \ CS_i\}$. Consider the two outcomes (note that $u(b, z) > k$ is not possible):

1. $u(b, z) = k$. Buy b and achieve utility k. So $Eu_{im} = k$.
2. $u(b, z) < k$. Then $k = \max\{E[\tilde{u}(CS)] \mid CS \in csc(\mathcal{B}_v) \setminus CS_i\}$. Since b would not be purchased in this case (and thus nothing from CS_i is purchased, so it expires) and by induction the expected utility for $csc(\mathcal{B}_v) \setminus CS_i$ is at least $\max\{E[\tilde{u}(CS)] \mid CS \in csc(\mathcal{B}_v) \setminus CS_i\}$, then $Eu_{im} \geq k$. □

Theorem 1 Let Eu_{im} and Eu_{na} denote the utilities expected when using the improved and naïve decision procedures, respectively. Then $Eu_{im} \geq Eu_{na}$.

Proof. By induction on the size of $csc(\mathcal{B}_v)$.

Base Case. Let $|csc(\mathcal{B}_v)| = 1$. Both procedures will choose the bundle purchase with the highest utility, so $Eu_{im} = Eu_{na}$.

Induction. Let $csc(\mathcal{B}_v)$ be of arbitrary size, let CS_i be the current comparison set, and let b be the bundle in CS_i available at cost z such that $u(b, z)$ is the maximum in CS_i. Assume that the expected utility when using the improved decision procedure is greater than or equal to that when using the naïve decision procedure for $csc(\mathcal{B}_v) \setminus CS_i$. By Lemma 1 there are three cases:

1. $k = u(b, z)$. Both procedures would choose b so $Eu_{im} = Eu_{na}$.
2. $j \leq u(b, z) < k$. The naïve procedure would choose b and the improved would not. So $Eu_{na} = u(b, z) < k$. Since, by Lemma 2 $Eu_{im} \geq k$, $Eu_{im} > Eu_{na}$.
3. $u(b, z) < j$. Neither procedure would choose b, and would thus allow CS_i to pass. By induction, $Eu_{im} \geq Eu_{na}$. □

6.2 Expected Utility Increase for the Improved Decision Procedure

The increase in utility that one would expect to achieve when using the improved procedure over the naïve procedure for a simple case is derived. Here we consider the situation where there are two comparison sets CS_1 and CS_2, there are no interdependencies between the utilities of the two comparison sets (thus $CS_1 \cap CS_2 = \phi$), and $ci(CS_1)$ is before $ci(CS_2)$. Let $j = \max\{E[\tilde{u}(b')] \mid b' \in CS_2\}$ and $k = E[\tilde{u}(CS_2)]$. Let an example probability density function for the unknown outcome of $\tilde{u}(CS_1)$, which is the highest utility over all bundles in CS_1, be as depicted in Figure 2. Example points for j and k are also displayed. Using these points, the area under the curve is divided into three regions. Let A_1, A_2, and A_3 represent the areas of these regions. Specifically,

$$A_1 = \int_0^j p(x)\, dx \qquad A_2 = \int_j^k p(x)\, dx \qquad A_3 = \int_k^1 p(x)\, dx$$

Fig. 2. Probability density function for the utility of purchasing an bundle in CS_1

Theorem 2 The expected increase in utility to be achieved by using the improved decision procedure is

$$\int_j^k (k-x)p(x)\, dx \tag{3}$$

where there are two independent comparison sets CS_1 and CS_2, $p(x)$ is the probability density function for the outcome of the highest utility over all bundles in CS_1, $j = \max\{E[\tilde{u}(b')] \mid b' \in CS_2\}$, and $k = E[\tilde{u}(CS_2)]$.

Proof. Let Eu_{na} be the expected utility of using the naïve decision procedure and Eu_{im} be the expected utility of using the improved procedure. Since the naïve decision procedure will choose to buy from CS_1 if it contains a bundle with utility higher than j, and let CS_1 expire otherwise, thus expecting utility k, then[3]

$$Eu_{na} = A_1 k + A_2 E[\tilde{u}(CS_1)|j < \tilde{u}(CS_1) < k] + A_3 E[\tilde{u}(CS_1)|\tilde{u}(CS_1) > k] \tag{4}$$

[3] Here $E[X|C]$ is the expected value of X given that condition C holds.

Since the improved decision procedure will choose b if its utility is higher than k, and let it expire otherwise, thus expecting utility k, then

$$Eu_{im} = A_1 k + A_2 k + A_3 E[\tilde{u}(CS_1)|\tilde{u}(CS_1) > k] \tag{5}$$

Subtracting gives

$$\begin{aligned} Eu_{im} - Eu_{na} &= A_2 k - A_2 E[\tilde{u}(CS_1)|j < \tilde{u}(CS_1) < k] \\ &= A_2 k - A_2 \cdot \frac{\int_j^k xp(x)\, dx}{A_2} \\ &= k \int_j^k p(x)\, dx - \int_j^k xp(x)\, dx \\ &= \int_j^k (k - x)p(x)\, dx \quad \square \end{aligned} \tag{6}$$

For the general case involving many comparison sets, if $j = \max\{E[\tilde{u}(b')] \mid b' \in csc(\mathcal{B}_v) \setminus CS_1\}$ and $k = \max\{E[\tilde{u}(CS)] \mid CS \in csc(\mathcal{B}_v) \setminus CS_1\}$, then this value is a lower bound on the expected increase, with the restriction that comparison set utilities are independent. If there exist some interdependencies, then a lower bound can still be determined if the interdependencies are removed such that j is unaltered. For example, if there exists an item i that resides in bundles in two or more comparison sets, then restrict the quote interval of i to $ci(CS)$ where CS is a comparison set that contains a bundle b such that $i \in b$ and $\tilde{u}(b) = \max\{\tilde{u}(b') \mid i \in b'\}$. Since j will be still be the true value but k could be an underestimate, then this value will be a lower bound.

6.3 An Example

Let $\{CS_1, CS_2\}$ be a comparison set cover with non-intersecting comparison intervals and let $CS_1 = \{b_1\}$, and $CS_2 = \{b_2, b_3\}$. For simplicity, let the bundle costs be independent and normally distributed, and consider the buyer to be risk neutral. Then the bundle purchase utilities will be normally distributed. Consider the parameters given in Table 2.

Table 2. Means and variances of bundle purchase utilities in the example

At time $t_r(b_1)$, the buyer will know $c(b_1)$ (and therefore $u(b_1, c(b_1))$), μ_2, σ_2, μ_3, and σ_3. A decision must be made at that time between purchasing b_1, which will achieve utility $u(b_1, c(b_1))$, and allowing b_1 to expire, which will achieve the higher of the two utility outcomes for b_2 and b_3.

Recall that the naïve decision procedure chooses b_1 iff $u(b_1, c(b_1)) \geq j$, where $j = \max\{E[\tilde{u}(b)] \mid b \in CS_2\} = .484$, and the improved procedure chooses b_1 iff $u(b_1, c(b_1)) \geq k$, where $k = E[\tilde{u}(CS_2)]$. Using MC simulation, we find that $k = .545$. Then by Theorem 2, the expected increase in utility is

$$\int_{.484}^{.545} (.545 - x)p(x)\, dx \approx .012 \tag{7}$$

where $p(x)$ is the normal probability density function with $\mu = .5$ and $\sigma = .06$.

This means that we expect to achieve .012 more utility by using the improved decision procedure. It is difficult to make any conclusions as to the significance of this increase without knowing the utility function, since it is the result of some combination of more highly preferred bundles and lower costs. But, for the sake of a simple example, assume that all bundles are preferred equally. If this is the case, then lower costs will be completely responsible for the increase in utility. Also, assume that the range of possible outcomes of bundle costs is

[$\$100, \$200$] and the range of utility is normalized to [0,1] (which is commonly done). Therefore $u(b, \$100) = 1$ and $u(b, \$200) = 0$. If the buyer is risk neutral then each .01 of bundle utility represents \$1, and therefore an increase in utility of .012 represents a savings of \$1.20. This is a significant result considering that this is such a small-scale example. Consider a more large-scale example, such as purchasing materials for a major construction project, where the range of values spans one million dollars instead of one hundred. In this case, the increase in utility represents \$12,000 in savings.

7 Conclusions and Future Work

This paper formalizes the online bundle purchasing problem (OBPP) as the problem faced by a buyer in need of purchasing one of possibly several bundles, in the setting where at any given time, the purchaser has price quotes for some items that are available for some fixed period of time, as well as the knowledge of incoming quotes, availabilities, sales, price fixing, etc., for other items available during definite future periods. Expected utility is used as the maximization goal. Initially, the Prequote-Quote-Rescind (PQR) protocol is introduced as a set of message passing rules that provide a framework defining the time intervals during which certain information is known to the purchaser regarding item availabilities and prices. A decision procedure that exploits time intervals during which many options will be available is then proposed for the OBPP, and is proven to yield a higher expected utility than a naïve decision procedure that simply pursues the best bundle. We substantiate this with an analysis of the value of considering future choices by deriving a measure of the improvement that a purchaser would expect to realize if future choices are considered when making decisions, compared with simply pursuing best bundle purchase, for a simple case.

If your market is completely known and invariable, then a naïve inflexible purchase decision procedure is appropriate. Our PQR protocol accommodates dynamic and volatile market conditions, and we provide a robust, flexible purchase decision procedure that deals with the uncertainty and exploits future options.

In current on-line shopping, the purchaser is offered quotes immediately upon asking for them, but we feel that this forces sellers to set high prices to make up for not knowing what the demand for their product will be. Our proposed PQR protocol allows sellers to gather quote requests, and gives sellers some idea of demand before they set a price, perhaps allowing them to set a more fair or less risky value. In B2B e-commerce, price setting in response to quote requests is already common. In future work, we will look at the effect of the PQR protocol on vendors, as well as the effects of protocols other than the PQR on both vendors and purchasers. This will likely involve the use of some game-theoretic techniques such as those described by Shubik [15].

In another project, we plan to explore the computational ramifications of relaxing the restriction that all items in a bundle must be available at the same time. It is very possible that the best solution consists of buying items in a bundle at different times. Making a partial bundle purchase is however quite risky, since the buyer may be forced to purchase expensive items to complete the bundle if cost outcomes turn out to be high. Expected utility theory is used again here to determine whether the added utility of making a partial purchase outweighs the risk. This is a very complex mathematical problem, since we not only consider the expected outcomes of these future items needed to complete bundles, but also what choices are to be made in the future and what new information will be known at those times. All of these facts affect the expected utility of a decision, and must be considered. Straight simulation of the outcomes therefore does not give an accurate result. We are currently researching both randomized and non-randomized algorithms that solve this new problem.

We also plan to extend the model to allow the buyer to participate in on-line auctions. While the addition of various auction mechanisms would greatly magnify the computational burden, it would certainly make the methods described much more useful. New techniques, based on those developed in this paper, would likely be needed to accomplish this goal.

References

1. Blum, A.; Sandholm, T.; and Zinkevich,M. Online algorithms for market clearing. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 971–980, 2002.

2. Boutilier, C.; Goldszmidt, M.; and Sabata, B. Continuous value function approximation for sequential bidding policies. In *the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-99)*, pp. 81–90, Stockholm, 1999.

3. Boutilier, C.; Goldszmidt, M.; and Sabata, B. Sequential auctions for the allocation of resources with complementaries. In *the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*, pp. 527–534, Stockholm, 1999.
4. Byde, A. A dynamic programming model for algorithm design in simultaneous auctions. In *WELCOM'01*, pp. 152–163, Heidelburg, Germany, 2001.
5. Byde, A.; Preist, C.; and Jennings, N.R. Decision procedures for multiple auctions. In *Proc. 1st Int Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS'02)*, pp. 613–620, Bologna, Italy, 2002.
6. Hammersley, J.M.; and Handscombe, D.C. *Monte Carlo Methods*. New York: Wiley, 1964.
7. Hammersley, J.M.; and Morton, K.W. A new Monte Carlo technique: antithetic variates. In *Cambridge Phil. Soc.*, volume 52 (1956), pp. 449–475.
8. Jacobson, R.; and Obermiller, C. The formation of expected future price: a reference price for forward-looking consumers. *Journal of Consumer Research*, 16 (March 1990), 420–432.
9. Kannan, P.K.; and Kopalle, P.K. Dynamic pricing on the internet: Importance and implications for consumer behaviour. *International Journal of Electronic Commerce*, 5, 3 (Spring 2001), 63–83.
10. Keenan, F. The price is really right. *BusinessWeek Online*, url="http://www.businessweek.com/magazine/content/03_13/b3826046.htm", Accessed on March 31, 2003.
11. Keeney, R.L; and Raiffa, H. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: Wiley, 1976.
12. Mendenhall, W.; and Scheaffer, R.L. *Mathematical Statistics with Applications*. Boston: Duxbury Press, 1973.

13. Metropolis, N.; and Ulam, S. The Monte Carlo method. *Journal of the American Statistical Association*, 44, 247 (1949), 335–341.

14. Preist, C.; Bartolini, C.; and Byde, A. Agent-based service composition through simultaneous negotiation in forward and reverse auctions. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, San Diego, pp. 55–63, 2003.

15. Shubik, M. *A Game-Theoretic Approach to Political Economy.* Massachusetts: MIT Press, 1984.

16. von Neumann, J.; and Morgenstern, O. *Theory of games and economic behaviour, 2nd ed.* Princeton: Princeton University Press, 1947.

Figures and Tables

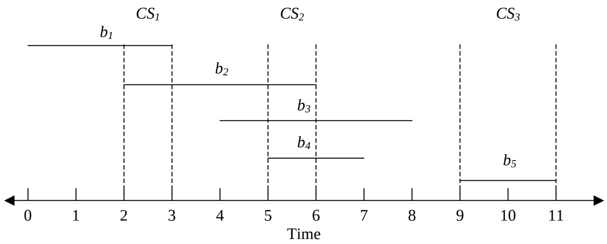

Fig. 1. Comparison set cover of $\mathcal{B}_v$ in Example 2

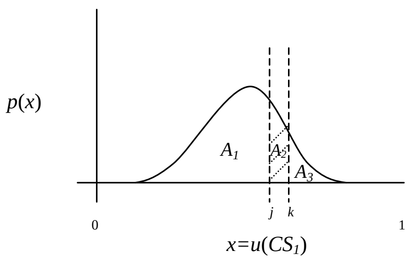

Fig. 2. Probability density function for the utility of purchasing a bundle in CS_1

Interval	Information
$[t_0, t_p(i)]$	nothing is known about i
$[t_p(i), t_n]$	$t_q(i)$ is known; $t_r(i)$ is known; a probability measure on $c(i)$ (and perhaps $c(i)$ itself) is known
$[t_q(i), t_r(i)]$	i is available for purchase
$[t_q(i), t_n]$	the actual price of i is known
$[t_r(i), t_n]$	i is subject to unavailability or price change

Table 1. Summary of time periods during which the buyer will have certain information about i

b_i	μ_i	σ_i
b_1	.5	.06
b_2	.475	.13
b_3	.484	.1

Table 2. Means and variances of bundle purchase utilities in the example