



Expected-Utility-Maximizing Price Search with Learning

Ronald M. Harstad; Andrew Postlewaite

Management Science, Volume 27, Issue 1 (Jan., 1981), 75-80.

Stable URL:

<http://links.jstor.org/sici?sici=0025-1909%28198101%2927%3A1%3C75%3AEP%3E2.0.CO%3B2-Z>

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

Management Science is published by INFORMS. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/informs.html>.

Management Science

©1981 INFORMS

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2003 JSTOR

EXPECTED-UTILITY-MAXIMIZING PRICE SEARCH WITH LEARNING*

RONALD M. HARSTAD[†] AND ANDREW POSTLEWAITE[‡]

We consider a model of search when the distribution of prices (wages) is unknown. The effect of changing the objective function from minimizing expected cost to maximizing expected utility is examined.

(DECISION ANALYSIS—SEQUENTIAL; DYNAMIC PROGRAMMING—BAYESIAN;
UTILITY/PREFERENCE—APPLICATIONS)

1. Introduction

A considerable literature¹ analyzes the behavior of an individual who is to buy one unit of a good in a market with differing prices for the good. The consumer is assumed to sample, at a fixed cost, from a known probability distribution over a finite set of potential prices, and to minimize the expected sum of price plus search costs. The optimal strategy is characterized by a reservation price, with search ceasing as soon as a price quotation no higher than the reservation price is elicited.

Rothschild [3] is rightly critical of the assumption that the searcher behaves as if he knew the distribution of prices, and looks instead at an environment where a searcher with prior beliefs about the unknown price distribution “learns” through Bayesian updating of beliefs as prices are observed.

For this environment, the characterization which would be analogous to a reservation price is the reservation-price property: the existence of a threshold price function which determines a reservation price for each history of observed price offers. Rothschild establishes the reservation-price property for an expected-cost minimizing search strategy, in the case where the searcher’s prior beliefs have a Dirichlet distribution (cf. below). Bayesian updating with a Dirichlet prior has the crucial neutrality property; observing price p_i does not affect the relative perceived likelihood of prices p_j and p_k , $j, k \neq i$.

This paper considers the extension to an expected-utility-maximizing searcher with the option of purchasing different quantities of the commodity. In particular, we consider a searcher who derives utility from leisure and a consumption good. A price quotation for the good can be obtained by foregoing a fixed amount of leisure. After any price is observed, the searcher may exchange any (or all) of his (or her) remaining leisure for the consumption good at this price. The quantity ultimately purchased will in general depend upon the prices he has observed.

In the case where the commodity in question is a durable good, and the quantity purchased is invariant with respect to the price, we might expect that the quality of the good is not invariant. If so, the quality may be treated as a proxy for quantity

* Accepted by Bennett L. Fox, former Departmental Editor; received August 14, 1979. This paper has been with the authors 6 months for 1 revision.

[†] University of British Columbia and University of Illinois at Urbana-Champaign.

[‡] Princeton University and University of Illinois at Urbana-Champaign.

¹ For a bibliography, see [2] or [3].

purchased. Thus even for durable consumer goods expected-cost minimization may yield different results from expected-utility maximization.

We show that for expected-utility maximization the reservation-price property may not hold. This is demonstrated by means of a simple non-pathological example where prior beliefs have a Dirichlet distribution, as in Rothschild.

We also show that if we assume a searcher's indirect utility function globally satisfies a particular inequality, then his expected-utility-maximizing stopping rule will exhibit the reservation-price property if he has a Dirichlet prior. The inequality that guarantees the reservation-price property is satisfied for all direct utility functions which are homogeneous.

Searchers are modelled here as solving a dynamic programming problem. Where search is observed in the economy, less complex rules of thumb may be followed, though some rules of thumb are presumably discarded when found to deviate too far from optimal behavior.

2. Preferences and Beliefs

A searcher faces one market where the numeraire good, leisure, is exchanged for a consumption good. A price quotation can be elicited by foregoing c units of leisure.

Suppose the individual has decided to cease searching upon observing the (relative) price p . The quantity to purchase is then determined by the standard (full-information) consumer problem,²

$$u(p, I) = \max_{x, L} w(x, L) \quad \text{subject to } I - L - px \leq 0 \quad (1)$$

where $w(\cdot) : R_+^2 \rightarrow R^1$ is the direct utility function, (x, L) the vector of consumption of the purchased good and leisure, and I the "current income" or the finite leisure endowment \bar{I} less leisure foregone for search. w is assumed to be concave and continuous. As properties of (1) have been well studied, x , L and $w(\cdot)$ will be submersed, and the analysis concerned with the indirect utility function $u(p, I)$ and the decision whether to search again with current income, and quoted price, given updated beliefs.

The environment is the formal extension of Rothschild's model to expected-utility maximization, altered only to maintain consistency with that objective. Search is without recall, eliciting each time an element of the finite set

$$P = \{p_1, p_2, \dots, p_n\}$$

labelled ascendingly, so $p_i < p_{i+1}$.³

²The searcher presumably has the option of purchasing some quantity at the elicited price, but continuing to search. (Whether to search again is a binary decision, how much to purchase is not; this is presumably irrelevant to housing, automobile and similar markets.) If, for example, the initial quotation was perceived to be moderate, the probability of lower prices may be believed substantial enough to justify continued search. Yet the perceived probability of continually finding higher prices may be significant enough to justify for a risk-averse searcher nonzero purchase at the initial price quoted. We must clearly make the realistic assumption that the good cannot be sold for leisure at any quoted price, to avoid an unusual type of speculation.

For the examples and characterizations below, we may, without loss of generality, ignore this option.

³If the set of potential price quotations is bounded, finiteness seems a quite tenable assumption. (Price quotations involving fractions of a penny are rare.) However, a large or complete potential price set (all integral dollar amounts between 300 and 500, say) suggests neutral updating of beliefs is a particularly heroic assumption.

The actual (but unknown) probability distribution over P , a multinomial, is characterized by $\Pi = (\pi_1, \dots, \pi_n) \in \Delta$, where Δ is the simplex in R_+^n , π_i the probability that any given price quotation will be p_i . The searcher updates by Bayes' rule a subjective prior distribution over Δ as price quotations are obtained. If this experience is summarized by $N = (N_1, \dots, N_n)$, the number of times each price has been observed, let $\rho = 1/\sum N_i$ and $\mu_i = \rho N_i$, $i = 1, \dots, n$. Then

$$(\mu, \rho) \in \Gamma = \Delta \times [0, 1]$$

completely characterizes experience. With an observation of p_i , (μ, ρ) becomes

$$h_i(\mu, \rho) = \left(\frac{\mu_1}{1 + \rho}, \dots, \frac{\mu_{i-1}}{1 + \rho}, \frac{\mu_i + \rho}{1 + \rho}, \frac{\mu_{i+1}}{1 + \rho}, \dots, \frac{\mu_n}{1 + \rho}; \frac{\rho}{1 + \rho} \right). \quad (2)$$

Let $\lambda(\mu, \rho)$ be the current prior over P for a searcher with history (μ, ρ) . Note that

$$\lim_{\rho \rightarrow 0} \lambda_i(\mu, \rho) = \mu_i, \quad i = 1, \dots, n. \quad (3)$$

We maintain the assumption that the searcher's prior is a Dirichlet distribution (in which case updated beliefs are also Dirichlet), which has been completely characterized.

PROPOSITION 1 (ROTHSCHILD [3]). *A searcher has a Dirichlet prior if and only if it is possible to parametrize his experience so that*

$$\lambda_i(\mu, \rho) = \mu_i \quad \text{for all } (\mu, \rho) \in \Gamma.$$

By (3) above, then, all searchers come with enough experience to behave as if their priors were Dirichlet.

3. The Optimal Strategy

Let

$$V_1(\mu, \rho, I) = \sum_i \lambda_i(\mu, \rho) u(p_i, I). \quad (4)$$

$V_1(\mu, \rho, I)$ is the expected (indirect) utility for a consumer allowed to search only once, assuming history (μ, ρ) , and holding current income I (after costs of this one allowed search have been paid). Designate

$$V_T(\mu, \rho, I) = \sum_i \lambda_i(\mu, \rho) \max\{u(p_i, I), V_{T-1}[h_i(\mu, \rho), I - c]\} \quad (5)$$

$V_T(\mu, \rho, I)$ is then the expected attainable utility for an individual who can search at most T times, has prior experience (μ, ρ) , and income I after costs of the first search have been paid. Clearly

$$V_{T-1}(\mu, \rho, I) \leq V_T(\mu, \rho, I) \leq u(p_1, I)$$

for all T , so the $V_T(\mu, \rho, I)$ converge.

Defining

$$V(\mu, \rho, I) = \lim_{T \rightarrow \infty} V_T(\mu, \rho, I), \quad (6)$$

it is clear that

$$V(\mu, \rho, I) = \sum_i \lambda_i(\mu, \rho) \max\{u(p_i, I), V[h_i(\mu, \rho), I - c]\}, \quad (7)$$

so the optimal stopping policy is to cease search, upon eliciting price p_i with experience (μ, ρ) and income I , if $u(p_i, I) \geq V[h_i(\mu, \rho), I - c]$, otherwise to search again.

Rothschild [3] proves that the optimal strategy for an expected-cost-minimizing searcher implies a finite number of searches. His argument extends to this model, although finite income implies finite search trivially.

4. Non-Reservation-Price Examples

In general, searchers who learn about the unknown distribution of prices do not exhibit the reservation-price property. Rothschild's example imagines that $P = \{1, 2, 3\}$, $c = 0.01$ ($\bar{I} \geq 0.5$), and prior beliefs admit only two possible price distributions: (a) $\text{prob}(3) = 1$, or (b) $\text{prob}(1) = 0.99$, $\text{prob}(2) = 0.01$. Then an initial quote of 3 will end search, but an initial price of 2 will cause further search, as a much better price is expected.

In his example, a lower price offer has a direct effect, higher attainable utility if search ceases ($u(2, I) > u(3, I)$). It also has an indirect effect, measured in utility terms as $V[h_2(\mu, \rho), I - c] - V[h_3(\mu, \rho), I - c]$. In this case, the indirect effect of a lower price quote swamps the direct effect. When that cannot happen, the reservation-price property obtains.

PROPOSITION 2. *If for all $(\mu, \rho) \in \Delta$, all $I > 0$,*

$$|u(p_k, I) - u(p_i, I)| \geq |V[h_k(\mu, \rho), I - c] - V[h_i(\mu, \rho), I - c]|, \quad (8)$$

then

$$V[h_i(\mu, \rho), I - c] \leq u(p_i, I) \leq u(p_k, I) \quad (9)$$

implies

$$V[h_k(\mu, \rho), I - c] \leq u(p_k, I). \quad (10)$$

That (9) implies (10) is the reservation-price property. We adopt and maintain the notational convention $p_i > p_k$.

PROOF. Trivial if $V[h_k(\mu, \rho), I - c] \leq V[h_i(\mu, \rho), I - c]$. If not, by (8),

$$u(p_k, I) - u(p_i, I) \geq V[h_k(\mu, \rho), I - c] - V[h_i(\mu, \rho), I - c],$$

so

$$u(p_k, I) - V[h_k(\mu, \rho), I - c] \geq u(p_i, I) - V[h_i(\mu, \rho), I - c] \geq 0.$$

When the searcher's objective is to minimize expected cost, the indirect effect of a lower price quote is solely the informational value of that price (the expected-cost evaluation of the effect upon updated beliefs). In the case of a Dirichlet prior, the informational value of a lower price cannot swamp its direct effect, and Rothschild is able to prove that expected-cost-minimizing searchers with Dirichlet priors exhibit the reservation-price property.

However, the Dirichlet assumption is not sufficient to establish the reservation-price property for expected-utility-maximizing search, as now computing the indirect effect

of observing a lower price must take into account the expected change in the marginal utility of leisure. That is, while the nominal cost of search is fixed, the expected cost in foregone utility will depend on price observations.

Consider a simple illustration. Let

$$w(x, L) = \begin{cases} x + 40, & \text{if } x \geq \frac{2}{3}, \\ 122x - 40\frac{2}{3}, & \text{if } 0 \leq x \leq \frac{2}{3}, \end{cases}$$

$$P = \left\{ \frac{1}{320}, \frac{1}{4}, 3 \right\}, \text{ and}$$

$$(\mu, \rho) = \left(\frac{1}{8}, \frac{1}{8}, \frac{3}{4}; \frac{1}{4} \right)$$

is the searcher's experience. Search cost, c , is 1, and income after the first search is 2.

It is obvious that one or two searches will be optimal. If the first search elicits $p = 3$, the posterior is $(\mu, \rho) = (\frac{1}{10}, \frac{1}{10}, \frac{4}{5}; \frac{1}{5})$. A second search would yield expected utility

$$EU = \frac{1}{10} w(320, 0) + \frac{1}{10} w(4, 0) + \frac{8}{10} w(\frac{1}{3}, 0) = 40.4.$$

Without search $w(\frac{2}{3}, 0) = 40\frac{2}{3}$ can be achieved, so search ceases.

Suppose instead the first search observed $p = \frac{1}{4}$. The posterior is then $(\mu, \rho) = (\frac{1}{10}, \frac{3}{10}, \frac{3}{5}; \frac{1}{5})$. A second search would yield expected utility

$$EU = \frac{1}{10} w(320, 0) + \frac{3}{10} w(4, 0) + \frac{3}{5} w(\frac{1}{3}, 0) = 49.2,$$

while stopping after the first search can only attain $w(8, 0) = 48$.

So, in this example, if the first search elicits the highest price, 3, search stops and 2 units are purchased. But if the first price quoted is lower, $\frac{1}{4}$, expected utility is maximized by searching again. Note that allowing for partial purchase after the first search would not alter the example.

Also, as the example relies only upon the value of the indirect utility function at a finite number of points, a smooth utility function approximating the w function above would not change the conclusion. Similarly any sufficiently small perturbation of this utility function will give rise to the same search behavior. In particular, the preferences over x and L could have been strictly monotonic and convex without altering the conclusion.

5. A Reservation-Price Theorem

The following property is useful for reservation-price behavior:

$$|u(p_k, I - c) - u(p_i, I - c)| \leq |u(p_k, I) - u(p_i, I)|, \tag{*}$$

any $p_i, p_k \in P$, and $I > c$.⁴

PROPOSITION 3. *If $w(x, L)$, the direct utility function, is homogeneous of degree h , property (*) obtains, for all $h \geq 0$.*

⁴(*) is a rather familiar property in nonlinear programming, demonstrable if $w(\cdot)$ is concave and p_i is sufficiently close to p_k . However, (*) is crucial globally, not just for nearby prices. For an arbitrary concave utility function, the following is true:

If a searcher with experience (μ, ρ) and income I optimally stops searching upon eliciting p_i , then there exists $\epsilon > 0$ such that, were the last price observed instead p_k , optimal searching also would stop, for all $p_k \in (p_i - \epsilon, p_i)$.

This is not a particularly useful result.

PROOF,

$$w(\alpha x, \alpha L) = \alpha^h w(x, L) \quad \text{for all } \alpha > 0$$

implies

$$u(p_i, \alpha I) = \alpha^h u(p_i, I) \quad \text{for all } \alpha > 0.$$

Thus

$$\begin{aligned} u(p_k, I - c) - u(p_i, I - c) &= \left(\frac{I - c}{I}\right)^h [u(p_k, I) - u(p_i, I)] \\ &\leq u(p_k, I) - u(p_i, I). \end{aligned}$$

THEOREM. *Assume property (*), and that prior beliefs are Dirichlet. Then the stopping rule for an expected-utility-maximizing searcher exhibits the reservation-price property.*

The proof, which is nearly parallel to Rothschild's argument, is omitted.⁵

Arai [1] adds one condition to this model to establish that increased perceived uncertainty lowers the reservation price function, and thus raises the expected number of searches. Specifically,

PROPOSITION 4 (ARAI [1]). *Assume (*) and that $\partial^2 u(p, I) / \partial p^2 \geq 0$. If an expected-utility-maximizing searcher's prior is Dirichlet, then increasing risk lowers the reservation price.*

Arai's result is then a generalization of Theorem III in Rothschild [3].

6. Concluding Remarks

From the theorem we see that the reservation-price property obtains if the searcher has a (direct) utility function which is homogeneous of degree h , any $h \geq 0$. Thus, marginal utility of income can diminish (at arbitrarily high rates) without destroying the reservation-price property so long as it diminishes "regularly." In the example, the marginal utility of income is constant for low income, then diminishes, then is once again constant. This suggests that third-order considerations are related to the question of reservation price behavior.⁶

⁵The proof, which is available from the first author in the working paper version of this research, closely follows Rothschild's line of argument, bounding differences of maxima analogously to his bounds on differences of minima. To use the inductive hypothesis, Rothschild must consider a searcher obtaining a lower price quotation p_k instead of a higher price p_i an arbitrary number of times. We must also adjust for differing levels of income across the inductive hypothesis, and so must use property (*).

⁶This is a minor revision of Faculty Working Paper #546, College of Commerce and Business Administration, University of Illinois at Urbana-Champaign, Feb. 1979 (same title). The working paper version contains the proof of the theorem. Support of revision and preparation by NSF SES 79-15356 is acknowledged. We thank, without implicating, Dave Cass and Rich Kihlstrom for helpful comments.

References

1. ARAI, KAZUHIRO, "Expected-Cost-Minimizing Search and Expected-Utility-Maximizing Search," unpublished, University of Illinois, 1979.
2. ROTHSCHILD, MICHAEL, "Models of Market Organization with Imperfect Information: A Survey," *J. Political Econom.*, Vol. 81 (1973), pp. 1283-1308.
3. ———, "Searching for the Lowest Price When the Distribution of Prices is Unknown," *J. Political Econom.*, Vol. 82 (1974), pp. 689-711.