An Insurance Style Model for Determining the Appropriate Investment Level against Maximum Loss arising from an Information Security Breach

Roger Adkins
School of Accountancy, Economics & Management Science
University of Salford
Greater Manchester M5 4WT UK

r.adkins@salford.ac.uk

The economic consequences of breaches in information security cannot be underestimated. According to the World Bank 2002 survey on reported cyber crime the US$ value of the effects from intrusions from a variety of sources has been increasing at an accelerating rate over the past decade. This survey contains a list of reported intrusions recording criminal and employee abuse with losses ranging from several thousand dollars to many million dollars and the victims of these security breaches are in the main players belonging to the global financial service industries. The KPMG 2002 Global Security Survey covering major international firms reported that the average expenditure on information security represented approximately 10% of the total spend on IT and that the level was expected to rise in the future. The average cost of a breach in information security was estimated at just over $100 thousand. It concludes that reporting procedures concerning hostile intrusions into the information bank were in several cases crude and the methods for formally measuring the economic consequences of breaches lacked the sophistication to provide guidance of any value. The Deloitte Touche Tohmatsu 2003 Global Security Survey makes similar conclusions. It points out that although there is a greater awareness on the increased sophistication of attacks on computer information and encouraging trends for financial institutions to treat information security seriously, greater all-round effort is still required on all aspects of information security.

In this paper, we introduce an insurance style model for investigating the investment outlay in security systems that guarantees an upper limit to the losses incurred through potential security breaches over the specified planning horizon. The proposed model differs from a typical investment model that compares the stream of future benefits (and costs) with the initial capital outlay or equates the marginal cost of additional security investment with the marginal benefits from thwarting security breaches. In these models, any variability in expected future cash flows caused by changes in the investment outlay in security systems is not considered to influence that rate at which the cash flows are discounted, or the discount rate is simply not treated as a relevant variable of analysis. In contrast, the proposed model uses the approach adopted by much of the risk management literature that represents potential future losses by a range of possible values and an accompanying probability distribution function, see for example Dowd (2002). Any investment in information security is then assumed to influence the properties of these potential future losses through changing the range of possible losses and / or the probability function, and thereby the expectation and variance of the probability function. Since the variability of future cash flows determines the appropriate discount rate used in evaluating their present value, any change in the range of possible losses and / or the accompanying probability function will necessarily entail a change in the appropriate discount rate. The proposed model

overcomes this difficulty through applying contingent claims analysis and the law of one price to adjust the cash flows so that they may be legitimately discounted at the risk-free rate

The paper is organised in the following way. The first section explains the conceptual foundation bedded in financial options theory that underpin the insurance style model for deriving the economic level of investment in information security that will limit the loss sustained from security breaches over the specified planning period. Since the solution is not in a closed analytical form, the following section explores the behaviour of this function using numerical methods. The developed model is based on the assumption that investment into information security limits the losses incurred through a breach in security. This basic model is then expanded to include the possibility that the investment will also act as a deterrence and reduce the probability that an intrusion will take place. Finally, the paper is concluded by a discussion of the results and some indications for future work.

The Model

The model defines that when a security breach actually occurs at some instant of time during the planning horizon, the loss incurred by the organisation is represented by a random variable, L. For the purposes of developing the model's formulation, we will assume that this random variable may take on only two binary outcomes at some instant of time: an upper loss value denoted by L^+ that occurs with a probability p, with $0<p<1$, and a lower loss value denoted by L^- that occurs with a probability (1-p). Given that a security breach actually occurs at some instant of time in the future, the expected loss generated by the breach is $\{p\ L^+ + (1-p)\ L^-\}$ at that instant of time. A critical issue that is addressed by the model is the way that this future expectation is discounted to the present time.

The probability that a security breach actually occurs over a small interval Δt is denoted by $\lambda\ \Delta t$ where λ is the probability of a security breach taking place per unit time period. The specified planning horizon T is assumed to be divided into a large number N of small intervals of time Δt, then $T = N\ \Delta t$. When a breach does take place at some instant of time with probability $\lambda\ \Delta t$, the incurred loss sustained by the organisation can take on one of the two binary outcomes as specified by the upper and lower limits of L. When no breach takes place with probability $(1- \lambda\ \Delta t)$, there is no incurred loss.

The formulation of this model is depicted in figure (1). This figure shows the possible outcomes of breach and no breach as well as any losses incurred for the first instant of time from $t = 0$ to $t = \lambda\ \Delta t$. The representation shown in figure (1) can be replicated N times to reveal the entire model over the specified planning horizon from $t = 0$ until $t = T$. initially, we will consider the case of a one period model, that is when $t=0$ and $t=\Delta t$, which can be conveniently represented by $t=1$ for simplicity. After completing the analysis for the one period model, it will be extended to cover the entire planning horizon.

To examine the economic implications of this model, we will treat the case of breach and no breach and the case of high level and low level of loss incurred separately. This implies that the probability of a breach λ and the probability of a high level loss

p are statistically independent so that the chance of a breach occurring does not influence the level of the loss incurred. Under this condition, we may examine the economic implications of the events contained within the box displayed in figure (1), which portray the consequences of a security breach.

In the current formulation, the loss incurred when a breach occurs is defined by the distribution of losses. Consider the introduction of either an insurance policy or a capital investment expenditure on information security that effectively limits the loss incurred by the breach to some specified maximum level. If the organisation is prepared to accept an upper loss limit of Q each time a breach occurs, we can ask how much the organisation is willing to spend on this insurance policy or investment opportunity that limits the maximum loss to Q. The benefits given presence of a breach from such an investment will be equal to (L – Q) whenever L is greater than Q and zero if otherwise. If the fair value of the investment at time t = 0 is denoted by F_0, then the value of this investment at time t=1 will be given by:

$$F_1 = \max\{L - Q, 0\} \quad (1)$$

Figure (2) shows the cost of the investment and the future benefits accruing from the investment at time t=1. Because the future benefits are asymmetric due to the use of the maximum function and do not follow the distribution of L, a critical question is how to discount these future benefits since a discount rate based on the riskiness implied by the distribution of L is clearly inappropriate. The solution is to recognise the similarity between the current formulation and the binomial option pricing model for financial options, see for example: Dubofsky and Miller jr (2003), Hull (2003), Trigeorgis (2000) and to apply contingent claims analysis and the law of one price to represent the cash flows by their risk-neutral counterparts and discount at the risk-free rate.

One of the assumptions of the financial options model is that the outcomes in the absence of the option follow the distribution of a known portfolio of traded securities. The known portfolio of traded securities may be composed of a single security or a combination of traded securities; we will refer to this portfolio as the twinning security. This implies that the distribution of losses incurred through a breach has to be completely correlated with a twinning security and deriving usable results from the model requires that an appropriate twinning security is available. Identifying the twinning security is critical, but it is also problematic.

In the financial options literature, the twinning security is typically the underlying asset, such as a financial asset like a government bond, a security of a company or a commodity like oil, which is continuously traded on a certain exchange. In these cases, the choice of the twinning security is often obvious. Financial options theory has also been applied to real assets in order derive the benefits from undertaking a project arising from managerial flexibility in the presence of uncertainty and investment irreversibility. When options theory is applied to real assets, the twinning security is identified as the security which is most strongly correlated with the project's uncertainty in the absence of flexibility. The twinning security may be a commodity when the main source of the project's uncertainty is derived from the commodity's volatility, or it may be the traded security of the firm in question or of the firm in the same risk class as the project. Following this approach, we may be able to argue that when the potential loss, or value at risk, is proportionate to the focal

firm's capital value, then its share price would be an eligible twinning security. More recent literature on real options is recognising that a natural or highly correlated twinning security may not exist in practice, Copeland and Antikarov (2001). In the absence of a naturally eligible traded security, the advice is to derive the necessary distributional properties of the loss, particularly its volatility, using subjective measures just as would be done when performing a project evaluation based on expected future cash flows. We will assume that whenever a naturally eligible twinning security is not available, then we are able to derive the necessary distributional properties from subjective measures. Whatever approach informs the distributional properties over time of loss arising when a security breach occurs, we will assume in keeping with options theory that the volatility of losses increases proportionately with time.

The existence of a traded twinning security is crucial to deriving the fair price of the option, or in our terminology to derive the price of the insurance premium or the initial investment outlay to limit the upper loss to Q. The financial options formulation defines a dynamic portfolio composed of riskless assets, B and the known portfolio of traded assets, S that can be constructed which exactly mirrors the distribution of value of F: this is shown in figure (3).

Our analysis follows the standard theory of the binomial option pricing model, see Dubofsky and Miller jr (2003), Hull (2003). In this formulation, the value of F equals a linear combination of B and S where the dynamic portfolio is composed of m units of S and minus the amount of B; alternatively, the riskless asset B is conceived as the sum of one unit of the option F and m units of the twinning security S. At time t = 0:

$$B = mS_0 + F_0 \quad (2)$$

At time t = 1:

$$(1+r)B = mS_1 + F_1 \quad (3)$$

reflecting that the value of the riskless asset increases by the riskless rate r. Since the value of F is determined from the maximum function (1) and the value of S is derived from the known distributional properties of the twinning security, the binomial formulation enables the two unknowns, the number of units of S and the amount of the riskless asset B in the dynamic portfolio to be derived:

$$m = \frac{F_1^+ - F_1^-}{S_1^+ - S_1^-}$$

$$B = \frac{mS_1^- - F_1^-}{1+r}$$

From the values of m and B, the value of F_0 can be derived using equation (2). Equivalently, it can be established that the value of F_0 is expected value of F_1 using risk neutral probabilities, q and 1-q, discounted at the riskless rate:

$$F_o = \frac{\hat{E}[F_1^s]}{1+r} = \frac{qF_1^+ + (1-q)F_1^-}{1+r}$$

where the risk neutral probabilities are given by:

$$q = \frac{(1+r)S_0 - S_1^-}{S_1^+ - S_1^-}$$

The value of F_0 is the expected value of F derived using risk neutral probabilities rather than those of the actual probability distribution and discounted at the riskless rate. This result has been obtained for a single time period t = 1. Repeating the representation depicted in figure (2) with n stages forms a replicating binomial lattice that facilitates the extension of the analysis to cover a longer time period up to $t = \tau$, with $h = \tau/n$. By applying the theory of financial options theory, the value of F_0 can be shown to be the cumulative probability distribution of the binomial distribution, which can be approximated using the log-normal approximation to:

$$F_0 = L_0 N(x) - Q(1+r)^{-\tau} N(x - \sigma\sqrt{\tau}) \quad (4)$$

where N(.) denotes the cumulative Normal distribution with:

$$x = \frac{\ln\left(L_0/Q(1+r)^{-\tau}\right)}{\sigma\sqrt{\tau}} + 0.5\sigma\sqrt{\tau}$$

and:

$$\sigma = \frac{\log_e(u)}{\sqrt{h}}$$

$$u = \frac{S_1^+}{S_0} \text{ and } \frac{1}{u} = \frac{S_1^-}{S_0}$$

The variable F_0 denotes the price of an option that enables the holder to limit the maximum loss to Q incurred by a security breach at the time $t = \tau$. In the language of financial options theory, F_0 is the value of a European style call option with exercise or strike price equal to Q and time to expiration $t = \tau$. From equation (4), the variable F_0 is dependent on the expected loss L_0 at t=0 if a breach occurs, the level of protection Q, the time when the breach occurs, τ and the volatility of the losses, σ. Without loss of generality, the value of the investment which restricts the maximum loss to Q at the time t when a breach occurs is:

$$F_0 = F_0(Q, t) \quad (5)$$

The breach in security can take place not only at any time from now until the planning horizon but it is also possible that there are multiple security breaches during that period. To build this possibility into the model, we now need to focus on the other half of figure (1) that was not considered by the preceding analysis, which is displayed in figure (3). By applying the financial options theory as before, we can deduce the risk neutral probabilities for this case to be equal to the actual probabilities because the probability of a security breach is very small. If we denote by q' as the risk neutral probability of a breach and (1-q') as the risk neutral probability of no breach, then the current value of S is given by the expected value of S at time $t = \Delta t$ derived using risk neutral probabilities discounted at the riskless rate r:

$$S_0 = \frac{E\left[S_1^s\right]}{(1+r\,\Delta t)} = \frac{(1-q')S_1^+ + q'S_1^-}{(1+r\,\Delta t)}$$

$$= \frac{q'S_1^-}{(1+r\,\Delta t)}$$

Since $\lambda\,\Delta t$ is very small, then approximately:

$$\lambda\,\Delta t = \frac{S_0}{S_1^-}$$

and:

$$q' = \lambda \Delta t \times (1 + r \Delta t) = \lambda \Delta t$$

We may now treat the probabilities of a breach and of a no breach, that is $\lambda \Delta t$ and $(1-\lambda \Delta t)$ respectively as if they were risk neutral probabilities. Replicating the binomial lattice for breach and no breach in figure (1) to two successive time periods, see figure (4), the probability of a breach at stage $t = 2 \Delta t$ is given by:

$$\lambda \Delta t \times (1 - \lambda \Delta t) + (\lambda \Delta t)^2 = \lambda \Delta t$$

More generally, the probability of a breach at stage n, preceded by j breaches and (n-j-1) no breaches in a specified order is:

$$(\lambda \Delta t)^j (1 - \lambda \Delta t)^{n-1-j} \lambda \Delta t$$

The probability of a breach at stage n, preceded by j breaches and (n-j-1) no breaches in any order is:

$$\binom{n-1}{j} (\lambda \Delta t)^j (1 - \lambda \Delta t)^{n-1-j} \lambda \Delta t$$

The probability of a breach at stage n is:

$$\sum_{j=0}^{n-1} \binom{n-1}{j} (\lambda \Delta t)^j (1 - \lambda \Delta t)^{n-1-j} \lambda \Delta t = \lambda \Delta t \qquad (6)$$

We can now combine the two halves of model together. The function $F_0(Q, t)$ from equation (5) represents the initial outlay cost of obtaining protection against an upper limit of Q at time $t = n \Delta t$ assuming that a security breach occurs at that time. The expected initial outlay cost of obtaining protection against an upper limit of Q at time $t = n \Delta t$ is:

$$\lambda \Delta t \times F_0(Q, t)$$

The expected initial outlay cost of obtaining protection against an upper limit of Q over the planning horizon from t = 0 until t = T is:

$$G_0 = \int_0^T \lambda F_0(Q, t) dt \qquad (7)$$

The function given by equation (7) does not have a closed analytical form because it involves an integral of the cumulative Normal distribution function so its properties can only be inferred through numerical analysis.

Numerical Analysis on the Model

Initially, we will investigate the variations in the value of F_0 for a various parametric values and then proceed to examine the expected initial outlay cost of obtaining protection over the specified planning horizon.

The principal parameters informing the variations in the value of F_0 are described by the right hand side of equation (4): these are the level of protection required by the organisation, the time measured in years at which the security breach occurs and the variability of the loss incurred by a breach in security. All the graphs displaying the change in the value of F_0 are measured against the time metric. Deriving numerical values of F_0 requires information on the expected loss L_0 and the riskless discount rate r; we will set $L_0 = 120$ and r = 5%.

Figure (5) displays the variations in F_0 arising from changes in the level of protection and time given that value of $\sigma = 2$. The graphs for different levels of protection increase monotonically towards an asymptotic value equalling the expected loss. This is because of the nature of the financial options model that assumes that the distribution of the underlying asset becomes increasingly more variable over time and this increased variability makes the cost of protection greater. In our model, since the value at risk follows a distribution whose volatility increases proportionately with time, we would expect the cost of protection to increase as the date of a possible security breach approaches the specified planning horizon. Figure (5) also shows that the cost of obtaining protection against a security breach starts at a positive level when, from equation (1), max $\{L-Q, 0\} > 0$. The critical role of the variability of possible losses incurred by a security breach can be observed from figure (6) that graphs the relationship between F_0 and different values of σ, which shows that the cost of protection increases with the variability of losses incurred from a security breach. With other relevant factors remaining constant, the probability of obtaining a loss exceeding the protection limit in the presence of a security breach increases monotonically with increases in the variability of those losses as measured by the standard deviation.

Figures (5) and (6) reveal the variations in the cost of providing protection against losses incurred through a security breach above an upper limit given that a breach in security takes place at some specified time, t. Critically, we need to know the cost of protection over a specified planning horizon from $t = 0$ until $t = T$, and this is defined by equation (7). For the purposes of understanding the variations in the cost of providing protection over the specified planning horizon, we will eliminate the role of the probability of an effective breach per unit of time, λ from that equation and instead consider:

$$H_0 = \int_0^T F_0(Q, t)dt \qquad (8)$$

Figure (7) displays the variations in H_0 arising from changes in the level of protection and in the time horizon given the value of $\sigma = 2$. Again, the cost of providing protection over the specified planning horizon is related positively with the level of protection afforded, with greater levels of protection defined by lower values of Q requiring greater costs of providing that protection. The various curves show that the graphs of the cost of protection increase with the planning horizon and tend in the long term to become parallel to a linear function of time passing through the origin with slope equalling the expected loss. This is only of theoretical interest. In the real world insurance is normally renewed every year and we would expect firms to upgrade their information security systems at regular intervals, possibly once per year.

In figure (8), the profiles of the cost of protection due to changes in the planning horizon and the variability of the losses incurred from security breaches show that the cost of providing protection are positively influenced by the variability of losses. The cost of providing protection is less when the variability in the losses sustained from a security breach is less severe.

Extensions to the Model

To obtain the cost of providing protection over a certain planning horizon is specified by G_0 defined by equation (7). The variable G_0 is a linear function of the expected number of breaches per unit of time, λ. However, the investment outlay on security protection may not only impact on the level of losses incurred through a breach but may also reduce the probability of a breach taking place. In the investment model of security breaches by Gordon and Loeb (2002), the authors apply the following function to relate the way that changes in the investment level affect the probability of an effective security breach:

$$\lambda = \lambda_0^{(\alpha G_0 + 1)} \quad (9)$$

where λ_0 represents the probability of a security breach per unit of time when no investment in security is made and α denotes the intensity of the investment. The value of α will be greater when the investment into security using a particular technology is more effective and this will result in a lower investment cost since the value of λ_0 is less than one. By combining equations (7), (8) and (9), the expected initial outlay cost for obtaining a certain degree of protection over a particular planning horizon when the initial investment level influences the probability of a security breach is given by:

$$G_0 = \lambda_0^{(\alpha G_0 + 1)} \times H_0 \quad (10)$$

The profiles of G_0 are exhibited in figure (9) for values of λ_0 and different levels of investment intensity given that H_0 is set at the constant value of 100. The linear line passing through the origin represents the benchmark case when $\alpha = 0$. It can be seen that greater levels of investment intensity reduce the level of investment required to a varying extent.

Discussion

This paper sets out to develop a model of information security investment grounded in the concepts of financial options and probability theory. This insurance style model, which is founded on the principle of economic rationality and fair pricing, relates the required level of investment to the variability in the losses sustained from a security breach and the probability of a breach actually taking place. Its merits lie in the explicit inclusion of the time value of money and the possibility that more than one breach may occur within the specified planning horizon. Unlike competing models that treat the loss incurred through a security breach as a constant, this model adopts the value-at-risk principle by allowing any consequential loss to be random and to follow a distribution with known expected loss and variability. Further, the planning horizon is considered to be a critical decision variable that influences the level of required investment rather instead of an arbitrary number that does not figure in the analysis. Although the preliminary model assumes that investment in security only influences the consequential losses, the model has sufficient flexibility to include the influence of the investment level on the probability of a security breach occurring.

It has to be recognised that the proposed model is founded on economic rationality since the derivation of the fair price for the financial option is based on rationale economic agents. This assumption of rationality is reasonable in many cases. In contrast, prospect theory, which is grounded on the behaviour of economic agents, purports that organisations may adopt a risk loving stance in the presence of losses through preferring uncertain higher losses than a certain lower losses, see Kahneman

and Tversky (1979), Kahneman, et al (1991). This attitude implies that decision makers would prefer to run the risk of incurring a large improbable loss such as sustained from a breach in security rather than investing in an information security technology that may not bring any tangible benefits at all. Further, the degree of risk aversion changes over time owing to the currency of events. A recent security breach is likely to raise the degree of risk aversion shared by decision makers whereas the absence of any recent security breaches is likely to instil a cavalier approach to deciding on the appropriate level of investment security.

The proposed model makes no distinction between the various forms of investment opportunities for reducing the threat of security breaches. These alternative ways of thwarting the danger of breaches in security can range from a combination of hardware and software ICT solutions to strategies involving the development of people and organisational capabilities. These various measures will differ not only in their relative effectiveness in deterring general criminal and hostile intrusions but also in their capacity to thwart specific threats such as impersonations, fraud or insider abuse. Some of these measures will require the funding of upfront investment outlays whereas others such as training will demand funding on a continual basis. Further, the model does not discriminate between various sources of loss, such as reputation losses, compensation, the deficit from the degradation of information integrity and the resources required to return the system to normality, but treats them collectively and measurable. Empirical research needs to focally address these areas so that the relative effectiveness of the various deterrence systems on offer as well as their efficacy in the face of specific threats can be identified and precisely measured.

The model has demonstrated its flexibility for dealing with both mitigating losses and probability reduction of a security breach. The original insurance style representation applies the concept of financial options to the investment in information security that limits the maximum amount at risk in order to derive the appropriate level of initial investment for a given degree of protection, a required planning horizon and variability in the loss distribution. This model was then adapted to permit the initial level of investment to bear on the probability of a security breach actually taking place. Clearly, further research on this basic model needs to done in order to identify the best mix of investment for mitigating loss and intruder probability reduction as well as considering more sophisticated forms of value-at-risk.

References

TE Copeland and V Antikarov (2001)
Real Options
Texere Publishing Ltd

K Dowd (2002)
An Introduction to Market Risk Measurement
John Wiley & Sons Ltd

DA Dubofsky and TW Miller jr (2003)
Derivatives: Valuation and Risk Management
Oxford University Press

A Gordon and MP Loeb (2002)
The Economics of Information Security Investment
ACM Transactions on Information Security, 5(4) 438-457

JC Hull (2003)
Options, Futures and Other Derivatives
Prentice-Hall

D Kahneman and A Tversky (1979)
Prospect theory: an analysis of decision making under risk
Econometrica 47, 263-291

D Kahneman, J Knetschen and R Thaler (1991)
Anomalies: the endowment effect, loss aversion and status quo bias
Journal of Economic Perspectives 1, 193-206

L Trigeorgis (2000)
Real Options: Managerial Flexibility and Strategy in Resource Allocation
MIT press

Figure 1
Representation of the Basic Model

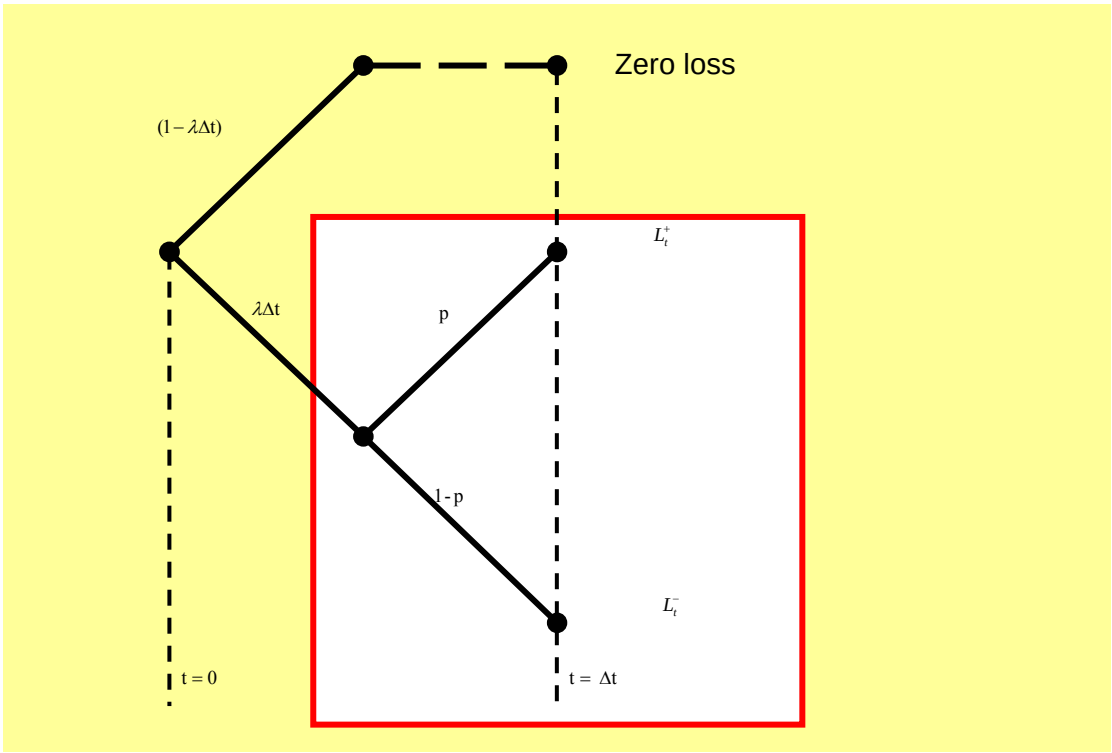

Figure 2
Representation of the value of investment F required conferring a degree of protection against a breach that limits the maximum loss to Q

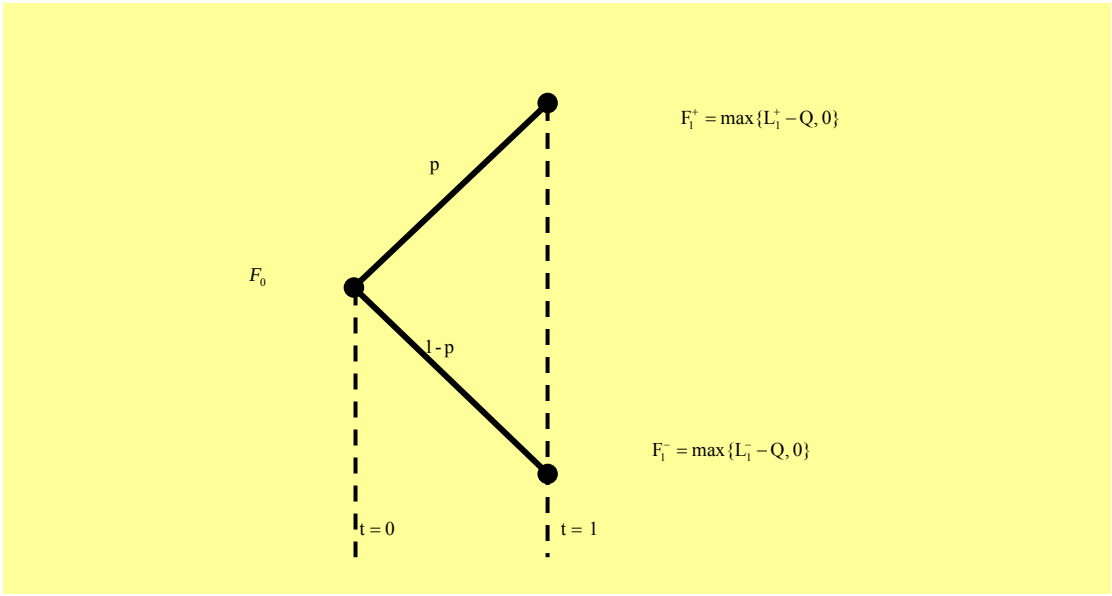

Figure 3
Construction of dynamic portfolio composed of the twinning security and riskless security that replicates the option's outcome profile

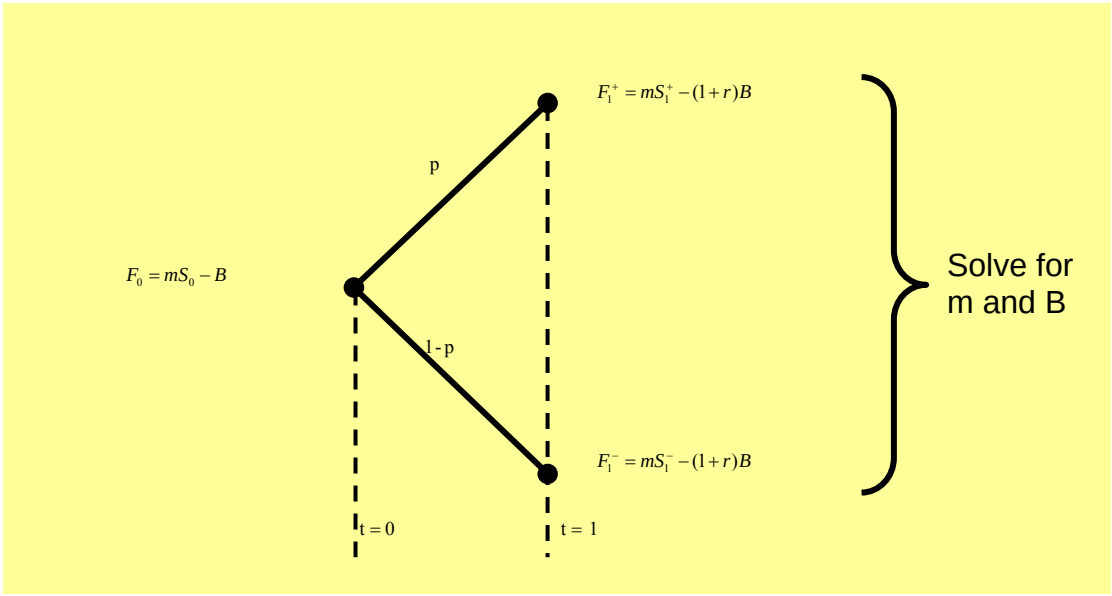

Figure 4
Binomial lattice for the case of breach and no breach for two stages

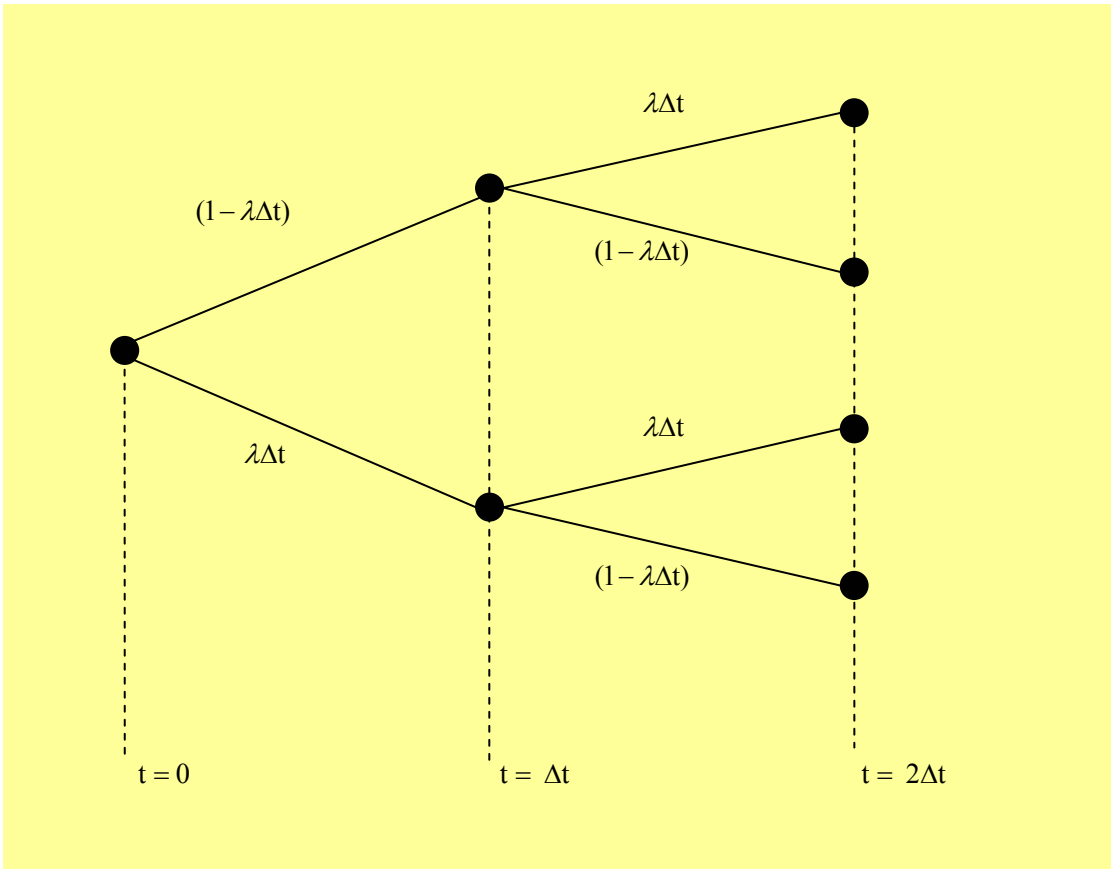

Figure 5
Profile of F_0 against time defined in years with
$L_0 = 120$
$r = 5\%$
$\sigma = 2$

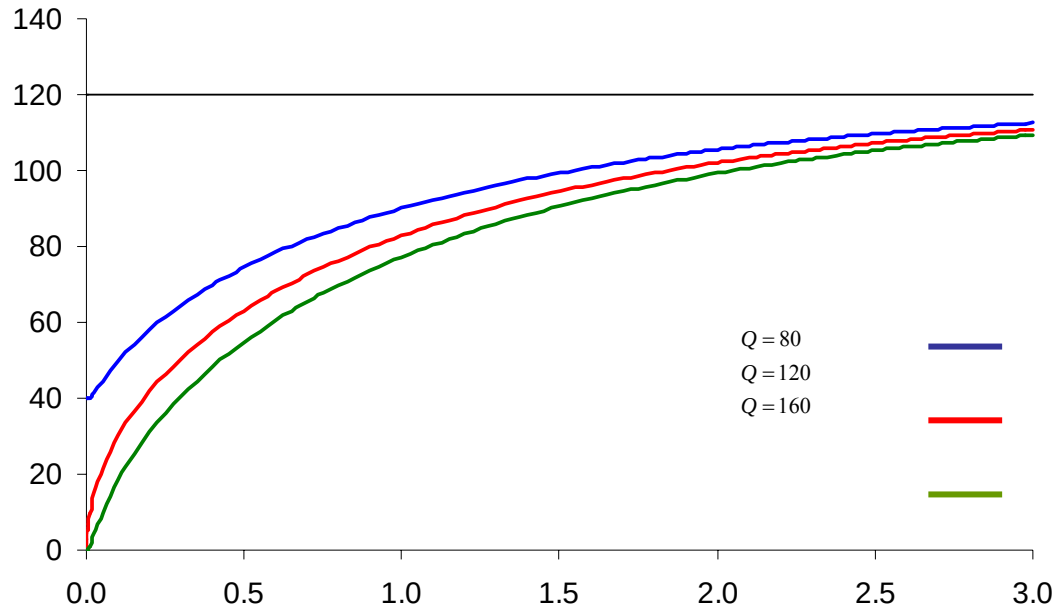

Figure 6
Profile of F_0 against time defined in years with
$L_0 = 120$
$r = 5\%$
$Q = 120$

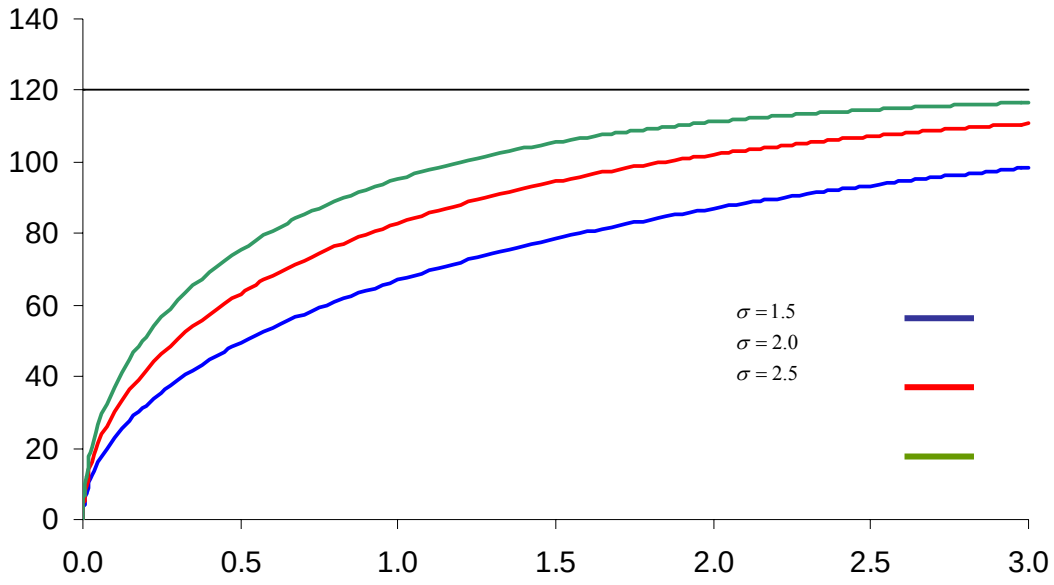

Figure 7
Profile of H_0 against time defined in years with
$L_0 = 120$
$r = 5\%$
$\sigma = 2$

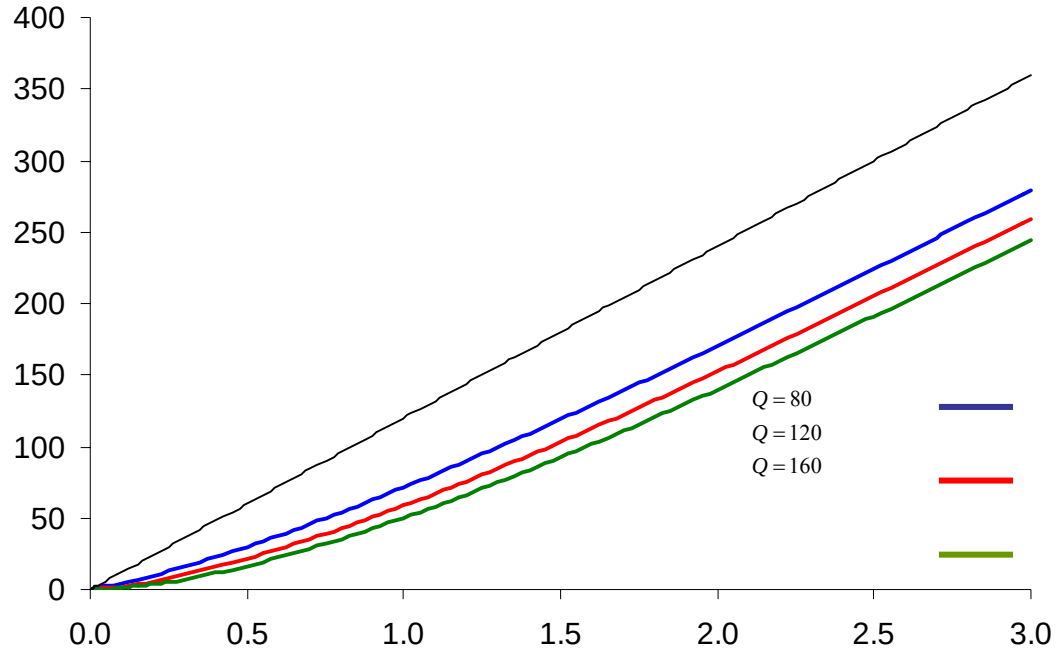

Figure 8
Profile of H_0 against time defined in years with
$L_0 = 120$
$r = 5\%$
$Q = 120$

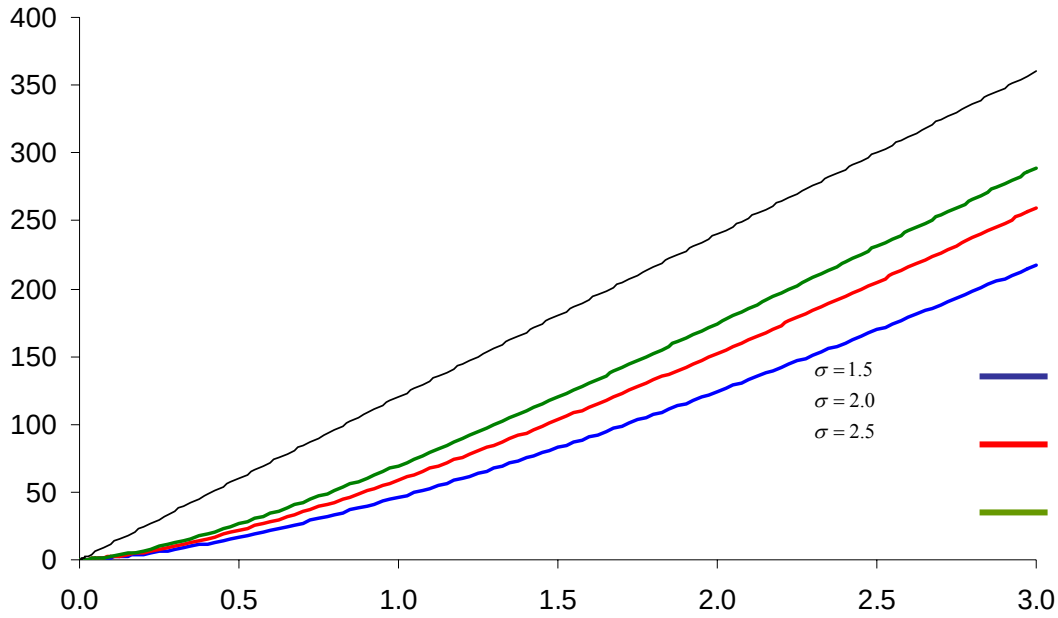

Figure 9
Profile of G_0 against values of λ_0
$H_0 = 100$