

# Timing When to Buy\*

Rakesh Agrawal  
Search Labs  
Microsoft Research  
rakesha@microsoft.com

Samuel Jeong  
Search Labs  
Microsoft Research  
saieong@microsoft.com

Raja Velu<sup>†</sup>  
School of Management  
Syracuse University  
rpvelu@syr.edu

August 2011

## Abstract

Most e-commerce sites to-date have focused on helping consumers decide what to buy and where to buy. We study the complementary question of helping consumers decide *when* to buy, focusing on consumer durables. We introduce a utility-based model for evaluating different approaches to this question. We focus on how best to make use of forecasts in making recommendations, and propose three natural strategies. We establish a relationship between these strategies, and show that one of them is optimal. We conduct a large-scale experimental study to test the performance and robustness of these strategies. Across a wide range of conditions, the best strategy obtains 90% of the maximum possible gains.

## Categories and Subject Descriptors

J.4[**Social and Behavioral Sciences**]: Economics;

F.2.0[**Analysis of Algorithms and Problem Complexity**]: General

## General Terms

Economics, Algorithms, Experimentation

## Keywords

Recommendation, Automated Shopping

## 1 Introduction

In recent times, there has been tremendous growth in the amount of commerce conducted over the Web. The yearly retail e-commerce sales in U.S. alone are estimated to have topped \$100 Billion. Nearly seven out of ten consumers say that the Internet has become important in providing them with information to help them make buying decisions and that they are likely to shop online before making an offline purchase [10]. Thus the Internet has started to play a crucial role in the decision making process of the consumers irrespective of whether they buy online or offline.

---

\*Microsoft Research Technical Report: MSR-TR-2011-97.

<sup>†</sup>Work done while author was visiting Search Labs.

Various e-commerce sites exist to help consumers make purchasing decisions. Some websites provide product reviews to help consumers decide what to buy among similar products. Others compare prices across different retailers to help consumers decide where to buy the product. However, few sites exist to help consumers decide *when to buy*.

In this paper, we study the design of decision aids that help consumers decide the timing of the purchase of durable products such as household appliances and electronic goods. Such decision aids are important as consumer durables are typically expensive, and consumers experience remorse when they find out they could have bought the products at lower prices had they waited a little longer [5].

Certain websites advice consumers about potential price changes in the domain of airfare ([bing.com/travel](http://bing.com/travel), formerly *farecast* [8]) and event tickets ([seatgeek.com](http://seatgeek.com)). There is a crucial difference between the domain of tickets and the domain of durable goods which we study in this paper. While both problems have to deal with the uncertainty of future prices, in the case of a durable product, there is a cost associated with waiting to buy the product due to loss of use, which has to be factored into the decision. Thus, while recommending when to buy a ticket is driven solely by potential price changes, recommending when to buy a consumer durable requires a more in-depth modeling of the utility to the consumer.

In order to make good recommendations, we must first be able to predict potential price changes. In [1], we compared and evaluated different forecasting methods incorporating prices, sales volume, seasonality, and competition for predicting prices in consumer durables. We found that methods that use both the current price level and the recent price changes can generally predict product prices well. In this paper, we set out to answer the following questions:

- What is a formal framework for modeling a recommendation system that advises consumers when to buy? How do we measure the performance of such a system and compare alternatives?
- How can we best take advantage of forecasts in coming up with recommendations to the users?
- How much utility can these approaches create for users under realistic settings?

We address the first question by proposing a utility-driven model that captures the trade-off a prospective consumer has to make when faced with the decision of when to buy, namely between potential price drops and the loss in utility due to waiting. An essential feature of the formulation is the *sequential* nature of the recommendations, inspired by Markov decision processes. We address the second question by introducing three natural strategies for recommendations treating forecasts as blackboxes. We conduct a theoretical analysis to compare these strategies, and establish optimality for one of them. We address the last question by conducting an empirical analysis that measures the potential gains users make by following the recommendations. The analysis provides a foundation for building a system that helps forward-looking consumers make more informed decisions.

## 1.1 Road Map

The rest of the paper is organized as follows. We begin with a discussion of related work in Section 2. We introduce a formal model for the recommendation problem by proposing a model of utility of the consumers in Section 3. We then describe three recommendation strategies in Section 4 and

analyze them in Section 5. In Section 6, we present a large-scale experiment using prices in the electronic goods category that quantifies the gains to the consumers from using our system. We discuss the implication of this study and future research directions in Section 7.

## 2 Related Work

The decision between present and future consumption has been studied in economics as the *intertemporal choice* problem. In his seminal work *Theory of Interest*, Fisher proposed that such trade-off should be captured by means of discounting [9]. More recent marketing studies on forward-looking consumer behavior has also been based on discounting, including [20] that estimates the discount rates for different types of products, and [17] that compares different forms of discounting. Inspired by these studies, we propose to model the loss of utility due to waiting by means of discounting. However, while these past proposals attempt to explain consumer behavior through discounting, we are interested in helping consumers that discount future consumption to make the best timing decision.

Our formulation of the problem of when to buy a consumer durable is related to optimal stopping problems studied in probability theory [4]. To our knowledge, no past work has proposed a framework for analyzing this decision using optimal stopping, which we believe provide a natural formulation that allows us to compare different approaches in a formal and rigorous way. One important distinction in our work, in contrast to prior work on optimal stopping, is that the distributions of the random variables are not assumed to be known; rather, their distributions are observed through forecasting algorithms that may be inaccurate.

There have been a number of studies in the marketing literature on the pricing of consumer durables [15, 18]. Price evolutions are often modeled from the product life-cycle point of view, where price changes are linked to stages of the life of a product. Dino applied this model for electronic products at a category level and the model fits well [6]. The methodology has also been extended to incorporating demand models [7]. These prior works suggest that prices for consumer durables follow definite patterns. See [2, 13] for surveys of forecasting methods.

Traditionally, recommender systems focus on suggesting to users items that they may like [14]. Recommendations may be made based on items that similar users like [19], similar items that a user likes [16], or hybrid methods that combine the two [3]. While our work is also about making recommendations to users, the recommendations are about when an item should be purchased. The recommendation is different in nature and is not amenable to a solution by existing techniques.

## 3 System Design

In this section, we present the design of our system. We first propose a formal framework for analyzing the problem of timing the purchase of consumer durables based on utility theory. We then describe how such a system may be implemented as part of a shopping service.

### 3.1 Model

Consider a consumer who has selected a particular product to buy, and is deciding whether to buy the product right away or wait for its price to drop. We treat time as discrete, and denote the latest time until the consumer would wait as time  $T$ , which we call the consumer’s *decision horizon*.

The problem of *timing when to buy* is to decide the time index  $\tau$  within consumer’s decision horizon at which she should buy the product. It is a *sequential decision problem*—at each time step  $s$ , she decides whether to buy the product or wait. If she chooses to wait, she will face another decision at the next time step. The problem repeats itself until she buys the product.

The consumer has a *value*  $\theta$  for the product, which measures the benefit she expects to derive from owning and using the product. In deciding whether to buy now or wait, the consumer faces a trade-off between starting to derive the benefit right away and the loss due to potential price drops in the future. This *intertemporal trade-off* is commonly modeled by *discounting* the value of the product if it is purchased in the future [17, 20]. We denote the discount rate of the consumer by  $\delta$  ( $0 < \delta \leq 1$ ), and take it to be time-invariant within the consumer’s decision horizon. If a consumer decides to buy the product at time  $t$ , she derives a utility of  $\theta\delta^t$ . The quantities  $T$ ,  $\theta$ , and  $\delta$  are specific to the consumer of a given product, and may vary for different consumers and products.

Let the prices of the product from start up to time  $T$  be  $(p_0, p_1, \dots, p_T)$ , which we call the *price path*  $\mathbf{p}$  of the product. The consumer first gets interested in buying the product at time  $t$ . For the purpose of this paper, we assume  $t = 0$ . This is without loss of generality since the prices before time  $t$  cannot change the utility of the consumer, which we define momentarily. At time  $s$ , the prices up to time  $s$  are called the *history*,  $\mathbf{h}_s = (p_0, p_1, \dots, p_s)$ , and prices after  $s$  are called the *future prices*,  $\mathbf{f}_s = (p_{s+1}, \dots, p_T)$ . We sometimes write  $\mathbf{p}$  as  $(\mathbf{h}_s, \mathbf{f}_s)$ .

Given a price path, a value, and a discount rate, the utility of the consumer if she buys the product at time  $s$  equals

$$u(s, \mathbf{p}, \theta, \delta) = \theta\delta^s - p_s \quad . \quad (1)$$

If a consumer possesses a crystal ball that informs her of the price path of the product  $\mathbf{p}$ , deciding when to buy is a simple matter of selecting the time  $t$  that maximizes the utility  $u(t, \mathbf{p}, \theta, \delta)$ . Unfortunately, at time  $s$ , the consumer has only access to the history  $\mathbf{h}_s$ , and has to make a decision without the full knowledge of  $\mathbf{p}$ .

The role of a forecasting algorithm is to make use of the history of a product to forecast its future prices in order to help the consumer make a more informed decision.

**Definition 1.** A *forecasting algorithm*  $F$  maps a history  $\mathbf{h}_s$  to a distribution over future prices  $\mathbf{f}_s$ . We denote by  $E$  the distribution over future prices predicted by the forecasting algorithm. We also write  $E$  as  $F(\mathbf{h}_s)$ .

Using the forecasts over future prices, a *recommendation strategy* makes a recommendation to buy or to wait taking into account the consumer parameters.

**Definition 2.** A *recommendation strategy*  $r$  maps a history  $\mathbf{h}_s$ , a distribution  $E$  over future prices, a value  $\theta$  of the product, and a discount rate  $\delta$  to  $\{0, 1\}$ , where 0 is a recommendation to wait and 1 is that to buy.

As the latest time the consumer will wait before buying the product is time  $T$ , we require that  $r(\mathbf{h}_T, E, \theta, \delta) = 1$  for any distribution  $E$ .

We refer to the pair  $(F, r)$  as a *buy-or-wait recommendation system*, or *system* when the context is clear. Given a system  $(F, r)$ , a price path  $\mathbf{p}$ , and a discount rate  $\delta$ , we can determine when a consumer following its advice will buy the product.

**Definition 3.** Given price path  $\mathbf{p}$ , value  $\theta$  of the product, discount rate  $\delta$ , and system  $(F, r)$ , the *purchase time*  $\tau$  is defined as

$$\tau(\mathbf{p}, \theta, \delta, F, r) = \min_{0 \leq s \leq T} \{s : r(\mathbf{h}_s, F(\mathbf{h}_s), \theta, \delta) = 1\} ,$$

i.e., this is the first time index the system recommends the consumer to buy.

Manufacturers actively manage their revenue through product pricing, and hence product prices often change according to some planned schedule, albeit with uncertainty due to supply, demand, and competition [6, 18]. This results in a distribution over price paths. We use  $D$  to denote this true distribution over price paths. The success of a system can be measured in terms of the expected utility the system generates for the consumer, where expectation is taken with respect to the distribution  $D$  over price paths.

**Definition 4.** Given distribution  $D$  over price paths, value  $\theta$  of the product, discount rate  $\delta$ , the *performance*  $\eta$  of system  $(F, r)$  for a given consumer equals

$$\eta(F, r, D, \theta, \delta) = \mathbb{E}_D \left[ u(\tau(\mathbf{p}, \theta, \delta, F, r), \mathbf{p}, \theta, \delta) \right] , \quad (2)$$

i.e., the expected utility to a consumer who follows the advice of system  $(F, r)$  over the distribution  $D$  of price paths.

Putting it together, our problem can be stated as

**Problem Definition.** Given a set of price paths sampled according to distribution  $D$ , value  $\theta$  of the product, and discount rate  $\delta$ , select a system  $(F, r)$  that maximizes performance  $\eta(F, r, D, \theta, \delta)$ .

### 3.2 Integration With Shopping Services

We envision the recommendation system for helping users to decide when to buy to be integrated as an alert service to existing product search portals. When a user issues a query for a product, a price chart will be displayed to the user to help her understand the historical evolution of the prices for the product. If the user is interested in buying the product, additional information related to the user choices, such as the decision horizon, the value for the product, and the discount rate, will be elicited from the user. The elicitation process can be based on techniques described in [11, 20]. The service will then monitor changes in the product prices, and alert the user when the system recommends buying the product.

## 4 Recommendation Strategies

In this section, we present three recommendation strategies as well as techniques for implementing them. These strategies take as input a consumer's value for a product  $\theta$ , her discount rate  $\delta$  for the product, and the distribution  $E$  of future prices forecasted by algorithm  $F$  at time  $s$  taking into account the price history  $\mathbf{h}_s$ . They output 0 if the recommendation at time  $s$  is to wait and 1 if it is to buy. These strategies differ in how they estimate the utility of waiting.

## 4.1 Expected Utility Strategy

A natural recommendation strategy is to choose to wait if there exists a future time for which the expected utility is greater than zero. It can be formally stated as follows.

**Definition 5.** The *Expected Utility Strategy*,  $\text{EUS}(\mathbf{h}_s, E, \theta, \delta)$ , outputs 1 (buy) if for all  $t, s < t \leq T$ , the expected utility buying at time  $t$  is no better than the expected utility from buying at time  $s$ . That is, it outputs 1 if

$$\begin{aligned} \max_{s < t \leq T} \mathbb{E}_E[u(t, (\mathbf{h}_s, \mathbf{f}_s), \theta, \delta)] &\leq \mathbb{E}_E[u(s, (\mathbf{h}_s, \mathbf{f}_s), \theta, \delta)] \\ &= (\theta\delta^s - p_s) , \end{aligned} \quad (3)$$

and 0 (wait) otherwise.

EUS is an easy strategy to implement. First, the RHS is a fixed quantity. To evaluate the LHS of Eq (3), by linearity of expectation,

$$\max_{s < t \leq T} \mathbb{E}_E[u(t, (\mathbf{h}_s, \mathbf{f}_s), \theta, \delta)] = \max_{s < t \leq T} \left( \theta\delta^t - \mathbb{E}_E[p_t] \right) . \quad (4)$$

The expensive part of the computation involves computing the expected price at time  $t$  according to the distribution  $E$ . For many distributions (e.g. normal distribution and uniform distribution), this value can be computed efficiently by evaluating a closed-form formula. In general, it can be computed through sampling and taking the sample average.

**Example 1.** We give an example of how EUS is evaluated. Suppose that  $s = 0, p_0 = 100, \theta = 100, \delta = 1, T = 2$ , and  $E$ , the distribution of future prices  $\{p_1, p_2\}$ , is given as

$$\begin{aligned} E(\{105, 105\}) &= 0.25 & E(\{105, 95\}) &= 0.25 \\ E(\{95, 105\}) &= 0.25 & E(\{95, 95\}) &= 0.25 \end{aligned} .$$

The RHS of Eq (3) equals  $(100 - 100) = 0$ . The LHS equals

$$\begin{aligned} \max_{0 < t \leq 2} \mathbb{E}_E[u(t, \mathbf{p}, 100, 1)] &= \max_{0 < t \leq 2} (100 - \mathbb{E}_E[p_t]) \\ &= \max_{0 < t \leq 2} (100 - 100) = 0 . \end{aligned}$$

Hence, EUS recommends to buy. □

## 4.2 Sample Path Strategy

Another natural recommendation strategy, motivated by options pricing theory [12], is to generate and evaluate sample price paths according to the forecasted distribution. If, in expectation, there exists a future time for which the expected utility is greater than the utility of buying the product now along the path, it recommends to wait. It can be formally stated as follows.

**Definition 6.** The *Sample Path Strategy*,  $\text{SPS}(\mathbf{h}_s, E, \theta, \delta)$ , outputs 1 if the expected maximum utility of buying at some future time  $t, s < t \leq T$ , is no better than the expected utility of buying at time  $s$ , i.e.,

$$\mathbb{E}_E\left[\max_{s < t \leq T} u(t, (\mathbf{h}_s, \mathbf{f}_s), \theta, \delta)\right] \leq (\theta\delta^s - p_s) , \quad (5)$$

and outputs 0 otherwise.

The difference between EUS and SPS is the order in which expectation and maximization is taken. Whereas under EUS, expectation is taken separately over each time index, and the time index with the highest utility is selected, under SPS, expectation of the highest utility along an entire price path is taken. Their values could be different for some distribution  $E$ . We explore this further in Section 5.

To implement SPS, one needs to evaluate the LHS of Eq (5). Unlike EUS, the value of the LHS of Eq (5) depends on the joint distribution of future prices  $\mathbf{f}_s$ , since the expectation is taken over price paths, and hence does not lend itself to the same decomposition as in Eq (4). A parsimonious approach to implement this strategy is to sample a set of price paths and evaluate the maximum on each price path, and estimate the expectation by averaging the sampled values.

**Example 1 (cont'd).** Under the same setting, we give an example of how SPS is evaluated. The RHS remains unchanged. The LHS of Eq (5) equals

$$\begin{aligned} \mathbb{E}_E[\max_{0 < t \leq 2} u(t, \mathbf{p}, 100, 1)] &= 0.25(100 - 105) + 0.25(100 - 95) \\ &\quad + 0.25(100 - 95) + 0.25(100 - 95) = 2.5 \ . \end{aligned}$$

Hence, SPS recommends to wait. □

### 4.3 Self Application Strategy

The final recommendation strategy we consider is defined recursively through self application. It makes a buy recommendation if the expected utility obtained from waiting at time  $s$  and following itself thereafter is no larger than that of buying the product at time  $s$ . In other words, the strategy estimates the expected utility of waiting by computing how much utility it would gain if it follows its own recommendation starting from time  $(s + 1)$ , i.e., applies itself at time  $(s + 1)$ , and compares the expected utility to that of buying at time  $s$ .

To give a formal definition of the strategy, we introduce a notation for conditioning the present forecast with new prices. Suppose that the present forecast over future prices  $\mathbf{f}_s$  is  $E$ . If the price at time  $(s + 1)$  equals  $p_{s+1}$ , then the *updated* forecast over future prices  $\mathbf{f}_{s+1}$ , denoted by  $E|p_{s+1}$ , is a distribution over  $\mathbf{f}_{s+1}$  that satisfies

$$\mathbb{P}_{E|p_{s+1}}(\mathbf{f}_{s+1}) = \mathbb{P}_E((p_{s+1}, \mathbf{f}_{s+1}) \mid p_{s+1}) \ ,$$

i.e., for any future prices  $\mathbf{f}_{s+1}$ , the probability of the future prices according to  $E|p_{s+1}$  equals the probability according to  $E$  conditioned on the price at time  $(s + 1)$  being  $p_{s+1}$ .<sup>1</sup>

**Definition 7.** The *Self Application Strategy*,  $\text{SAS}(\mathbf{h}_s, E, \theta, \delta)$ , outputs 1 if and only if the expected utility obtained from waiting at time  $s$  and following SAS thereafter is no larger than the expected utility of buying at time  $s$ , i.e.,

$$\mathbb{E}_E[w(p_{s+1}, |\mathbf{h}_s| + 1, E|p_{s+1}, \theta, \delta)] \leq (\theta\delta^s - p_s) \ , \tag{6}$$

---

<sup>1</sup>More formally,  $\mathbb{P}_E$  is a measure over the elements  $\mathbf{f}_s$ , and  $\mathbb{P}_{E|p_{s+1}}$  is a measure over the elements  $\mathbf{f}_{s+1}$ . For any event  $e$  measurable over the  $\sigma$ -field generated by  $\mathbf{f}_{s+1}$ ,  $\mathbb{P}_{E|p_{s+1}}(e) = \mathbb{P}_E(e|p_{s+1})$ . For exposition, we simplify the statement and state the results as if the distribution is discrete.

where  $w(p, t, E, \theta, \delta)$  is defined as the utility obtained from following the recommendations of SAS after having waited for the first  $(t - 1)$  time steps. It can be written recursively as the maximum between the utility of buying the product at time  $t$ , or waiting at time  $t$  and following SAS thereafter,

$$w(p, t, E, \theta, \delta) = \max \begin{cases} \theta\delta^t - p \\ \mathbb{E}_E[w(p_{t+1}, t + 1, E|p_{t+1}, \theta, \delta)] \end{cases} \quad (7)$$

and for time  $t = T$ ,

$$w(p, T, E, \theta, \delta) = \theta\delta^T - p . \quad (8)$$

**Remark 1.** SAS depends on history  $\mathbf{h}_s$  only through the current price  $p_s$  and the length of the history  $|\mathbf{h}_s|$ . We separate these factors in the recursion to highlight how they evolve. This dependence relationship is also true for EUS and SPS. We make use of this relationship in Section 5.

Like SPS, SAS can be implemented by sampling and estimating the expectations via sample averages. However, a direct sampling algorithm will run in time exponential in  $T$  due to its recursive nature. This is because each evaluation of  $w(p, t, E, \theta, \delta)$  according to Eq (7) requires evaluating an expectation that calls the function  $w$  with different arguments. If  $k$  samples are drawn for estimating the expectation by sample average, a total of  $O(k^{T-s})$  number of evaluations of Eq (8) will be needed.

Based on experiments reported in [1], autoregressive model of order 1 is often a good fit for product prices. Under this model, prices follow a *Markov process*, i.e., the distribution of future prices  $\mathbf{f}_s$  depend only on the price  $p_s$  but not the prices before. In this case, we can speed up the computation exponentially by *backward induction*.

Starting from time  $T$ , we draw  $k$  samples, and evaluate the utility of buying at time  $T$ . We then compute the value of waiting till time  $(T - 1)$  by using these values in evaluating the expectations, and taking the maximum according to Eq (7). We proceed backwards in time until we reach time  $s$ . A formal description is given in Algorithm 1.



---

**Algorithm 1:** SAS Evaluation by Backward Induction

---

**Input** : History  $\mathbf{h}_s$ ; Dist.  $E$ ; Discount rate  $\delta$ ; Value  $\theta$

**Output**: LHS of Eq (6),  $\mathbb{E}_E[w(p_{s+1}, s + 1, E|p_{s+1}, \theta, \delta)]$

**Params**: Samples per time step,  $k$

/\* The array entry  $p[t, i]$  stores the  $i$ -th sampled price for time  $t$ , and the array entry  $w[t, i]$  stores the value estimated for  $w(p[t, i], t, E|p[t, i], \theta, \delta)$ . SAMPLE and PROB are two auxiliary functions that samples uniformly at random from the domain of  $E$  and evaluates transition probabilities according to  $E$ . \*/

```
for  $i \leftarrow 1$  to  $k$  do
  |  $p[T, i] \leftarrow \text{SAMPLE}(E, T)$ 
  |  $w[T, i] \leftarrow \theta\delta^T - p[T, i]$ 
end
for  $t \leftarrow (T - 1)$  downto  $(s + 1)$  do
  | for  $i \leftarrow 1$  to  $k$  do
    | |  $p[t, i] \leftarrow \text{SAMPLE}(E, t)$ 
    | |  $v \leftarrow \sum_{j=1}^k \text{PROB}(E, t, p[t, i], p[t + 1, j]) \times w[t + 1, j]$ 
    | |  $s \leftarrow \sum_{j=1}^k \text{PROB}(E, t, p[t, i], p[t + 1, j])$ 
    | | 1 |  $w[t, i] \leftarrow \max\{\theta\delta^t - p_t, v/s\}$ 
    | end
  | end
  |  $v \leftarrow \sum_{j=1}^k \text{PROB}(E, s + 1, p_s, p[s + 1, j]) \times w[s + 1, j]$ 
  |  $s \leftarrow \sum_{j=1}^k \text{PROB}(E, s + 1, p_s, p[s + 1, j])$ 
2 return  $v/s$ 
```

---

The algorithm depends on two auxiliary functions. The function  $\text{SAMPLE}(E, t)$  draws a price *uniformly at random* from the domain of prices at time  $t$  under distribution  $E$ . The function  $\text{PROB}(E, t, p, p')$  evaluates the transition probability from price  $p$  to price  $p'$  at time  $t$  according to the distribution  $E$ . Prices are drawn uniformly at random from their domains to avoid factoring in the probabilities twice, as the samples are subsequently weighted by the transition probabilities. Lines 1 and 2 estimate the expectations needed in Eqs (6) and (7) by sample averages ( $v/s$ ).

Algorithm 1 produces an unbiased estimate of the LHS of Eq (6) when prices follow a Markov process, since the probability of transitioning from price  $p_t$  to  $p_{t+1}$  is independent of prices  $p_{t'}$  for  $t' < t$ . Hence, for price paths  $\mathbf{p}$  and  $\mathbf{p}'$ ,  $E|(p_s, p_{s+1}, \dots, p_t)$  equals  $E|(p'_s, p'_{s+1}, \dots, p'_t)$  if  $p_t = p'_t$ .

The running time of the algorithm is  $O(k^2T)$ , since each computation of  $w[t, i]$  requires a summation over  $k$  terms, which is computed once for each of the  $k$  samples per time step, and there are  $T$  time steps. Though the time is polynomial in  $k$  and  $T$ , the algorithm implicitly considers an exponential number of price paths  $\mathbf{p}$ , as the final expectation takes into account each combination of prices  $p[t, i]$  from every time index  $t$  and  $i$  from 1 to  $k$ .

**Example 1 (cont'd).** Under the same setting, we give an example of how SAS is evaluated. The RHS remains unchanged. The LHS of Eq (6) equals

$$\begin{aligned} \mathbb{E}_E[w(p_1, 1, E|p_1, 100, 1)] &= \mathbb{E}_E[\max\{100 - p_1, \mathbb{E}_{E|p_1}[100 - p_2]\}] \\ &= \mathbb{E}_E[\max\{100 - p_1, 0\}] = 0.5(5) + 0.5(0) = 2.5 . \end{aligned}$$

Hence, SAS recommends to wait. □

## 5 Analysis

We now analyze the strategies presented in Section 4. First, we study whether the recommendations made by the three strategies are related. Second, we consider an idealized situation where the forecasting algorithm is perfect, and determine which of the strategies works best. For exposition, we present the results as if the distributions are discrete. The statements can be made formal for continuous distributions.

### 5.1 Relationship Among the Strategies

Without loss of generality, we take  $s = 0$ . When the decision is between buying now or the next period, i.e.,  $T = 1$ , all three strategies makes the same recommendation.

**Theorem 1.** *When  $T = 1$ , for any distribution  $E$ , value  $\theta$ , and discount rate  $\delta$ , EUS, SPS, and SAS make the same recommendation.*

*Proof.* When  $T = 1$ , the decision is between buying at time 0 and at time 1. It is clear from Eq (3), (5), and (6) that the RHS of the three strategies are equal. The LHS of Eqs (3), (5), and (6) are also equal, since for discrete  $t$ , if  $0 < t \leq 1$ ,  $t = 1$ . Hence, the LHS are all equal to

$$\mathbb{E}_E[u(1, (p_0, p_1), \theta, \delta)] = \mathbb{E}_E[\theta\delta - p_1] . \quad \square$$

When  $T \geq 2$ , the strategies may make different recommendations. Nonetheless, the recommendations are related, and we can show that SPS is more optimistic about the value of waiting than SAS, and SAS is more optimistic about the value of waiting than EUS.

**Theorem 2.** *For any history  $\mathbf{h}_t$ , distribution  $E$ , value  $\theta$ , and discount rate  $\delta$ , if EUS recommends to wait, then SAS will also recommend to wait. If SAS recommends to wait, then SPS will also recommend to wait.*

To prove this theorem, note that the recommendations made by these strategies are determined by the values of the LHS of Eqs (3), (5), and (6). These values can be interpreted as the values of waiting estimated by the strategies. Let  $v_{\text{EUS}}(\mathbf{h}_s, E, \theta, \delta)$ ,  $v_{\text{SPS}}(\mathbf{h}_s, E, \theta, \delta)$ , and  $v_{\text{SAS}}(\mathbf{h}_s, E, \theta, \delta)$  denote the respective LHS of the strategies for history  $\mathbf{h}_s$ . We claim that for all history  $\mathbf{h}_s$ ,

$$v_{\text{EUS}}(\mathbf{h}_s, E, \theta, \delta) \leq v_{\text{SAS}}(\mathbf{h}_s, E, \theta, \delta) \leq v_{\text{SPS}}(\mathbf{h}_s, E, \theta, \delta) . \quad (9)$$

Let  $v_{\text{buy}}(\mathbf{h}_s, \theta, \delta)$  denote the value of the RHS of the strategies for history  $\mathbf{h}_s$ . If the claim is true, then if EUS recommends to wait, so will SAS, since  $v_{\text{SAS}}(\mathbf{h}_s, E, \theta, \delta) \geq v_{\text{EUS}}(\mathbf{h}_s, E, \theta, \delta) \geq v_{\text{buy}}(\mathbf{h}_s, \theta, \delta)$ . Similarly, if SAS recommends to wait, so will SPS, since  $v_{\text{SPS}}(\mathbf{h}_s, E, \theta, \delta) \geq v_{\text{SAS}}(\mathbf{h}_s, E, \theta, \delta) \geq v_{\text{buy}}(\mathbf{h}_s, \theta, \delta)$ .

To show the claim, we examine the order in which the expectation and the maximum is taken under the three strategies. The relationship between EUS and SPS can be easily established by Jensen's Inequality. The difficulty is with SAS. The proof makes use of an inductive argument and Jensen's inequality to establish the desired relationship, and can be found in the Appendix.

Recall that Remark 1 states that the three strategies depend on history  $\mathbf{h}_s$  only through the current price  $p_s$ , and its length  $|\mathbf{h}_s|$ . Further, the value of waiting does not depend on the current price  $p_s$ . Henceforth, we rewrite the function  $v_{\text{EUS}}(\mathbf{h}_s, E, \theta, \delta)$  as  $v_{\text{EUS}}(s, E, \theta, \delta)$ , and likewise for SPS and SAS.

## 5.2 Performance Under Perfect Forecasts

We now consider the performances of the three strategies when the forecasting algorithm is *perfect*, which informally means that the forecasted distribution coincides with the true distribution  $D$ . Note that it does *not* mean that the price path itself is known with certainty. While it is not possible to design a forecasting algorithm that is perfect, this analysis helps us to isolate the difference in performance that is solely explained by the choice of recommendation strategy. It also helps to develop insights as to why some of the strategies may fail to make the best recommendation despite very good forecasts.

We first define a property of forecasting algorithm about how current forecasts are related to future ones.

**Definition 8.** A forecasting algorithm  $F$  is *consistent* if for all histories  $\mathbf{h}_s$  and  $\mathbf{h}_t$  where  $\mathbf{h}_t$  includes  $\mathbf{h}_s$  as its first  $s$  prices, i.e.,  $\mathbf{h}_t = (\mathbf{h}_s, p_{s+1}, \dots, p_t)$ ,  $F(\mathbf{h}_t)$  produces a distribution identical to the *conditional* distribution induced by  $F(\mathbf{h}_s)$ , i.e., for all  $\mathbf{f}_t$ ,

$$\mathbb{P}_{F(\mathbf{h}_t)}(\mathbf{f}_t) = \mathbb{P}_{F(\mathbf{h}_s)|(p_{s+1}, \dots, p_t)}(\mathbf{f}_t) .$$

Intuitively speaking, a consistent forecasting algorithm is one where there are “no surprises” when it updates its forecast. Given the current forecast, when a new price arrives, one just needs to condition on the latest observation to obtain the updated forecast.

A consistent forecasting algorithm is related to a collection of distributions over price paths. Consider a consistent forecasting algorithm  $F$ . Given any history  $\mathbf{h}_t$ , it produces a distribution  $F(\mathbf{h}_t)$  over the price paths  $\mathbf{f}_t$ , which by consistency equals  $F(p_0)|(p_1, \dots, p_t)$ . Therefore, it can be viewed as a collection of distributions of price paths indexed by  $p_0$ . Conversely, given a distribution  $E$  of price paths  $\mathbf{f}_t$ , it can be turned into a forecasting algorithm for  $\mathbf{f}_{t+1}$ , such that given any price  $p_t$ , it produces a conditional distribution over  $\mathbf{f}_{t+1}$ . This duality between distributions and consistent forecasting algorithms will be exploited further in this section. For semantic clarity, we denote the forecasting algorithm for prices  $\mathbf{f}_{t+1}$  induced by distribution  $E$  over price paths  $\mathbf{f}_t$  as  $[E]$ .

A forecasting algorithm  $F$  is perfect for distribution of price paths  $D$  when  $F$  and  $D$  are “interchangeable”.

**Definition 9.** A forecasting algorithm  $F$  is *perfect* for distribution of price paths  $D$  if  $F$  is consistent, and  $F(\mathbf{h}_0)$  produces the same distribution as  $D$  conditional on  $p_0$ , i.e., for all price paths  $\mathbf{p}$ ,

$$\mathbb{P}_{F(p_0)}(\mathbf{f}_0) = \mathbb{P}_D((p_0, \mathbf{f}_0) | p_0) .$$

Another way to interpret a perfect forecasting algorithm is that  $[D]$  gives the same forecasts for any  $p_0$  as  $F$ . Note that if  $F$  is perfect, by property of consistency, for all history  $\mathbf{h}_t$ ,

$$\mathbb{P}_{F(\mathbf{h}_t)}(\mathbf{f}_t) = \mathbb{P}_D((\mathbf{h}_t, \mathbf{f}_t) | \mathbf{h}_t) .$$

Hence, other than not knowing the distribution of the first price  $p_0$ ,  $F$  has as much information as  $D$ .

We now consider the performance of the three strategies under perfect forecasts. We start with an example that shows that EUS is not optimal.

**Example 2** (Suboptimality of EUS). Suppose that  $s = 0$ ,  $p_0 = 100$ ,  $\theta = 100$ ,  $\delta = 1$ ,  $T = 2$ . Let the true distribution  $D$  of price paths be

$$\begin{aligned} D(\{100, 105, 105\}) &= 0.25 & D(\{100, 105, 95\}) &= 0.25 \\ D(\{100, 95, 105\}) &= 0.25 & D(\{100, 95, 95\}) &= 0.25 \end{aligned} .$$

This setup is identical to Example 1, and assume that  $E = F(\{100\})$  in Example 1 is a perfect forecast for  $D$ . We have shown that

$$\begin{aligned} v_{\text{EUS}}(0, F(\{100\}), 100, 1) &= 0 \\ v_{\text{SPS}}(0, F(\{100\}), 100, 1) &= 2.5 \\ v_{\text{SAS}}(0, F(\{100\}), 100, 1) &= 2.5 \end{aligned}$$

and EUS will recommend to buy at time 0, whereas SPS and SAS will recommend to wait. One can verify that waiting is better, since there is a 0.5 chance that the price will drop to 95 at time 1, at which point the consumer can buy and earn a utility of 5, and if the price goes to 105 at time 1, the consumer can choose to wait and earn an expected utility of 0 from buying at time 2.  $\square$

This example reveals a weakness in the estimation procedure used in EUS, namely that it fails to take into account the sequential nature of the decisions to be made. Even though waiting and buying at the next time step may earn zero utility in expectation, one may be able to choose differently depending on the outcome at the next time step. Hence EUS may systematically underestimate the value of waiting.

Our next example shows that SPS is not optimal either.

**Example 3** (Suboptimality of SPS). Consider the same setup as Example 2, but with a change to the true distribution  $D$  of price paths

$$\begin{aligned} D(\{100, 110, 125\}) &= 0.25 & D(\{100, 110, 95\}) &= 0.25 \\ D(\{100, 95, 125\}) &= 0.25 & D(\{100, 95, 95\}) &= 0.25 \end{aligned} .$$

Under a perfect forecast, the value of the three strategies are:

$$\begin{aligned} v_{\text{EUS}}(0, F(\{100\}), 100, 1) &= \max\{100 - 102.5, 100 - 110\} = -2.5 \\ v_{\text{SPS}}(0, F(\{100\}), 100, 1) &= 0.25(100 - 110) + 0.25(100 - 95) \\ &\quad + 0.25(100 - 95) + 0.25(100 - 95) = 1.25 \\ v_{\text{SAS}}(0, F(\{100\}), 100, 1) &= 0.5(\max\{100 - 95, 100 - 115\}) \\ &\quad + 0.5(\max\{100 - 110, 100 - 110\}) = -1.25 . \end{aligned}$$

Now, SPS will recommend to wait at time 0, whereas EUS and SAS will recommend to buy. In this case, buying is better than waiting at time 0. If one waits, at time 1, if  $p_1 = 95$ , one can at most gain a utility of 5 by buying at time 1. But if  $p_1 = 110$ , then both waiting and buying at time 1 lead to an expected utility of  $-10$ . Combining the two, it is better to buy at time 0.  $\square$

This example reveals a weakness in SPS, namely that it assumes it can pick out the best time to buy as if the price path is known. In the above example, it assumes that if the price path is  $\{100, 110, 125\}$ , it will buy at time 1; but if the price path is  $\{100, 110, 95\}$ , it will buy at time 2. This is an inconsistent choice, since the price histories at time 1 in both cases are identical. As a result, SPS overestimates the value of waiting.

**Remark 2.** In Example 3, all three strategies produce different estimated values for waiting. Hence, there exists distributions for which the inequality in Eq (9) holds strictly.

Through these two examples, we observe that SAS accounts for the sequential nature of the decision problem, and ensures consistency in the choices through self application. It turns out that these are the key properties for a recommendation strategy to be optimal in expectation.

**Theorem 3.** *For any distribution of price paths  $D$ , value  $\theta$ , discount rate  $\delta$ , if the forecasting algorithm  $F$  is perfect, there does not exist a recommendation strategy  $r$  that has higher performance than SAS, i.e., for all strategies  $r$ ,  $\eta(F, r, D, \theta, \delta) \leq \eta(F, SAS, D, \theta, \delta)$ .*

The key to establishing this theorem is that the value SAS estimated for waiting precisely measures the performance it will obtain if it decides to wait. This is provided by the following ‘‘Lemma’’. We put this lemma in quotes because the statement of the Lemma is not sufficiently formal. A formal statement of the Lemma involves some subtlety, and is presented as Lemma 6 in the Appendix.

‘‘**Lemma.**’’ *For any  $\delta, \theta, \mathbf{h}_s$  for  $s \in [0, T)$ ,*

$$v_{SAS}(s, F(\mathbf{h}_s), \theta, \delta) = P([F(\mathbf{h}_s)], SAS, F(\mathbf{h}_s), \theta, \delta) .$$

Using this lemma, we can establish using an inductive argument that SAS is optimal in expectation. The proof is similar in spirit to the optimality proof of using dynamic programming (Bellman equations) to solve Markov decision processes, but with added subtlety due to the semantic differences between a forecasting algorithm and a distribution. The proof can be found in the Appendix.

### 5.3 Implications

We have established two theoretical results. The first relates the recommendations made by the three strategies, and holds for *any* forecasting algorithm. Theorem 1 shows that the choice of recommendation strategies does not matter when the decision horizon is over a single period. Theorem 2 establishes a relationship useful for post-hoc analysis of empirical performances. For example, if we observe that SPS outperforms EUS, we can infer that the system under EUS is erring on the side of making too many buy recommendations, as for any data set and forecasting algorithm, the set of situations where EUS makes a buy recommendation is a superset of ones for SPS.

The second result shows that when forecasts are perfect, SAS is the optimal recommendation strategy. We believe that SAS continues to be the best strategy when forecasts are imperfect, although we cannot directly establish its optimality theoretically, as a poor recommendation may be either due to the strategy or the forecasts. We address this question by means of experiments in the next section.

## 6 Experiments

We now conduct an experimental study to evaluate the performance of the three proposed recommendation strategies. The objective of this study is to quantify the performances under realistic settings of price changes and forecasting errors. It also complements the theoretical analysis, which shows the optimality of SAS, by quantifying the differences in performances among EUS, SPS, and SAS.

## 6.1 Setup

To generate realistic distribution of price paths, we analyzed price changes in four categories of electronic products—Televisions, Digital cameras, Camcorders, and Printers. The data was obtained for prices between January 2005 and September 2008 from NPD, a firm that tracks prices registered at major retail stores. We found that the average price decrease per time period for these products was about 3.5%, with a standard deviation of about 5%. We generate synthetic data based on these statistics, and systematically vary different factors to measure the performances of the strategies under different settings.

We consider a family of distributions to model the price paths of these products based on how persistent are unexpected price changes, which we call *shocks*. Formally, suppose that the average decrease per time period is  $\mu$ , we define the shock at time  $t$  as the difference between the actual price  $p_t$  and the expected price  $p_0\mu^t$ . We select a parameter  $\alpha \in [0, 1]$  that controls *shock persistence*. In more detail, from time period  $t$  to time period  $(t + 1)$ , the change in prices is sampled from a normal distribution  $N(\mu, \sigma)$ . Let  $p_t$  and  $q_t$  denote the price and the shock in time period  $t$ . The price  $p_{t+1}$  in time period  $(t + 1)$  equals

$$p_{t+1} = d(p_t - (1 - \alpha)q_t)$$

where  $d$  is sampled from  $N(\mu, \sigma)$ .

Shock persistence is important as it models whether product prices exhibit definite patterns. When its value is low, the prices follow definite paths; when its value is high, the prices appear random. At the extremes, when  $\alpha = 0$ , shocks are *transient* and do not carry over from one time period to the next. This setting gives rise to a distribution where the prices of one period is independent of another, since the price  $p_{t+1}$  at time period  $(t + 1)$  depends on  $(p_t - q_t) = p_0\mu^t$ , a quantity that is independent of  $p_t$ . It can be shown that the price at time period  $t$  is sampled from  $N(p_0\mu^t, p_0\mu^{t-1}\sigma)$ . In contrast, when  $\alpha = 1$ , shocks are *persistent* and affect all future prices. This gives rise to a distribution of price path that resembles random walks, since the price  $p_{t+1}$  at time period  $(t + 1)$  depends on  $p_t$ , the sampled price of time period  $t$ .

The first decision starts at time 0 (i.e.  $s = 0$ ), and the decision horizon  $T$  of the consumer is selected to be between four and eight time periods, which we believe is a reasonable timeframe for purchasing consumer durables. The discount rate  $\delta$  is selected to be between 95.5% and 97.5%, and is comparable to the study of [20], after adjusting for technological improvements. Finally, the value of the product  $\theta$  is selected to be a multiple  $\rho$  of the initial price  $p_0$ , i.e.,  $\theta = \rho p_0$ . We vary  $\rho$  between 1.0 and 1.1, our best guess on when a user starts to be interested in purchasing the product.

To evaluate the effect of imperfect forecasting algorithms on the performances of the strategies, we introduce two parameters,  $\text{err}_\mu$  and  $\text{err}_\sigma$ , that control the accuracy of the estimated mean and that of the estimated standard deviation of the forecasts. We assume that the forecasting algorithm correctly identifies the level of shock persistence, but may make (normally-distributed) errors in estimating the parameters to the distribution. For example, with  $\mu = 3.5\%$ , it estimates the mean to be  $\mu'$ , where  $\mu'$  is sampled from the distribution  $N(\mu, \text{err}_\mu)$ . Likewise, a mistake may be made in estimating the standard deviation. Based on our price prediction experiments conducted in [1], values of  $\text{err}_\mu$  around 3% and  $\text{err}_\sigma$  around 1% are consistent with the errors made by the best forecasting algorithms we evaluated.

The simulation parameters are summarized in Table 1, where we give the default choices for these parameters along with their minimum and maximum values in our sensitivity analysis. In

Parameter	Default	Min	Max
$\mu$ (mean price decrease)	3.5%	2.5%	4.5%
$\sigma$ (stdev. of price decrease)	5%	3%	6%
$\alpha$ (shock persistence)	0	0	1
$T$ (decision horizon)	6	4	8
$\delta$ (discount rate)	96.5%	95.5%	97.5%
$\rho$ (ratio of init. value to price)	1.05	1.00	1.10
$\text{err}_\mu$ (stdev. of estimated $\mu$ )	0%	0%	4%
$\text{err}_\sigma$ (stdev. of estimated $\sigma$ )	0%	0.5%	2%

Table 1: Summary of simulation parameters.

the subsequent experiments, we vary one parameter at a time between its minimum and maximum values to understand its impact on the performances. For each combination of parameters, we sampled 100 price paths according to the true distribution, and for each path, evaluated each strategy 100 times on the path to obtain its average performance. For imperfect forecasting algorithms, the estimated mean and the estimated standard deviation are sampled once for each run.


In reporting the results, we present the *relative* performances. This is done by normalizing the performance of a strategy by that of an *omniscient* strategy that knows all the prices on the price path and selects the optimal time to buy. The omniscient strategy is not constrained to make the same choice for the same price history. Thus, its value constitutes an upper bound on the utility that can be gained on each path. As it requires knowing the price path in advance, its performance is unachievable in practice. Nonetheless, it helps us to gauge how close to optimal we can achieve. We also report the performance of always buying at time 0 as a baseline. Since relative performance is scale-free and does not depend on  $p_0$ ,  $p_0$  is not varied in the experiments.

## 6.2 Perfect Forecasts with Transient Shocks


In the first set of experiments, we focus on the case where the forecasting algorithm is perfect and shocks are transient. While we already know from the theoretical analysis that SAS is optimal for this setting, the experiments serve two important purposes. First, they serve to quantify the differences in performances between SAS and the other two strategies. Second, they identify factors that influence the performances of these strategies.

The performances with a sensitivity analysis on each of the five variables are shown in Figure 1. Key findings are


- (1) Across all conditions, all three strategies outperform the baseline of always buying at time 0. The performances of SAS and SPS are consistently higher than that of EUS. In 70% of the settings, SAS and SPS achieves about 90% of the maximum possible savings. The differences in performances are statistically significant under a *paired t-tests* with  $p$ -values  $< 0.001$  for all but one cases. Despite the visual similarity, SAS outperforms SPS in all of the cases.
- (2) In most of the cases, EUS recommends buying at time 0, hence its performance overlaps with the baseline. The cases where the recommendations are different are when  $\mu > (1 - \delta)$  (Figures 1(a) and (d)).
- (3) As Figures 1(a) and (d) show, varying  $\mu$  and  $\delta$  leads to almost identical results. This is because  $\mu$  and  $(1 - \delta)$  are closely related. Generally speaking, when  $\mu > (1 - \delta)$ , it is better to wait, since




(a) Relative performance vs.  $\mu$




(b) Relative performance vs.  $\sigma$



(c) Relative performance vs.  $T$



(d) Relative performance vs.  $\delta$



(e) Relative performance vs.  $\rho$

Figure 1: Relative performances under different simulation parameters






Figure 2: Relative perf. vs. shock persistence

prices are dropping faster than utility loss due to discounting; conversely, when  $\mu < (1 - \delta)$ , it is better to buy. However, when the standard deviation of the decrease is high, one could gain by waiting due to price fluctuations even when  $\mu$  is smaller than  $(1 - \delta)$ . Both SAS and SPS (correctly) recommend waiting in such cases, leading to large outperformance compared to EUS.

(4) As Figure 1(b) shows, as the standard deviation of the price decrease ( $\sigma$ ) increases, the relative performance of EUS decreases, highlighting the importance of considering more than just the expected utility. The performances of SAS and SPS keep pace with that of the omniscient strategy.

(5) As Figure 1(c) shows, the relative performances are almost constant across different decision horizons ( $T$ ). This suggests that  $T$  is not a critical factor in the simulation.

(6) As Figure 1(e) shows, the relative performances improve as the ratio of initial value to price ( $\rho$ ) increases. The shape of the curve is concave (downwards) due to the following reason. As  $\rho$  increases, the absolute performances for all strategies increases due to the larger difference between  $\theta (= \rho p_0)$  and  $p_0$ . On the other hand, the expected utility obtained from timing the purchase better, which applies to SAS, SPS, and the omniscient strategy, grows at a much slower pace. Hence, the curves are shaped concave downwards.

In summary, in the majority of the experimental conditions, SAS performs within 90% of the maximum utility that can only be obtained knowing the entire price path. A consumer following SAS does better than the baseline of always buying by about 40% of the maximum utility on average.

### 6.3 Perfect Forecasts with Persistent Shocks

The next set of experiments evaluates the relative performances of the strategies under different shock persistence level ( $\alpha$ ). We single out the shock persistence as a separate experiment since its level has a profound impact on the behavior of the price paths. As mentioned, at  $\alpha = 0$ , shocks are transient and the distribution is independent from time period to time period; at  $\alpha = 1$ , the distribution of price paths resembles that of a random walk. The results are presented in Figure 2.

We observe that SAS outperforms the other two strategies. However, its relative performance decreases as  $\alpha$  increases. This is because as the distribution of price paths becomes more like that of a random walk, timing the purchase without knowledge of the price path becomes difficult. For an increasing fraction of the price paths, the price shocks stay either positive or negative over the entire path. While upto  $\alpha = 0.8$ , a strategy may yield a positive expected utility by recommending to wait at time 0, as  $\alpha$  increases, a larger fraction of the gains are offset by losses on price paths for

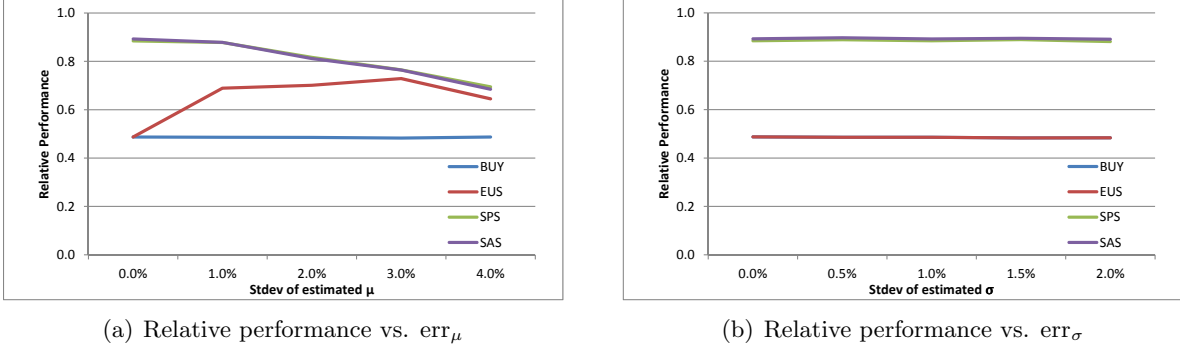


Figure 3: Relative perf. under imperfect forecasts

which the price shocks stay positive throughout. On the other hand, since the omniscient strategy makes its choices knowing the price paths, it can selectively wait on ones for which waiting yields positive expected utilities, and buy on ones where waiting would have yielded negative expected utilities. These together lead the relative performances of SAS to become lower as  $\alpha$  increases.

Note that as  $\alpha$  increases, SPS performs worse and worse, and this is one of few cases where we can separate visually the relative performances of SAS and SPS. The problem here is that under a random-walk like price path, SPS will systematically overestimate the value of waiting as highlighted in Example 3, leading to its recommending the user to wait when it is not advisable to do so. Hence its relative performance deteriorates faster than SAS as  $\alpha$  increases.

Based on the marketing literature surveyed in Section 2, we believe that the price evolution of durable products exhibit definite patterns and do *not* resemble random walks. Thus, the regime in which our system operates is likely to have small  $\alpha$ . Nonetheless, these experiments highlight the importance of this assumption. The shock persistence levels determine how well we can do for a given category of products.

## 6.4 Imperfect Forecasts

The final set of experiments evaluates the relative performances of the strategies under imperfect forecasting algorithms. The degree of imperfection is measured by the size of the errors in estimating the mean and the standard deviation of the distribution. These experiments help to quantify the true performances of the system, as forecasting algorithms in real application do make errors. Our parameters are chosen to reflect realistic levels of errors as measured by related experiments conducted in [1]. The results under transient shocks are presented in Figure 3. We observe similar trends for persistent shocks.

The results depend on what kind of errors is made. When errors involve estimating the means of the price decreases, both SAS and SPS perform worse than when forecasts are perfect, and their performances get worse as the errors get larger. On the other hand, the performance of EUS *improves* for small errors. This is because at the default parameter values,  $\mu = (1 - \delta) = 3.5\%$ , EUS will always recommend to buy at time 0. When there are errors in the estimated mean  $\mu'$ , some fraction of the time EUS will recommend to wait at time 0 when  $\mu'$  happens to be large. EUS earns higher expected utilities for such cases, resulting in higher relative performance of EUS. However, for large errors, the noise takes over and the performance of EUS decreases.

In contrast, when errors involve estimating the standard deviation of the price decreases, all

three strategies perform almost identically as if the forecasts are perfect. We found that the recommendations made by the strategies were the same as the ones made under perfect forecasts in many cases. This suggests that the strategies are tolerant to this form of estimation errors.

## 7 Concluding Remarks

We addressed the question of how to design a system that helps forward-looking consumers decide *when* to buy durable products. We proposed a utility-based framework that discounts the value of the product if its purchase is deferred and incorporates the savings accruing from the potential drop in the price of the product. Rather than designing forecasting algorithms for which a rich body of work already exists, we focused on developing recommendation strategies that employ forecasting algorithms as blackboxes. We devised three strategies as well as techniques for implementing them. These strategies differ in how they estimate the value of waiting. The Expected Utility Strategy (EUS) recommends a buy only if the maximum expected utility from buying the product in future is no better than the expected utility from buying it now. The Sample Path Strategy (SPS) recommends a buy only if the expected maximum utility of buying at some future time is no better than the expected utility of buying the product now. The difference between EUS and SPS is the order in which expectation and maximization is taken. Whereas under EUS, expectation is taken separately over each time index, and the time index with the highest utility is selected, expectation of the highest utility along an entire price path is taken under SPS. The Self Application Strategy (SAS) makes a buy recommendation only if the expected utility obtained from waiting now and following the strategy thereafter is no better than the expected utility of buying now.

We conducted a theoretical analysis of these strategies, and showed that the recommendations made by the three strategies are related—when SPS recommends to buy, so will SAS; when SAS recommends to buy, so will EUS. This is due to the estimated value of waiting is related through an inequality. We also established that SAS is optimal when the forecasts are perfect. We conducted an extensive experimental study to measure their performances both when forecasts are perfect and imperfect, and under different assumptions about the persistence of price shocks. The model parameters were derived from real data obtained from a data provider for durable goods. Based on the experimental study, we conclude that the optimal strategy SAS can improve the utility of consumers under most conditions. The improvements are especially significant when the standard deviations of price changes are high, or when the value the consumer placed on the product is close to the initial price. The present work complements our past study of price prediction for consumer durables [1]. Together they suggest that across a wide range of conditions, the proposed system can deliver significant savings to prospective users.

This work opens up several research directions. First, it is a challenging and interesting problem to establish an approximation bound on the proposed SAS evaluation (Algorithm 1). Second, we have assumed that the user has already selected which product to buy. A natural extension is when a user is interested in one of a few products (e.g., different models of TVs), and needs to decide which one to purchase as a function of time. Another question is how to improve user experience of product search through ranking functions that factor in whether it is a good time to buy the product.

## References

- [1] R. Agrawal, S. Ieong, and R. Velu. Ameliorating buyer’s remorse. In *Proc. KDD*, 2011.
- [2] G. Box, G. M. Jenkins, and G. Reinsel. *Time Series Analysis: Forecasting and Control*. Prentice Hall, 1994.
- [3] R. Burke. Hybrid recommender systems: Survey and experiments. *User Model. & User-Adapted Inter.*, 12(4), 2002.
- [4] Y. S. Chow, H. Robbins, and S. D. *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin Co, 1971.
- [5] W. H. Cummings and M. Venkatesan. Cognitive dissonance and consumer behavior: A review of the evidence. In M. J. Schlinger, editor, *Advances in Consumer Research Volume 2*, pages 21–32. Association for Consumer Research, 1975.
- [6] R. N. Dino. Forecasting the price evolution of new electronic products. *J. Forecasting*, 4:39–60, 1985.
- [7] R. J. Dolan and A. P. Jeuland. Experience curves and dynamic demand models: Implications for optimal pricing strategies. *J. Marketing*, 45:52–62, 1981.
- [8] O. Etzioni, R. Tuchinda, C. A. Knoblock, and A. Yates. To buy or not to buy: mining airfare data to minimize ticket purchase price. In *Proc. KDD*, pages 119–128, 2003.
- [9] I. Fisher. *Theory of Interest*. Augustus M. Kelly, 1930.
- [10] G. Fulgoni. State of the U.S. retail economy in Q2 2009. Technical report, Comscore, August 2009.
- [11] C. Gonzales and P. Perny. GAI networks for utility elicitation. In *Proc. 9th Int. Conf. Principles of Knowledge Representation and Reasoning*, 2004.
- [12] J. C. Hull. *Options, Futures, and Other Derivatives*. Prentice Hall, 5th edition, 2002.
- [13] R. J. Hyndman, A. B. Koehler, J. K. Ord, and R. D. Snyder. *Forecasting with Exponential Smoothing*. Springer, 2008.
- [14] D. Jannach, M. Zanker, A. Felfering, and G. Friedrich. *Recommender Systems: An Introduction*. Cambridge University Press, 2010.
- [15] S. S. Oren. Comments on “pricing research in marketing: The state of the art”. *J. Business*, 57:61–64, 1984.
- [16] P. Pazzani and D. Billsus. Content-based recommendation systems. In *The Adaptive Web*, 2007.
- [17] D. Prelec and G. Loewenstein. Beyond time discounting. *Marketing Letters*, 8:97–108, 1997.
- [18] V. R. Rao. Pricing research in marketing: The state of the art. *J. Business*, 57:39–60, 1984.

- [19] B. Sarwar, G. Karypis, J. Konstan, and J. Reidl. Item-based collaborative filtering recommendation algorithms. In *WWW*, 2001.
- [20] R. S. Winer. Discounting and its impact on durables buying decision. *Marketing Letters*, 8:109–118, 1997.

## A Proofs

We make use of the following lemmas to prove Theorem 2.

**Lemma 1** (Jensen’s Inequality). *For any concave function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$ ,  $\mathbb{E}[h(\mathbf{x})] \geq h(\mathbb{E}[\mathbf{x}])$ .*

Note in particular that  $\max$  is a concave function.

**Lemma 2.** *For a distribution  $E$  over future prices  $\mathbf{f}_s$ , and any function  $h : \mathbb{R}^{|\mathbf{f}_s|} \rightarrow \mathbb{R}$ ,*

$$\mathbb{E}_E[\mathbb{E}_{E|p_{s+1}}[h(p_{s+1}, \mathbf{f}_{s+1})]] = \mathbb{E}_E[h(\mathbf{f}_s)]$$

*Proof of Lemma 2.* This is a consequence of the Law of Total Expectation, which states that  $\mathbb{E}[\mathbb{E}[X | Y]] = \mathbb{E}[X]$ .  $\square$

*Proof of Theorem 2.* First, observe that for history  $\mathbf{h}_T$ , all three strategies must recommend buy, hence the theorem is true for all  $\mathbf{h}_T$ .

As mentioned in Section 5, the focus of this proof is on establishing the following relationship between the values of waiting estimated by the three strategies.

$$v_{\text{EUS}}(s, E, \theta, \delta) \leq v_{\text{SAS}}(s, E, \theta, \delta) \leq v_{\text{SPS}}(s, E, \theta, \delta) .$$

**Part 1.** We show that for all history  $\mathbf{h}_s$ , any distribution  $E$ ,  $v_{\text{EUS}}(s, E, \theta, \delta) \leq v_{\text{SAS}}(s, E, \theta, \delta)$ . We prove this by induction on  $s$ , starting from  $(T - 1)$  down to 0, hence proving it true for all history. We start by writing the function  $v_{\text{SAS}}$  in a more convenient form that exposes its recursive nature.

$$v_{\text{SAS}}(s, E, \theta, \delta) = \mathbb{E}_E[w(p_{s+1}, E|p_{s+1}, \delta, \theta, s + 1)] \tag{10}$$

$$= \mathbb{E}_E \left[ \max \left\{ \begin{array}{l} \theta\delta^{s+1} - p_{s+1} \\ \mathbb{E}_{E|p_{s+1}}[w(p_{s+2}, E|p_{s+2}, \delta, \theta, s + 2)] \end{array} \right. \right]$$

$$= \mathbb{E}_E \left[ \max \left\{ \begin{array}{l} \theta\delta^{s+1} - p_{s+1} \\ v_{\text{SAS}}(s + 1, E|p_{s+1}, \theta, \delta) \end{array} \right. \right] \tag{11}$$

where the last step in Eq (11) is based on the earlier definition of  $v_{\text{SAS}}$  in Eq (10).

*Base case.* When  $s = T - 1$ , for any history  $\mathbf{h}_s$ ,

$$v_{\text{EUS}}(s, E, \theta, \delta) = v_{\text{SAS}}(s, E, \theta, \delta) = \mathbb{E}_E[\theta\delta^T - p_T] ,$$

and the hypothesis is true. Note that this is analogous to the case proved in Theorem 1.

*Inductive case.* Suppose that the hypothesis is true for  $s = k + 1$ . We show that the hypothesis is true for  $s = k$ . For any history  $\mathbf{h}_s$ , we can express  $v_{\text{EUS}}(s, E, \theta, \delta)$  in a similar form to Eq (11).

$$v_{\text{EUS}}(s, E, \theta, \delta) = \max \left\{ \begin{array}{l} \mathbb{E}_E[\theta\delta^{s+1} - p_{s+1}] \\ \max_{s+1 < t \leq T} \mathbb{E}_E[\theta\delta^t - p_t] \end{array} \right.$$

From Eq (11), by Jensen's Inequality,

$$v_{\text{SAS}}(s, E, \theta, \delta) \geq \max \left\{ \begin{array}{l} \mathbb{E}_E[\theta\delta^{s+1} - p_{s+1}] \\ \mathbb{E}_E[v_{\text{SAS}}(s+1, E|p_{s+1}, \theta, \delta)] \end{array} \right.$$

Focusing on the second term in the maximum,

$$\begin{aligned} & \mathbb{E}_E[v_{\text{SAS}}(s+1, E|p_{s+1}, \theta, \delta)] \\ & \geq \mathbb{E}_E[v_{\text{EUS}}(s+1, E|p_{s+1}, \theta, \delta)] && \text{(Ind. hyp.)} \\ & = \mathbb{E}_E\left[\max_{s+1 < t \leq T} \mathbb{E}_{E|p_{s+1}}[\theta\delta^t - p_t]\right] && \text{(Eq (3))} \\ & \geq \max_{s+1 < t \leq T} \mathbb{E}_E[\mathbb{E}_{E|p_{s+1}}[\theta\delta^t - p_t]] && \text{(Jensen's)} \\ & = \max_{s+1 < t \leq T} \mathbb{E}_E[\theta\delta^t - p_t] && \text{(Lemma 2)} \end{aligned}$$

**Part 2.** For all history  $\mathbf{h}_s$ , any distribution  $E$ ,  $v_{\text{SAS}}(s, E, \theta, \delta) \leq v_{\text{SPS}}(s, E, \theta, \delta)$ . We prove this by induction on  $s$ .

*Base case.* When  $s = T - 1$ , for any history  $\mathbf{h}_s$ ,

$$v_{\text{SAS}}(s, E, \theta, \delta) = v_{\text{SPS}}(s, E, \theta, \delta) = \mathbb{E}_{f(\mathbf{h}_s)}[\theta\delta^T - p_T] .$$

*Inductive case.* Suppose that the hypothesis is true for  $s = k + 1$ . We show that the hypothesis is true for  $s = k$ . For any history  $\mathbf{h}_s$ , applying Lemma 2,

$$\begin{aligned} v_{\text{SPS}}(s, E, \theta, \delta) &= \mathbb{E}_E \left[ \max \left\{ \begin{array}{l} \theta\delta^{s+1} - p_{s+1} \\ \max_{s+1 < t \leq T} \theta\delta^t - p_t \end{array} \right. \right] \\ &= \mathbb{E}_E \left[ \mathbb{E}_{E|p_{s+1}} \left[ \max \left\{ \begin{array}{l} \theta\delta^{s+1} - p_{s+1} \\ \max_{s+1 < t \leq T} \theta\delta^t - p_t \end{array} \right. \right] \right] \end{aligned}$$

and  $v_{\text{SAS}}(s, E, \theta, \delta) =$

$$\mathbb{E}_E \left[ \mathbb{E}_{E|p_{s+1}} \left[ \max \left\{ \begin{array}{l} \text{term}_1 = \theta\delta^{s+1} - p_{s+1} \\ \text{term}_2 = v_{\text{SAS}}(s+1, E|p_{s+1}, \theta, \delta) \end{array} \right. \right] \right] \right]$$

The inner expectation, conditional on price  $p_{s+1}$ , is a function of  $p_{s+1}$ . If  $\text{term}_1 \geq \text{term}_2$ , then the corresponding inner expectation in  $v_{\text{SPS}}$  must be at least as large. Otherwise,

$$\begin{aligned} & \mathbb{E}_{E|p_{s+1}}[v_{\text{SAS}}(s+1, E|p_{s+1}, \theta, \delta)] \\ & = v_{\text{SAS}}(s+1, E|p_{s+1}, \theta, \delta) && (*) \\ & \leq v_{\text{SPS}}(s+1, E|p_{s+1}, \theta, \delta) && \text{(Ind. hyp.)} \\ & = \mathbb{E}_{E|p_{s+1}}\left[\max_{s+1 < t \leq T} \theta\delta^t - p_t\right] \end{aligned}$$

where (\*) follows by the fact that once  $p_{s+1}$  is fixed,  $v_{\text{SAS}}$  is a scalar value. Finally, note that this last inequality shows that the inner expectation of  $v_{\text{SPS}}$  must be as large for the case when  $\text{term}_1 < \text{term}_2$ .  $\square$

*Proof of Theorem 3.* First, without loss of generality, we assume that the choices of strategy  $r$  are deterministic, i.e., for each history  $\mathbf{h}_t$ , it will output either 0 or 1 all the time. This is because the performance of any randomized strategy can be expressed as a convex combination of the performances of deterministic strategies, and hence the best among these must have a performance at least as high as that of the randomized strategy.

Second, recall the definition of  $w$  in Eqs (7) and (8),

$$w(p, t, E, \theta, \delta) = \max \begin{cases} \theta\delta^t - p \\ \mathbb{E}_E[w(p_{t+1}, t+1, E|p_{t+1}, \theta, \delta)] \end{cases}$$

and for time  $t = T$ ,

$$w(p, T, E, \theta, \delta) = \theta\delta^T - p .$$

As noted in Remark 1,  $\text{SAS}(\mathbf{h}_s, E, \theta, \delta)$  depends on  $\mathbf{h}_s$  only through  $p_s$  and  $|\mathbf{h}_s|$ . To highlight this dependence and its relationship to  $w$ , we define  $\text{SAS}'$  as a function that satisfies  $\text{SAS}(\mathbf{h}_s, E, \theta, \delta) = \text{SAS}'(p_s, |\mathbf{h}_s|, E, \theta, \delta)$ . Note that the information to determine the value of the LHS and the RHS of Eq (6), the criteria for recommendations under SAS.

Third, as mentioned after Definition 8, there is some kind of a duality between distributions and consistent forecasting algorithm. Recall the notation  $[E]$  that denotes the forecasting algorithm induced by distribution  $E$  over price paths. This usage of a distribution as a forecasting algorithm becomes important in the proof when working with the purchase time function without a forecasting algorithm. In the definition of purchase time, the forecasting algorithm is always applied to the first element of the first argument. Therefore, for a consistent forecasting algorithm  $F$ , to determine the value of  $\tau(\mathbf{p}, \theta, \delta, F, r)$ , we only need to know the distribution of  $F(p_0)$ . We define a purchase time function  $\lambda$  that operates on distribution as

$$\lambda(\mathbf{p}, \theta, \delta, E, r) = \min_{0 \leq t \leq |\mathbf{p}|} \{r(\mathbf{h}_t, E|(p_1, \dots, p_t), \delta, \theta)\}$$

Note that for any price path  $\mathbf{p}$ , for consistent  $F$ ,

$$\tau(\mathbf{p}, \theta, \delta, F, r) = \lambda(\mathbf{p}, \theta, \delta, F(p_0), r) \tag{12}$$

$$\tau(\mathbf{f}_0, \theta, \delta, [F(p_0)], r) = \lambda(\mathbf{f}_0, \theta, \delta, F(p_0, p_1), r) \tag{13}$$

To prove the hypothesis, we need a number of lemmas. The first is related to a change in the length of the price path argument to the utility function.

**Lemma 3.** *For any  $\delta, \theta$ , if  $t > s \geq 0$ ,*

$$u(t, (\mathbf{h}_s, \mathbf{f}_s), \theta, \delta) = u(t - s, \mathbf{f}_s, \theta\delta^s, \delta) .$$

*Proof of Lemma 3.* Direct from the definition of utility. □

The second is related to what forecasting algorithm is perfect for a distribution  $D'$  of price paths of length  $T$  is generated from conditioning on distribution  $D$ .

**Lemma 4.** *If forecasting algorithm  $F$  is perfect for a distribution  $D$  of price paths of length  $T + 1$ , and  $D'$  is defined as a distribution of price paths of length  $T$ , where for all  $\mathbf{f}_T$ ,*

$$\mathbb{P}_{D'}(\mathbf{f}_T) = \mathbb{P}_D((p_0, \mathbf{f}_T) \mid p_0) ,$$

*then  $[F(p_0)]$  is perfect for  $D'$ .*

*Proof of Lemma 3.* Direct from Definition 9. □

The third is related to a memoryless property of purchase time under recommendation strategy SAS.

**Lemma 5.** *Let  $E$  be a distribution over future prices  $\mathbf{f}_s$ , and treat it as a forecasting algorithm. Let  $p$  be an arbitrary price. If  $\tau((p, \mathbf{f}_s), \delta, \theta, E, SAS) > 1$ , then*

$$\lambda((p, \mathbf{f}_s), \theta, \delta, E, SAS) = \lambda(\mathbf{f}_s, \theta\delta, \delta, E|p_{s+1}, SAS) + 1$$

To improve notational clarity, Let  $\mathbf{f}_{t,t'}$  denote the future prices between  $t$  and  $t'$ , i.e.,  $\mathbf{f}_{t,t'} = (p_{t+1}, p_{t+2}, \dots, p_{t'})$ .

*Proof of Lemma 5.* First, using the auxiliary function SAS', note that for history  $\mathbf{h}_s$  where  $|\mathbf{h}_s| > 0$ , for any distribution  $E$ ,

$$\begin{aligned} \text{SAS}(\mathbf{h}_s, E, \theta, \delta) &= \text{SAS}'(p_s, |\mathbf{h}_s|, E, \theta, \delta) \\ &= \text{SAS}'(p_s, |\mathbf{h}_s| - 1, E, \theta\delta, \delta) = \text{SAS}(\mathbf{h}_{1,s}, E, \theta\delta, \delta) \end{aligned}$$

If  $\tau((p, \mathbf{f}_s), \theta, \delta, E, SAS) > 1$ ,

$$\begin{aligned} &\lambda((p, \mathbf{f}_s), \theta, \delta, E, SAS) \\ &= \min_{1 \leq t \leq |\mathbf{f}_s|+1} \{\text{SAS}((p, \mathbf{f}_{s,s+t}), E|(\mathbf{f}_{s,s+t}), \theta, \delta) = 1\} \\ &= \min_{0 \leq t \leq |\mathbf{f}_s|} \{\text{SAS}(\mathbf{f}_{s,s+t}, (E|p_{s+1})|\mathbf{f}_{s+1,s+t}), \theta\delta, \delta) = 1\} \\ &= \lambda(\mathbf{f}_s, \theta\delta, \delta, E|p_{s+1}, SAS) \end{aligned} \quad \square$$

The next lemma is the key ‘‘lemma’’ that was stated in the main text, but now stated formally using  $\lambda$  rather than  $\tau$ , as it does not explicitly refer to a forecasting algorithm. It shows that the value  $v_{\text{SAS}}$  is exactly equal to the expected utility that SAS obtains conditional on history  $\mathbf{h}_t$ .

**Lemma 6.** *For any  $\delta, \theta, \mathbf{h}_s$  for  $s \in [0, T)$ ,*

$$v_{\text{SAS}}(s, F(\mathbf{h}_s), \theta, \delta) = \mathbb{E}_{F(\mathbf{h}_s)}[u(\lambda(\mathbf{f}_s, \theta\delta^{s+1}, \delta, F(\mathbf{h}_s)|p_{s+1}, SAS), \mathbf{f}_s, \theta\delta^{s+1}, \delta)]$$

*Proof of Lemma 6.* We show this by induction on  $s$ , starting from  $T - 1$  down to 0.

*Base case.* When  $s = T - 1$ , for any history  $\mathbf{h}_{T-1}$ , since all recommendation strategies must purchase at time  $T$ ,

$$\text{RHS} = \mathbb{E}_{F(\mathbf{h}_{T-1})}[u(0, \mathbf{f}_{T-1}, \theta\delta^T, \delta)] = \mathbb{E}_{F(\mathbf{h}_{T-1})}[\theta\delta^T - p_T] .$$



By the terminal condition of SAS given in Eq (8)

$$\text{LHS} = \mathbb{E}_{F(\mathbf{h}_{T-1})}[\theta\delta^T - p_T] = \text{RHS} .$$

*Inductive case.* Suppose the hypothesis is true for all history  $\mathbf{h}_{s+1}$ . We show that it is true for any history  $\mathbf{h}_s$ . We follow a path similar to Part 2 of Proof of Theorem 2. Recall that for consistent  $F$ ,  $F(\mathbf{h}_s)|_{p_{s+1}} = F(\mathbf{h}_{s+1})$ . We write LHS as

$$v_{\text{SAS}}(\theta, s, \delta, F(\mathbf{h}_s)) = \mathbb{E}_{F(\mathbf{h}_s)}[\mathbb{E}_{F(\mathbf{h}_s)|_{p_{s+1}}} [ \max \left\{ \begin{array}{l} \text{term}_1 = \theta\delta^{s+1} - p_{s+1} \\ \text{term}_2 = v_{\text{SAS}}(s+1, F(\mathbf{h}_s)|_{p_{s+1}}, \theta, \delta) \end{array} \right\} ]]$$

Likewise, by Lemma 2, we write RHS as

$$\mathbb{E}_{F(\mathbf{h}_s)}[\mathbb{E}_{F(\mathbf{h}_s)|_{p_{s+1}}} [ u(\lambda(\mathbf{f}_s, \theta\delta^{s+1}, \delta, F(\mathbf{h}_s)|_{p_{s+1}}, \text{SAS}), \mathbf{f}_s, \theta\delta^{s+1}, \delta) ]]$$

We focus on the inner conditional expectation, which is a function of  $p_{s+1}$ , bounding it per term for each value of  $p_{s+1}$ .

If  $\text{term}_1 \geq \text{term}_2$ , then at time  $(s+1)$ , SAS recommends buying, and  $\lambda(\mathbf{f}_s, \theta\delta^{s+1}, \delta, F(\mathbf{h}_s)|_{p_{s+1}}, \text{SAS}) = 0$ , hence

$$\mathbb{E}_{F(\mathbf{h}_s)|_{p_{s+1}}} [u(0, \mathbf{f}_s, \theta\delta^{s+1}, \delta)] = \theta\delta^{s+1} - p_{s+1} = \text{term}_1$$

If  $\text{term}_1 < \text{term}_2$ , SAS recommends waiting at time  $(s+1)$ , and  $\lambda(\mathbf{f}_s, \theta\delta^{s+1}, \delta, F(\mathbf{h}_s)|_{p_{s+1}}, \text{SAS}) > 0$ . Focusing on the RHS, and noting that by consistency of  $F$ ,  $F(\mathbf{h}_s)|_{p_{s+1}} = F(\mathbf{h}_{s+1})$ ,

$$\begin{aligned} \alpha &\stackrel{\text{def}}{=} u(\lambda(\mathbf{f}_s, \theta\delta^{s+1}, \delta, F(\mathbf{h}_s)|_{p_{s+1}}, \text{SAS}), \mathbf{f}_s, \theta\delta^{s+1}, \delta) \\ &\stackrel{(1)}{=} u(\lambda(\mathbf{f}_{s+1}, \theta\delta^{s+2}, \delta, (F(\mathbf{h}_s)|_{p_{s+1}})|(p_{s+2}), \text{SAS}) + 1, \mathbf{f}_s, \theta\delta^{s+1}, \delta) \\ &\stackrel{(2)}{=} u(\lambda(\mathbf{f}_{s+1}, \theta\delta^{s+2}, \delta, F(\mathbf{h}_{s+1})|_{p_{s+2}}, \text{SAS}), \mathbf{f}_{s+1}, \theta\delta^{s+2}, \delta) \\ &\Rightarrow \mathbb{E}_{F(\mathbf{h}_s)|_{p_{s+1}}} [\alpha] \stackrel{(3)}{=} v_{\text{SAS}}(s+1, F(\mathbf{h}_{s+1}), \theta, \delta) \end{aligned}$$

where (1) is by Lemma 5, (2) is by Lemma 3, and (3) is by the inductive hypothesis. Since all the inner conditional expectation agrees in value, the hypothesis holds.  $\square$

Using these lemmas, we now establish the hypothesis by induction. For any  $k > 0$ , we show that it holds for any distribution  $D$  of price paths of length  $k$ .

*Base case.* For  $k = 1$ , distribution  $D$  is over price paths of length 1, i.e, paths of the form  $(p_0)$ , in other words,  $T = 0$ . Since all recommendation strategies must buy at time  $T$ , they all have the same performance, and the hypothesis holds.

*Inductive case.* Suppose the hypothesis is true for any distribution  $D$  of price paths of length  $k$ . We now prove that it is also true for any distribution  $D$  of price paths of length  $k + 1$ , in other words,  $T = k$ .

We group the performance of any recommendation strategy  $r$  by the first price  $p_0$  using Lemma 2.

$$\begin{aligned} P(F, r, D, \theta, \delta) &= \mathbb{E}_D[u(\tau(\mathbf{p}, \theta, \delta, F, r), \mathbf{p}, \theta, \delta)] \\ &= \mathbb{E}_D[\mathbb{E}_{D|p_0}[u(\tau(\mathbf{p}, \theta, \delta, F, r), \mathbf{p}, \theta, \delta)]] \stackrel{\text{def}}{=} \mathbb{E}_D[h(p_0)] \end{aligned}$$

where we use  $h(p_0)$  to denote the inner conditional expectation, which is a function of  $p_0$ . If  $\tau(\mathbf{p}, \theta, \delta, F, r) = 0$ , then

$$\begin{aligned} h(p_0) &= \mathbb{E}_{D|p_0}[u(\tau(\mathbf{p}, \theta, \delta, F, r), \mathbf{p}, \theta, \delta)] = \theta - p_0 \\ &\leq w(p_0, 0, F(p_0), \theta, \delta) \end{aligned}$$

If  $\tau(\mathbf{p}, \theta, \delta, F, r) > 0$ ,

$$\begin{aligned} h(p_0) &= \mathbb{E}_{D|p_0}[u(\tau(\mathbf{p}, \theta, \delta, F, r), \mathbf{p}, \theta, \delta)] \\ &= \mathbb{E}_{D|p_0}[u(\lambda(\mathbf{p}, \theta, \delta, F(p_0), r), \mathbf{p}, \theta, \delta)] && \text{(Eq (12))} \\ &= \mathbb{E}_{D|p_0}[u(\lambda(\mathbf{p}, \theta, \delta, F(p_0), r) - 1, \mathbf{f}_0, \theta\delta, \delta)] && \text{(Lemma 3)} \end{aligned}$$

Now consider a problem instance where the distribution of price paths is given by  $D|p_0$ , value is  $\theta\delta$ , and discount rate is  $\delta$ . Note that  $D|p_0$  is a distribution over paths of length  $k$ . Let recommendation strategy  $r'$  be defined such that for all price paths  $\mathbf{x}$  of length at most  $T - 1$  and any distribution  $E$ ,

$$r'(\mathbf{x}, E, \delta, \theta\delta) = r((p_0, \mathbf{x}), [F(p_0)], \delta, \theta)$$

Then for all  $\mathbf{f}_0$ , by Eq (13),

$$\begin{aligned} \tau(\mathbf{f}_0, \delta, \theta, [F(p_0)], r') &= \lambda(\mathbf{f}_0, \delta, \theta, F(p_0, p_1), r') \\ &= \lambda((p_0, \mathbf{f}_0), \delta, \theta, F(p_0), r) - 1 \end{aligned}$$

And hence  $h(p_0)$  can be interpreted as the performance of  $r'$  in the new problem. Note that by Lemma 4,  $[F(p_0)]$  is a perfect forecasting algorithm for  $D|p_0$ . Hence,

$$\begin{aligned} h(p_0) &= P([F(p_0)], r', D|p_0, \theta\delta, \delta) \\ &\leq P([F(p_0)], \text{SAS}, D|p_0, \theta\delta, \delta) && \text{(Ind. hyp.)} \\ &= \mathbb{E}_{D|p_0}[u(\tau(\mathbf{f}_0, \theta\delta, \delta, [F(p_0)], \text{SAS}), \mathbf{f}_0, \theta\delta, \delta)] \\ &= \mathbb{E}_{D|p_0}[u(\lambda(\mathbf{f}_0, \theta\delta, \delta, F(p_0, p_1), \text{SAS}), \mathbf{f}_0, \theta\delta, \delta)] && \text{(Eq (13))} \\ &= v_{\text{SAS}}(0, F(p_0), \theta, \delta) && \text{(Lemma 6)} \\ &\leq w(p_0, 0, F(p_0), \theta, \delta) \end{aligned}$$

Now consider the performance of SAS on  $D$ ,

$$\begin{aligned} P(F, \text{SAS}, D, \theta, \delta) &= \mathbb{E}_D[u(\tau(\mathbf{p}, \theta, \delta, F, \text{SAS}), \mathbf{p}, \theta, \delta)] \\ &= \mathbb{E}_D[\mathbb{E}_{D|p_0}[u(\tau(\mathbf{p}, \theta, \delta, F, \text{SAS}), \mathbf{p}, \theta, \delta)]] \stackrel{\text{def}}{=} \mathbb{E}_D[h'(p_0)] \end{aligned}$$

If  $\tau(\mathbf{p}, \theta, \delta, F, \text{SAS}) = 0$ , then SAS recommends buying at the time 0, hence

$$\begin{aligned} w(p_0, 0, F(p_0), \theta, \delta) &= \max\{\theta - p_0, v_{\text{SAS}}(0, F(p_0), \theta, \delta)\} \\ &= \theta - p_0 = h'(p_0) \end{aligned}$$

If  $\tau(\mathbf{p}, \theta, \delta, F, \text{SAS}) > 0$ , then SAS recommends to wait at time 0, hence

$$\begin{aligned}
h'(p_0) &= \mathbb{E}_{D|p_0}[u(\tau(\mathbf{p}, \theta, \delta, F, \text{SAS}), \mathbf{p}, \theta, \delta)] \\
&= \mathbb{E}_{D|p_0}[u(\lambda(\mathbf{p}, \theta, \delta, F(p_0), \text{SAS}), \mathbf{p}, \theta, \delta)] && \text{(Eq (12))} \\
&= \mathbb{E}_{F(p_0)}[u(\lambda(\mathbf{f}_0, \theta\delta, \delta, F(p_0)|p_1, \text{SAS}), \mathbf{f}_0, \theta\delta, \delta)] && \text{(Lemma 5)} \\
&= v_{\text{SAS}}(0, F(p_0), \theta, \delta) && \text{(Lemma 6)} \\
&= w(p_0, 0, F(p_0), \theta, \delta)
\end{aligned}$$

where the last equality follows since SAS recommends to wait at time 0. Putting it together,

$$\begin{aligned}
P(F, \text{SAS}, D, \theta, \delta) &= \mathbb{E}_D[w(p_0, 0, F(p_0), \theta, \delta)] \\
&\geq \mathbb{E}_D[h(p_0)] = P(F, r, D, \theta, \delta)
\end{aligned}$$

□