Chapter 2

Inventory Optimization using a Renewal Model for Sales

Dan Calistrate[1], Marc Paulhus[2], Miro Powojowski[3], Gordon Sick[4],

Report written by Marc Paulhus

2.1 Introduction

The problem posed by Dr. Greg Robel from The Boeing Company was to optimize inventory control and production rates of a certain item in order to take advantage of unexpected surges in demand. Certainly any optimization will require a realistic and simple model to predict future sales. The problem presenter had suggested that sales might arrive according to a Poisson distribution.

We suggest that the company look at *renewal theory* for models of future sales orders. They have some very distinct advantages. They are powerful and reasonably easy to use, in fact the Poisson distribution is a special case. Renewal models are flexible enough to incorporate a variety of characteristics, such as clustering or regularity, upward or downward drift, and mean reversion. Also, some arguments can be made to justify the renewal models based on financial intuition.

The next section will define and discuss the renewal models that we are suggesting. In Section 2.3 we set up a realistic (albeit naive) problem for which we have an analytic solution. The last section will more directly address the problem presented and will discuss possible solution methods.

2.2 Renewal Models for Future Sales

An informal definition of the renewal process is now given. For a more rigorous treatment of the renewal theory the reader is referred to any standard text on the subject, e.g., Cox [1962]. Some events - in this case sales of a product - will be assumed to occur at point times in the interval $[0, \infty)$. Thus any sales history can be summarized in a (finite or infinite) sequence $\{T_i\}_{i \in J}$ (e.g., $\{10, 14, 31\}$ is a short sales history with three sales at times 10, 14 and 31). A probability density function $f(t)$ will be specified and it will be assumed that

Condition 1

$$P(T_1 \leq t) = \int_0^t f(\tau) d\tau$$

[1] University of Calgary
[2] University of Calgary
[3] Université de Montréal
[4] University of Calgary

as well as for any $k > 0$

$$P(T_{k+1} - T_k \leq t) = \int_0^t f(\tau)d\tau$$

Simply put, the first sale arrives randomly at T_1, following the density $f(t)$, and given that a sale occurs at some time T_k, the additional time till the next sale $T_{k+1} - T_k$ is also random with density $f(t)$. At each sale, the system is renewed and the history prior to the last sale is of no consequence.

For example, look at the probability density function given in Figure 2.2. Their precise definitions will be given later. For now notice that, with the middle density, most of the mass under the curve lies near the origin. Hence, it is most likely that another sale will occur soon after a sale has been made. If a sale is not made in a long period of time we can expect to have to wait even longer before the next sale occurs. This should produce a clustering effect in the sales predictions.

With the bottom density, there is very little mass near the origin and hence it is unlikely that a sale will occur immediately following another sale. This should cause the sale predictions to be somewhat regular.

The top density corresponds to the Poisson distribution.

Figure 2.2 displays a single sample realization of future sales under each of the corresponding density functions. You will notice that the middle density did produce clustered sale predictions whereas the bottom density predicted sales on a quite regular basis. The top density produced sale predictions somewhere in between.

These remarks can be given a precise formulation as follows: let

$$p(t, \Delta t) = P(T_{k+1} - T_k \in (t, t + \Delta t] \,|\, T_{k+1} - T_k > t)$$

This corresponds to the probability of a sale arriving at some time in the interval $(t, t + \Delta t]$ after the last sale, conditional on the event that there was no sale in the interval $(0, t]$ after the last sale. Let Δt be an arbitrary, but fixed, small positive number. In the case of the density $f(t)$ as in the middle figure, one finds that $p(t, \Delta t)$ is a decreasing function of t, so as the time since the last sale goes by, the probability of a sale arriving in a fixed time step Δt decreases. For the bottom density, $p(t, \Delta t)$ is initially an increasing function and it reaches a peak for some positive value of t. For the Poisson process, the function $p(t, \Delta t)$ is constant in t, hence this process is sometimes described as memoryless.

The mean of $f(t)$, μ, is the expected inter-sale waiting time. The amount of clustering is determined by the shape of the density function and is separate from the mean. Although any density function leads to a renewal process, it is proposed here that the Gamma distribution family be used. The Gamma family can be parametrized with just two parameters as follows:

$$\Gamma(\alpha, \rho; x) = \frac{(\frac{x}{\rho})^{\alpha-1} e^{-\frac{x}{\rho}}}{\rho\Gamma(\alpha)}$$

with $\rho > 0$, $\alpha > 0$. The mean of this distribution is $\rho\alpha$ and the variance is $\rho^2\alpha$. The parameter ρ is the scale parameter, controlling the mean, whereas α is the shape parameter, controlling the amount of clustering, regardless of the value of ρ. The Gamma family is convenient since it contains members exhibiting the clustered ($\alpha < 1$) and the regular ($\alpha > 1$) behaviours, as well as the Poisson process ($\alpha = 1$). Finally, the family is well known and many estimators for its parameters have been proposed and investigated.

If historical sales data is available then μ should be easy to estimate. Otherwise, μ should be determined by management based on future sales expectations. The choice of α should be made based on sales history tempered by managerial intuition.

We can model drift by including it in the mean of the renewal density function. That is, if we expect that demand should increase by 8% a year, then rather than renewing with an identical function after a sale, we renew with a density function with a correspondingly higher mean. Note that the mean of the renewal density function will depend on the absolute time, not on the time passed since the last sale. An example of a realization of sales with increasing drift is given in

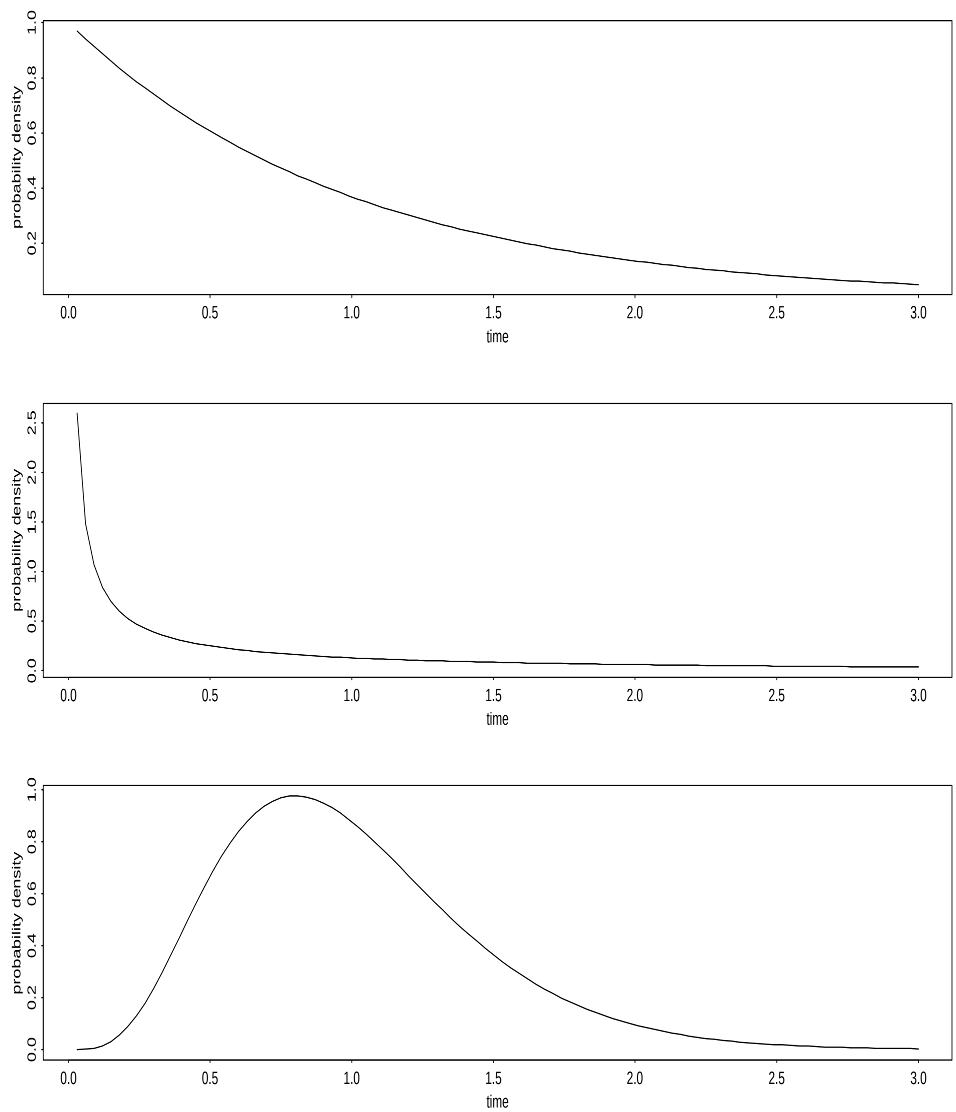

Figure 2.1: Examples of Gamma density functions. The top corresponds to the Poisson Distribution i.e. $\alpha = 1$. The middle corresponds to $\alpha = 1/5$ and will produce clustered sales predictions. The bottom corresponds to $\alpha = 5$ and will produce regular sales predictions.

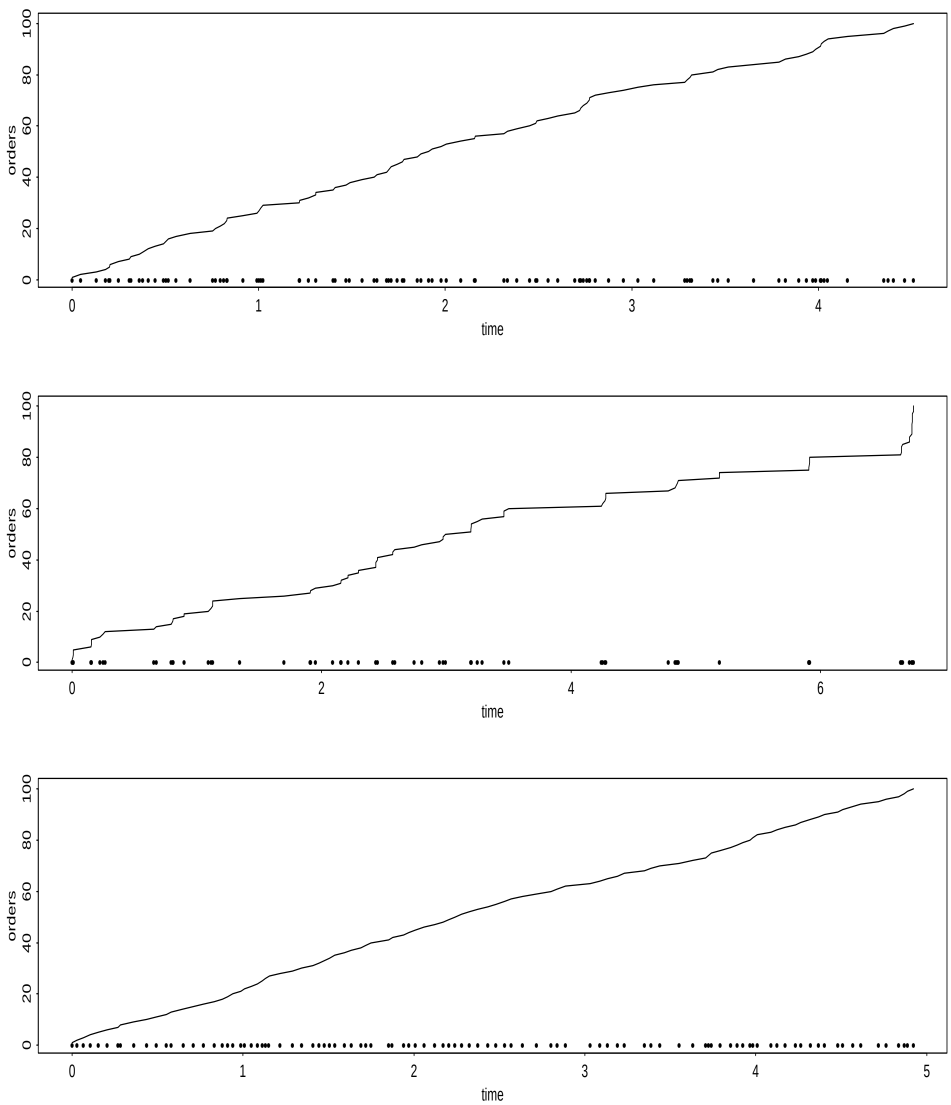

Figure 2.2: A simulated set of orders under the corresponding density functions with shape parameters as in Figure 2.2 and identical means. The dots along the bottom represent the simulated sales events. The line indicates the cumulative sales.

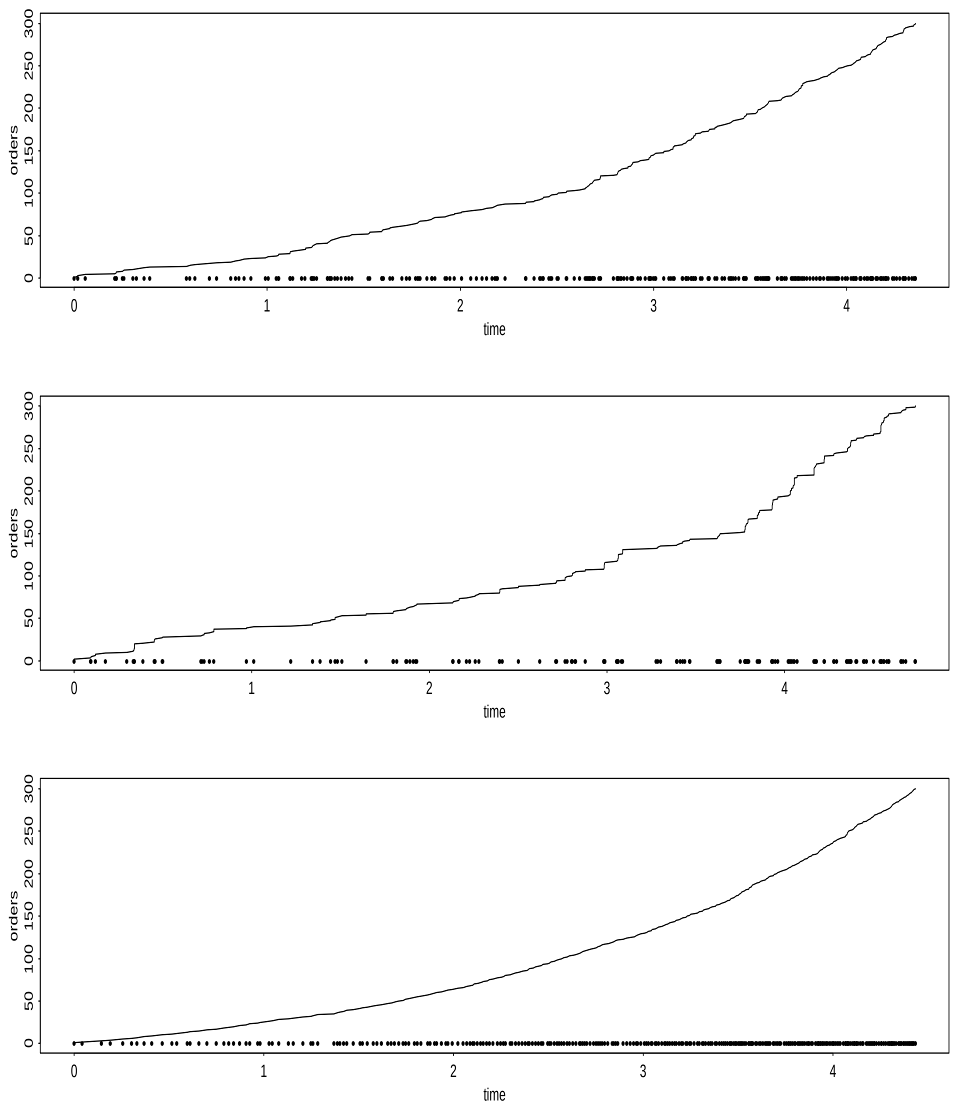

Figure 2.3: A simulated set of orders under the corresponding density functions with shape parameters as in Figure 2.2 with drift included. The dots along the bottom represent the simulated sales events. The line indicates the cumulative sales.

Figure 2.2. The density functions used to produce the figures are the same as those in Figure 2.2 except we included an upward drift in μ.

Mean reversion can be modeled in a similar way. If, in our simulated sales predictions, we are above/below the predicted mean then we can decrease/increase μ accordingly. One could also make α a function of the distance from the expected mean.

The process can become more complex if desired. For example, the management may have identified two independent customer types, one of which buys the product on a regular basis and one of which is more sporadic. Then, the sales could be modeled by two simultaneous and independent renewal processes, each with a different μ and α.

More complex models are easy to state but will complicate an analytical approach. If we are using simulation to answer our questions then the more complicated renewal processes should be no more difficult to use.

2.3 Analytical Solution to a Simpler Problem

To help the reader understand the procedure, we propose a hypothetical scenario which uses the renewal process and can be solved analytically. This scenario is much simpler than the problem posed to the workshop but serves as a useful example.

Suppose there is a software company that has developed a piece of software which we will call *product A*. Copies of the software are produced instantaneously and are sold to the customer for P dollars. Note that it is the lack of need for inventory control which makes this problem simpler than the one presented to the workshop. The company has already invested money in the development of product A and must now concentrate on optimizing the profits from the sale of the product. We will discretize the problem by defining a unit step of time. For the company to keep the product line alive for another unit of time will cost c dollars. This might reflect cost of advertising and the cost of a sales team.

Furthermore, the company has assigned the development team to another project and has no interest in developing further versions of product A. Like all software they know that the product will have a limited life span and hence we must answer the question, when is the optimal time to abandon the product (stop advertising and layoff the sales team). The residual value at abandonment will be K dollars which can be positive or negative. We are also interested in pricing the product line for the case that we have the opportunity to sell it.

Suppose management, based on past sale histories of similar products, concluded that the sales of the product will be clustered. This may be because a single buyer is likely to buy more than one copy or maybe if one corporation upgrades its software then it puts pressure on their competitors to do the same or maybe company B is copying the strategy of company A since company A has devoted some effort in researching the available products. In any case, μ and $\alpha < 1$ are chosen. Realistically we should include a downward drift but we will omit it for simplicity.

We will say that the product line is worth A_0 dollars at the start, which is an unknown that we are interested in computing. Start by looking at the tree of possible outcomes, Figure 2.3. Each level of the tree represents the next unit of time forward in our discrete model. Starting at the root of the tree, move right if a sale is made in the next time step and left if a sale is not made. The nodes of the tree are labeled with the value of the product line at that state. Due to the properties of the renewal process, if the product line was worth A_0 at time zero, then the product line will be worth A_0 just after any sale. We denote by A_t the value of the product line under the assumption that we have not made a sale in the last t time steps.

We will assume a risk-neutral investor and a riskless interest rate of r. Fix an arbitrary time, T, such that if a sale is not made in the last T time steps then we will abandon the product line. Hence we have assumed $A_T = K$. We now have the following set of linear equations:

$$A_t = \frac{\pi_t(A_0 + P) + (1 - \pi_t)A_{t+1} - c}{1 + r} \quad \text{for} \quad t = 0 \dots T - 1$$

$$A_T = K$$

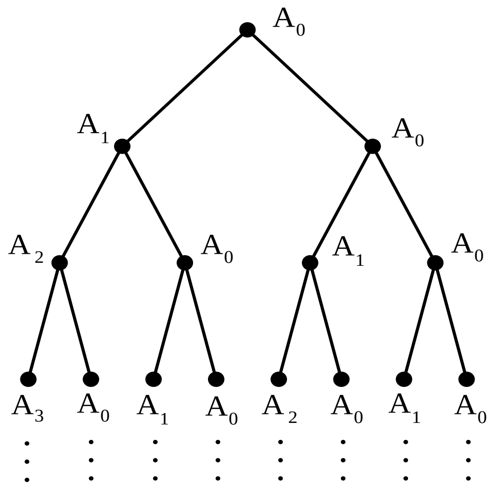

Figure 2.4: Tree of possible outcomes. A left move means no sale is made on the last time step, a right move means that a sale was made on the last time step. The labels represent the value of the product line at that state.

where π_t is the sale attenuation function, that is the probability that a sale is made at time t given that the last sale was made at time zero (a discrete version of the function $p(t, \Delta t)$ defined in section 2.2).

This is a linear system of $T + 1$ equations and $T + 1$ unknowns. Solve it to determine the value of the project under the assumption that the optimal time to abandon the project is if a sale has not been made in T time steps. Iterate this procedure with different choices for T and maximize A_0. This way we determine the current day value of the product line and the optimal time to abandon.

2.4 The Inventory Problem

The last section solved a problem in which we ignored the aspect of inventory. Here we shall set up the problem as it might be posed where the speed of production and the level of inventory are important factors.

Let us concentrate on a single product, in this case, a single line of planes. Looking at past sales history and estimating the future performance, management should establish which renewal density function should be used (include drift and mean reversion if desired). In light of the expected "unexpected surges in demand" it might be wise to consider a clustered sales model.

Again we will discretize the problem, maybe into one–day time steps. In what follows, please note the difference between a fixed one–time *cost*, and a *fee* which must be paid every time step.

We need to know the initial state of production, the possible states of production, the fee required to maintain each state, and the cost of switching between the states.

We suppose, for our hypothetical scenario, that a single plant can operate at the levels given in Table 2.1. Every time step we have the option to either open a new plant at a cost of D dollars, close a plant at a cost of A dollars, or pay $e_{i,j}$ dollars to change the production level at any existing plant from State i to State j.

We must also keep track of the number of planes in inventory at any time, which we will call I. At a given time step, if $I > 0$ then we must pay a fee equal to the cost of storing the planes, say cI. If $I < 0$ then we have over sold our stock and must pay a fee, say CI, which represents customer dissatisfaction. Note that we make no assumption about C and c. In fact, c might well be greater than C to a monopolist.

Another option, useful where competition is present, would be to have $I < 0$ affect the mean of the renewal function. That is, forcing the customer to wait for the delivery of the product may

	Airplane Production Rate	Fee to Maintain
State 1	100	15
State 2	200	5
State 3	400	4

Table 2.1: Production levels of a hypothetical assembly plant. The first column is the number of time steps required to produce an airplane. The second column is the fee required to maintain the plant at that level of production.

cause a loss in sales.

The optimal strategy for this problem is going to be a function of I, the state of production and the state of sales with respect to the density function. We do not have an analytic solution to this problem, but simulation can be used to price any fixed strategy (to some statistical accuracy) and then various methods can be used to try to optimize the strategy function.

2.5 References

[1] D.R. Cox, 1962. Renewal Theory. Science Paperbacks and Methuen & Co. Ltd.