

Option Pricing for Inventory Management and Control

Bryant Angelos, McKay Heasley, and Jeffrey Humpherys

Abstract—We explore the use of option contracts as a means of managing and controlling inventories in a retail market. Specifically, merchants can buy option contracts on unsold inventories of retail goods in an effort to hedge, pool, or transfer risk. We propose a new kind of European put option on an inventory where the holder is allowed to freely adjust the original sale price of the underlying good throughout the contract period as a means of controlling demand. Assuming the retailer will chose the profit maximizing pricing policy, we can price the option accordingly.

I. INTRODUCTION

The capital markets have been recently scrutinized in the press, congressional hearings, talk shows, etc., due in part to the active trading of highly complex and yet poorly understood over-the-counter derivatives and other investment vehicles used to transfer cash flows and risks in the consumer debt and mortgage industries. The recent financial crisis and global recession caused by large-scale holdings of *toxic paper* involving collateralized debt obligations, mortgage-backed securities, and subprime mortgages, demonstrates that we still have much to learn about financial engineering and the proper quantification of risk and reward. Nonetheless, despite the ever-growing list of blunders and debacles, businesses in the aggregate enjoy increasingly greater access to investment capital, less exposure to market risk, increased liquidity, and higher productivity as a result of derivative securities and investment banking [5], [7].

In the retail markets, merchants hold inventories of goods and services in a manner not terribly different than investors holding portfolios of investment securities. Indeed an investment portfolio can be viewed as an inventory and vice versa. Our goal is to identify the commonalities between these two areas and help tie these two industries together, at least mathematically, under the common umbrella of systems theory.

There are some obvious differences between the capital and retail markets. One notable example is that the former is highly instrumented and transactions are usually cleared centrally by financial intermediaries on a trading floor, whereas the latter has wildly varying trading practices and corresponding risks, thus resulting in relatively high transaction costs, brand effects, and geographically localized points of sale. Nonetheless as the information age continues to evolve and consumer purchases become more electronic, we may expect to see a similar commoditization in the retail markets over time. There are already web services such

as Paypal/eBay, PriceGrabber, and Amazon that centralize and/or clear numerous retail storefronts thus removing the need for a direct relationship between the consumer and the retailer. Similar transaction clearinghouses likewise exist in the wholesale and shipping industries.

The idea of an option in retail can be mathematically motivated by the classical newsvendor problem, which we review in Section II. In this problem, the retailer seeks to determine the optimal inventory needed to maximize expected profit when the demand is uncertain, the sale price is fixed, and the inventory is perishable so that any remaining items at the end of the day (or whatever the timeframe) are sold for scrap at some known salvage value; see for example [8]. Typical examples of perishable goods include bakery items in a grocery store and magazines on a newsstand, but electric power, theater tickets, and airplane seats also fit this description. In fact one could extend this to virtually any good by using perishability as a proxy for depreciation, where the future depreciated value of the item is its “salvage value” under the newsvendor paradigm. We remark that the solution of the newsvendor problem has many similarities to the Black-Scholes pricing of an option in the stock market.

Recall that a put option contract gives the owner the right, but not the obligation, to sell a particular underlying asset at a given price and at a specified instance (or period) of time. Hence, the value of a put option depends on the (expected) future value of the underlying asset. If the asset price rises beyond a certain level, called the strike price, the put option becomes worthless. If it goes below the strike price, the put option can be exercised and the writer pays the buyer the difference between the strike price and the asset price. In the case of a retail option, the merchant buys the option from the writer who then, at the end of the period, pays the difference between the strike price and the salvage value on any remaining inventory, or alternatively the contract could allow or require the writer to pay the strike price and take physical possession of the unsold inventory; see [2], [3]. One can price such options using the risk-neutral valuation, where the premium paid by the buyer at the beginning of the contract period is the future value of the writer’s expected payout.

In an effort to make retail options more useful, we relax the constraint on a fixed sale price u , thus allowing the retailer to freely adjust the sale price of the underlying good throughout the contract period in order to control demand, and thus maximize profit. In Section III, we formulate this by considering demand as a non-homogenous Poisson process where the expected arrival rate depends on the sale price of the underlying good. We model our inventory state as

Bryant Angelos and McKay Heasley are research assistants under the advisement of Jeffrey Humpherys, an assistant professor in the Department of Mathematics at Brigham Young University, Provo, UT, USA 84602. Correspondence should be sent to `jeffh@math.byu.edu`.

a continuous Markov process, which can be described as a system of ordinary differential equations (ODEs) that depend in part on the retailers pricing policy, which if known gives the risk-neutral valuation of the option contract. In Section IV, we determine the optimal pricing policy by solving a dynamic program; see for example [1]. We consider specifically linear and log-linear demand rates, although our method is quite general. We likewise reduce our dynamic program to a system of ODEs, which can be solved to provide the optimal pricing policy. We assume that the retailer will use the optimal pricing policy, and hence we set the option price accordingly. Finally, we perform some simulations in Section V and discuss future directions and offer additional ideas for using portfolio optimization tools in retail in Section VI.

II. MOTIVATING EXAMPLE

In this section we review the classic newsvendor problem. We consider the task of determining the inventory policy (or production policy) that will maximize expected profit, given that one has a good forecast of future demand. More precisely, we want to know how much inventory I_0 our retailer should buy from the wholesaler at cost C_0 to maximize expected profit, knowing that at the end of the sales period $[0, T]$ the remaining inventory I_T will be salvaged at a price of C_T per unit. We assume a fixed sale price u and that the demand X is a random variable with known statistics. Since the demand may exceed the available inventory, we relate the quantity sold as the random variable

$$Q = \min(X, I_0).$$

Hence, the profit Π is the revenue minus the cost of inventory plus the salvage, or rather

$$\Pi = Qu - I_0C_0 + I_T C_T.$$

where I_T is the remaining inventory, that is, the amount of product that does not sell and is thus salvaged at a price of C_T at the end of the sale period. This is given by

$$I_T = I_0 - Q = \max(I_0 - X, 0) := (I_0 - X)^+. \quad (1)$$

The profit Π is then expressible as the return one would get from selling out, minus the lost revenue from excess inventory that is only partially recovered from salvage. This yields

$$\Pi = I_0(u - C_0) - (u - C_T)I_T. \quad (2)$$

Hence, the expected profit is

$$E[\Pi] = I_0(u - C_0) - (u - C_T)E[I_T]. \quad (3)$$

We remark that in (2), the only random variable in computing the profit is I_T . Hence it suffices in much of our analysis to compute $E[I_T]$.

A. Optimal Inventory Policy

In the classical newsvendor problem, we seek to determine the optimal initial inventory I_0^* that maximizes profit, given a retail price u and marginal (or wholesale) cost C_0 .

For convenience, we assume X is a continuous random variable. Let $F(x)$ be the (usually strictly monotone) cumulative distribution function (cdf) and $f(x) := F'(x)$ be the probability density function (pdf) for the demand X . To find local extrema of $E[\Pi]$, we take its derivative with respect to I_0 and set it to zero. Hence

$$\frac{d}{dI_0} E[\Pi] = u - C_0 - (u - C_T) \frac{d}{dI_0} E[I_T] = 0,$$

or equivalently,

$$\frac{d}{dI_0} E[I_T] = \frac{u - C_0}{u - C_T}. \quad (4)$$

It is easy to show that the initial inventory I_0 that satisfies (4) is profit maximizing. To compute $E[I_T]$, we define the indicator random variable

$$J = \begin{cases} 1 & X \leq I_0 \\ 0 & X > I_0. \end{cases} \quad (5)$$

Then $I_T = J(I_0 - X)$ and

$$E[I_T] = I_0 E[J] - E[JX]. \quad (6)$$

Hence, the expected remaining inventory takes the general form

$$E[I_T] = I_0 \int_{x \leq I_0} f(x) dx - \int_{x \leq I_0} x f(x) dx. \quad (7)$$

Taking the derivative yields

$$\frac{d}{dI_0} E[I_T] = E[J] + I_0 \frac{d}{dI_0} E[J] - \frac{d}{dI_0} E[JX].$$

It is relatively straightforward to show that $E[J] = F(I_0)$,

$$\frac{d}{dI_0} E[J] = f(I_0) \quad \text{and} \quad \frac{d}{dI_0} E[JX] = I_0 f(I_0).$$

Hence

$$\frac{d}{dI_0} E[I_T] = F(I_0).$$

Combining with (4), the optimal inventory policy is

$$I_0^* = F^{-1} \left(\frac{u - C_0}{u - C_T} \right). \quad (8)$$

This is the solution to the classic newsvendor Problem. We remark that mathematically, the above analysis is very similar to the derivation of the Black-Scholes pricing of a European call option; see for example [10].




Fig. 1. Option value at expiration of (a) a stock option and (b) a retail option.

B. Fixed-Price Retail Option

We now explore the newsvendor Problem from the point of view of an option contract; see also [2], [3]. Recall that a European put option gives the holder the right, but not the obligation to sell an underlying good at a specified price and time.

Consider a European put option on the remaining inventory I_T with strike price $K \geq C_T$, initial inventory I_0 , and fixed sale price u . In other words, we have a derivative contract where the retailer pays p dollars to the option writer at the beginning of the contract period and then receives payment of KI_T in return for the remaining inventory. If the retailer sells out, the retailer pays nothing, and the net option value is $-p$ corresponding to the initial payment at the beginning of the contract period; see Figure 1. According to the risk-neutral valuation of the contract, the option is valued at

$$p = (K - C_T)E[I_T] \quad (9)$$

dollars (modulo interest). Moreover the profit for the retailer is given by

$$\Pi = I_0(u - C_0) - (u - K)I_T - KE[I_T]. \quad (10)$$

We remark that the expected profit in (10) is the same as that of (2). Indeed options that are based on the risk-neutral valuation do not improve the expected outcome, but they do reduce the uncertainty of the outcome. Specifically, the variance of (10) is given by $(u - K)^2 \text{Var}(I_T)$, whereas the variance of (2) is larger and given by $(u - C_T)^2 \text{Var}(I_T)$.

The problem with fixed-price retail options is that retailers may want to change prices throughout the contract period. If the underlying product is selling faster than expected, the retailer may be able to increase profits by raising prices. If the retailer is over stocked, he or she may be better off lowering prices to stimulate demand. In the next section, we introduce a new type of European put option that will allow for dynamic pricing.

III. A NEW KIND OF EUROPEAN PUT OPTION

We present a new type of European put option for retail, which allows the option holder to adjust prices throughout the contract period $[0, T]$ in an attempt to maximize profits. We present a method for computing the risk-neutral valuation of this retail option, where the initial inventory is m units. We assume that the demand is a nonhomogeneous Poisson process, where customer purchases arrive at the rate $\lambda(u(t))$, which depends on the time-dependent sale price $u(t)$ set by the merchant. Let $I_j(t)$ be the probability that the inventory level is j units at time t . Then the vector $I(t) = (I_0(t), \dots, I_m(t))$ is a continuous-time Markov process, where the transition probabilities likewise depend on $\lambda(u(t))$.




Fig. 2. Single period branch describing (11).

A. Discrete-Time Formulation

We discretize the time interval $[0, T]$. Let $u_{i,j}$ denote the sales price when the inventory is at j units at time $i\Delta t$, where $i = 0, \dots, n$, $j = 0, 1, \dots, m$, and $\Delta t = T/n$. For small intervals of time $[t, t + \Delta t)$, we assume that an item sells with probability $\lambda(u_{i,j})\Delta t$. Let $I_{i,j} = I_j(i\Delta t)$ denote the probability that the inventory level is j units at time $i\Delta t$. Clearly $\sum_{j=0}^m I_{i,j} = 1$, $I_{i,j} \geq 0$ for all i, j . Hence, following Figure 2, the transition probabilities satisfy

$$I_{i+1,j} = \lambda(u_{i,j+1})\Delta t I_{i,j+1} + (1 - \lambda(u_{i,j})\Delta t)I_{i,j}. \quad (11)$$

Hence, given a pricing policy $\{u_{i,j}\}$, we can determine the expected remaining inventory by computing the finite-difference equations forward to time $t = T$. Then the expected remaining inventory is given by

$$E[I_T] = \sum_{j=1}^m j I_j(T).$$

Then the option valuation comes from (9), yielding

$$p = (K - C_T) \sum_{j=1}^m j I_j(T). \quad (12)$$

This is the risk-neutral valuation of the option.

B. Continuous-Time Formulation

We can now take the limit as Δt goes to zero, thus turning this problem into a system of ODEs. Simplifying (11), yields

$$\frac{I_{i+1,j} - I_{i,j}}{\Delta t} = \lambda(u_{i,j+1})I_{i,j+1} - \lambda(u_{i,j})I_{i,j},$$

which, as $\Delta t \rightarrow 0$, gives the ODE

$$\frac{d}{dt}I_j(t) = \lambda(u_{j+1}(t))I_{j+1}(t) - \lambda(u_j(t))I_j(t), \quad (13)$$

where $I_j(0)$ is the initial probability distribution.

Hence, given the pricing policy $\{u_j(t)\}_{j=0}^m$, we can compute the state vector $I(t)$ at time T , by integrating the system of ODEs in (13) on the interval $[0, T]$. Figure 3 shows a sample solution to this differential equation for initial inventory $I_{10} = 1$, and $I_9 = I_8 = \dots = I_0 = 0$.




Fig. 3. Inventory probabilities over time.




Fig. 4. Single period branch describing (14).

IV. OPTIMAL PRICING

In the previous section, we showed how to price an option for a given pricing policy. We now show how to compute the optimal pricing policy given a known demand rate $\lambda(u)$. We first compute the optimal price by solving a dynamic program, and then, as in the section above, demonstrate that the problem can be computed as a solution to a system of ODEs by taking the subinterval length to zero. We assume that the retailer will choose the profit maximizing pricing policy, so that we can give the (Nash optimal) risk-neutral valuation of the option.

A. Pricing as a Dynamic Program

By partitioning the contract period $[0, T]$ into n equal subintervals, as above, we can design a dynamic program to determine the optimal pricing policy for the underlying

inventory. On each subinterval, the retailer chooses the price u^* that will maximize expected remaining revenue on that subinterval, given the inventory level j . We proceed by the principle of optimality. Let $E_{i,j}$ denote the expected remaining revenue earned during the i^{th} subinterval of time when the inventory level is j units. If an arrival occurs, the expected remaining revenue is $E_{i+1,j-1} + u_{i,j}$, and if no arrival occurs, the expected remaining revenue is $E_{i+1,j}$; see Figure 4. Thus at each point, we have

$$E_{i,j} = (1 - \lambda_{i,j}\Delta t)E_{i+1,j} + \lambda_{i,j}\Delta t(E_{i+1,j-1} + u_{i,j}), \quad (14)$$

where $\lambda_{i,j} = \lambda(u_{i,j})$. By taking a derivative with respect to u , we find that the expected remaining revenue is maximized when $u_{i,j}^*$ satisfies

$$0 = \lambda'(u_{i,j}^*)(E_{i+1,j-1} - E_{i+1,j} + u_{i,j}^*) + \lambda(u_{i,j}^*). \quad (15)$$

Hence, by solving (15) for $u_{i,j}^*$, we can compute the optimal prices and the expected remaining revenues at time $i\Delta t$ by using the optimal prices and expected remaining revenues at time $(i+1)\Delta t$; see Figure 5. This gives us a complete description of the optimal pricing policy and expected remaining revenues for all subintervals of time and for all inventory levels.




Fig. 5. Multiple periods for computing the expected remaining revenue. The rightmost column corresponds to the final period, where the option pays out K times the remaining inventory. By iterating backwards in time, we can determine the rest of the grid values for both the optimal price u^* and the corresponding expected remaining revenue values.

B. Linear and Loglinear Demand Rates


To more directly connect the expected remaining revenues and optimal prices, we compute optimal pricing policies for linear and log-linear demand rates.

1) *Linear Demand*: If the parameter of demand $\lambda(u)$ is linear with respect to price, that is $\lambda(u) = b - au$ for some $a, b > 0$, then (15) gives


$$u_{i,j}^* = \frac{E_{i+1,j} - E_{i+1,j-1}}{2} + \frac{b}{2a}.$$

2) *Loglinear Demand*: We repeat the process described above but with a different demand rate. We let $\lambda(u) = au^{-b}$, $b > 1$. Then (15) gives

$$u_{i,j}^* = \frac{b}{b-1}(E_{i+1,j} - E_{i+1,j-1}).$$



(a)



(b)

Fig. 6. Realization of (a) the optimal price of the underlying asset throughout contract period and (b) the option price throughout the contract period.

C. Pricing in Continuous Time

We conclude this section by showing that the optimal prices and expected remaining revenues can be computed by solving a certain system of ODEs much as we did in the previous section. By simplifying (14), we get

$$\frac{E_{i+1,j} - E_{i,j}}{\Delta t} = \lambda(u_{i,j})(E_{i+1,j} - E_{i+1,j-1} - u_{i,j}),$$

which, by letting $\Delta t \rightarrow 0$, yields the ODE

$$\frac{d}{dt} E_j(t) = \lambda(u_j(t))(E_j(t) - E_{j-1}(t) - u_j(t)), \quad (16)$$

where $E_0(t) = 0$ and $E_j(T) = jK$.

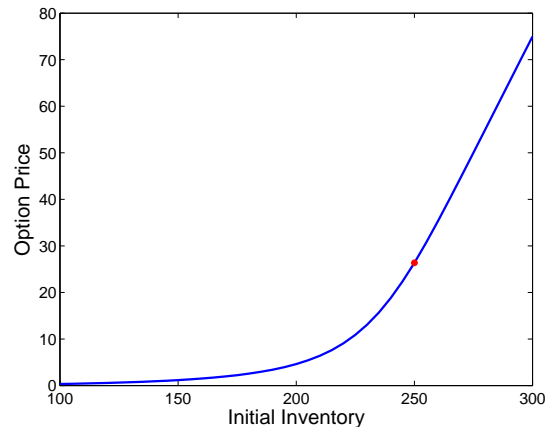
V. SIMULATIONS

We can code up our two systems of ODEs in MATLAB in just a few lines of code. For the integration, we use `ode45`. This variable-step routine allows for large steps when the errors are small, thus significantly speeding up computation time. Most of our runs executed in under a second. For large inventories or contract periods, the runtime increased proportionately. In the case that the inventory


becomes very large, one could treat inventory levels as a continuous variable, and resort to stochastic PDEs, as is done with option pricing in the stock market (despite the fact that the underlying also takes on discrete prices).

As a check, we generated simulations to realize consumer purchases and examine how the price of both the underlying and the option move over time; see Figure 6. We randomly generated numbers on the unit interval at each time step, and then dropped the inventory by one if the number generated was less than $\lambda(u_{i,j})$. We then computed the option price at that moment. Notice that both images seem to behave like a random walk, and look similar to stock-market data.

Finally, we describe how the option price varies as we toggle parameters. In Figure 7 we plot the option price against the initial inventory and the strike price, respectively. These vary in a monotonic, nonlinear fashion, as expected, growing linearly in the infinite limit and converging to zero for small values. We remark that the option price for varying contract periods has an inverse relationship to how the option price varies with initial inventory, that is, small contract periods price as if the initial inventory were large, and large contract periods price as if the initial inventory were small.



(a)



(b)

Fig. 7. Option price as we vary (a) the initial inventory and (b) strike price.

VI. INVENTORY VS. PORTFOLIO MANAGEMENT

From an inventory manager's point of view, it can be difficult, when looking at raw transactions data, to know whether low sales are due to low demand (e.g., nobody is buying) or low supply (e.g., retailer is out of stock). A retail option allows one to easily differentiate between low sales and low supply. Indeed the price of the option goes to zero when the inventory runs out whereas low sales would result in the option value being high. By using options and other financial instruments in a retail environment, merchants can better understand their markets by treating their inventories as money managers do portfolios. However in retail, there is the added advantage of being able to dynamically control the price of the underlying goods in an attempt to maximize profits.

There are many directions to take this portfolio view of inventory theory and control. Several recent results exist for managing portfolios of stock options over multiple periods; see for example [9]. We could also explore risk-management strategies ranging from a myriad of techniques including value-at-risk, cash-flow-at-risk, etc. Another direction is the further use of more exotic options. Variations on floors, caps, Asian options, Bermuda options, etc., should all be considered to see if there are "retail possibilities".

Another area of interest is how to price an option when dealing with an incompetent retailer. Note that the price of the option proposed here requires that the retailer continually adjust prices to maximize profits. However, if the retailer isn't adjusting them optimally, then the writer will be forced to buy more unsold inventory than necessary at the end of the contract period. While the existence of a Nash equilibrium should preclude such events theoretically assuming symmetric information, market efficiency, etc., in practice prices may be set badly. Room for such incompetence may need to be priced in the option.

We remark that options have been around for centuries in the financial markets [11] (most notably over the past 30 years), and they have only recently been suggested for use in a retail environment [2], [3]. There is, however, an extensive *supply chain contracts* literature going back for decades; see [12] for a review, as well as a *real options* literature in business strategy; see [6] and references within. It would be worthwhile to pull all of this literature together under the common umbrella of systems theory. Indeed recent events demand better analysis of financial modeling and a deeper understanding of the dynamics of the markets.

Finally, it is unclear whether retail option contracts could ever materialize into exchange-traded instruments like options on stocks, futures, etc. Surely some could find a natural home in certain niche retail markets, but overall the lack of standards, trust, third-party or government regulation, and/or difficulties in instrumentation may make it difficult to reliably create an active market for retail options in a general setting. We remark, however, that the use of retail options as a means of centralized inventory management and control in a retail chain, or more generally an organized network of

manufactures, wholesalers, distributors, retailers, speculators, and financiers, may prove to be valuable. Indeed, retailers can essentially operate as portfolio managers. Retail chains can buy and sell options internally, between their distribution centers, individual stores, and their own corporate finance department can act as the hedge fund for the company by pooling risks much like an insurance company. Moreover, by using the put option proposed in this paper, where the prices can be adjusted dynamically, the retailers can then use the price to control the returns in their portfolio, both for an individual store and in the aggregate.

VII. FUTURE WORK

We are currently considering variations on this work. To make our retail option viable, we need to further explore the degree to which the option price is sensitive to input parameters. We also need to flesh out the details about the variance in the option payoff.

We are also interested in seeing if a set of options can be written on an underlying collection of inventory, whether it is reasonable to have variable payoffs depending on how much inventory is remaining, and how the option might work with multiple periods and multiple echelons. Other issues include state estimation of the arrival rate, the effects of substitutes and complements, and the effects of coupons and other manufacturer-based promotions on the option price.

VIII. ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation, Grant No. DMS-0639328, BYU mentoring funds, and a gift from Park City Group, Inc. The authors thank PCG President and CEO Randall Fields and his associates for useful and stimulating conversations throughout this work.

REFERENCES

- [1] D. Bertsekas. *Dynamic Programming and Optimal Control*, (Volumes 1 and 2). *Athena Scientific*, 2005.
- [2] A. Burnetas and P. Ritchken. Option pricing with downward-sloping demand curves: The case of supply chain options. *Management Science*, 51(4):566–580, 2005.
- [3] F. Chen and M. Parlar. Value of a put option to the risk-averse newsvendor. *IIE Transactions*, 39(5):481–500, 2007.
- [4] J. Cox, S. Ross, and M. Rubinstein. Option Pricing: A Simplified Approach. *International Library of Critical Writings in Economics*, 143:461–495, 2002.
- [5] A. Demircuc-Kunt and R. Levine. Stock markets, corporate finance, and economic growth: An overview. *World Bank Economic Review*, 10(2):223–39, May 1996.
- [6] A. Dixit and R. Pindyck. *Investment under uncertainty*. Princeton University Press Princeton, NJ, 1994.
- [7] R. Levine and S. Zervos. Stock markets, banks, and economic growth. *American Economic Review*, 88(3):537–58, June 1998.
- [8] E. L. Porteus. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, 2002.
- [9] J. Primbs. Dynamic hedging of basket options under proportional transaction costs using receding horizon control. Preprint, 2007.
- [10] S. M. Ross. *An introduction to mathematical finance*. Cambridge University Press, Cambridge, 1999.
- [11] B. Smith. *A History of the Global Stock Market: From Ancient Rome to Silicon Valley*. University Of Chicago Press, 2004.
- [12] A. A. Tsay, S. Nahmias, and N. Agrawal. Modeling supply chain contracts: a review. In S. R. Tayur, R. Ganeshan, and M. J. Magazine, editors, *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers, Norwell, MA, USA, 1998.