Real options in system design: a methodological compromise and its implications

DRAFT

Konstantinos Kalligeros* Richard de Neufville
Massachusetts Institute of Technology, Cambridge MA 02139

10th Annual International Conference On Real Options: Theory Meets Practice
New York City, USA, June 14-17, 2006

Abstract: This paper proposes a simulation methodology for the evaluation of flexibility in certain classes of engineering systems for the purpose of design. The intended audience for this process is the engineering community, whose priority and goal is not the absolute valuation of the system for diversified investors in the market, but the relative, ordinal valuation of the alternative engineering designs from a flexibility perspective. The proposed methodology sacrifices some technical rigor in the options valuation in order to be compatible with current evaluation practices in engineering organizations. We propose a step process that involves simulation of the exogenous uncertainties, valuation of one design alternative per the organization's practices, and finally, relative valuation of the other design alternatives and options based on the organization's inferred risk aversion. We use certainty-equivalence arguments for valuation and avoid risk-neutral dynamics. The method is accurate for valuing plain vanilla options, but strictly speaking, incorrect for options on multiple assets. The method is applied to a case of alternative uses of a land parcel and mixed-use development, which demonstrates the design implications from using this approach versus a traditional, rigorous real options analysis.

Introduction

Engineering systems are traditionally designed to fixed specifications, even when uncertainty regarding the system's development or operating environment is acknowledged and examined in post-design sensitivity analyses. However, the idea of deliberately designing engineering systems so that they enable managerial flexibility is slowly gaining momentum in practice. Moreover, this flexibility is often conceptualized as a collection of real options, i.e., the right, but not an obligation to change, expand, shrink or otherwise evolve a system. Even so, however, the use of real options has not managed to penetrate into the practice of system design evaluation.

There have been academic efforts recently for the optimization of engineering systems based on the option value they create. Applications include phased deployment of communication satellite constellations under demand uncertainty (Chaize et al, 2003); decisions on component commonality

* Room E40-252, Massachusetts Institute of Technology, Cambridge MA 02139. Email kkall@alum.mit.edu

between two aircraft of the same family (Willcox and Markish, 2004); building design under rent and space utilization uncertainty (Zhao and Tseng, 2004; Kalligeros and de Weck, 2004; Greden et al, 2005; and de Neufville et al., 2005). These applications involve essentially the valuation of the options associated with a particular design and the external optimization of this design.

Some of these case studies focus on the engineering problem without attempting to provide a rigorous option analysis: their contribution is explicitly about how engineering systems can be designed to enable flexibility, and the link to the options literature is found only at a conceptual level. When a rigorous options valuation is attempted in engineering applications, it usually runs into theoretical errors, practical difficulties in application, and logical arguments that do not convince engineering audiences. In short, real options applications for engineering design have so far mostly been either intuitive and deficient from an economics perspective, or correct but simplistic and unconvincing for the engineering community.

We postulate that the main reasons for such low penetration lie in the attempt to transfer a "pure" and positive economic theory to the "production floor" or engineering design. Adopting the theory requires understanding the assumptions behind it, which in turn, require a level of economic literacy found in very few engineers. Even when the assumptions are understood, they often do not apply to the situation at hand. When the assumptions that justify an options analysis of the flexibility in a system are reasonable, the application often requires conceptual leaps of faith (e.g., using the risk-free rate for discounting) and changes in well-established valuation practices within the organization. Experience has shown that improving the financial sophistication of engineering organizations has been slow and difficult.

The audience for these applications, i.e., the engineering organization, is far more interested in a process for the selection of a design over another one, than a process for the calculation of the exact value of a system from the viewpoint of a diversified investor. To address this audience, we propose a compromise between a rigorous option valuation of the flexibility in a system, and the current practice of discounted-cash-flow analysis, simulation of the important uncertainties and sensitivity analysis of a system's financial performance.

We suggest that alternative engineering designs define operational states for the engineering organization. These operational states have intrinsic value (due to the cash flow they generate) as well as value from the managerial flexibility they enable. This flexibility can be expressed as real options to choose and transition to other operational states. The valuation of the initiating operational state amounts to the valuation of its intrinsic cash flows and these real options. In the proposed methodology, a single representative operational state is valued using the developing organization's standard practices. This valuation implicitly provides the organization's risk tolerance towards the uncertainty in the system's performance. This risk tolerance is implicitly used for the valuation of alternative system designs (operational states) and options to reconfigure the system in the future (transitions between states). The valuation of options and operational states is achieved using an original simulation method which builds on the stratified state aggregation method (Barraquand and Martineau 1995) and the Generalized Multi-Period Option Pricing model (Arnold and Crack 2004).

Figure 1 shows graphically the "placement" of the proposed valuation scheme along four dimensions: state-space modeling, uncertainty modeling, decision rules and valuation. We use Monte-Carlo simulation to model the state space, as it is far more appealing to the engineering community than both lattices and continuous-time formulations. For the same reason, we propose simulating the external uncertainties to the engineering system and deducing the value of the system along each simulation event; alternative approaches involve identifying a portfolio of securities that is believed to perfectly track the value of the system, or explicitly assuming the value of the system follows a Geometric Brownian Motion (e.g., see Copeland and Antikarov 2000). We have found that both these alternatives are not convincing to the engineering community or imply absurd assumptions for many technical systems. The decision rules in our method are calculated recursively using stratified state aggregation (Barraquand and Martineau 1995). An alternative for this could have been the arbitrary specification of decision rules or their parameterization (e.g., see Longstaff and Schwarz 2000 or Andersen 2000). Finally, discounting of future values is performed across alternative designs and options so that the price of risk is retained constant. We achieve this using the Generalized Multi-Period Option Pricing model (see Arnold and Crack 2004) and avoiding the need to simulate risk-neutral paths.

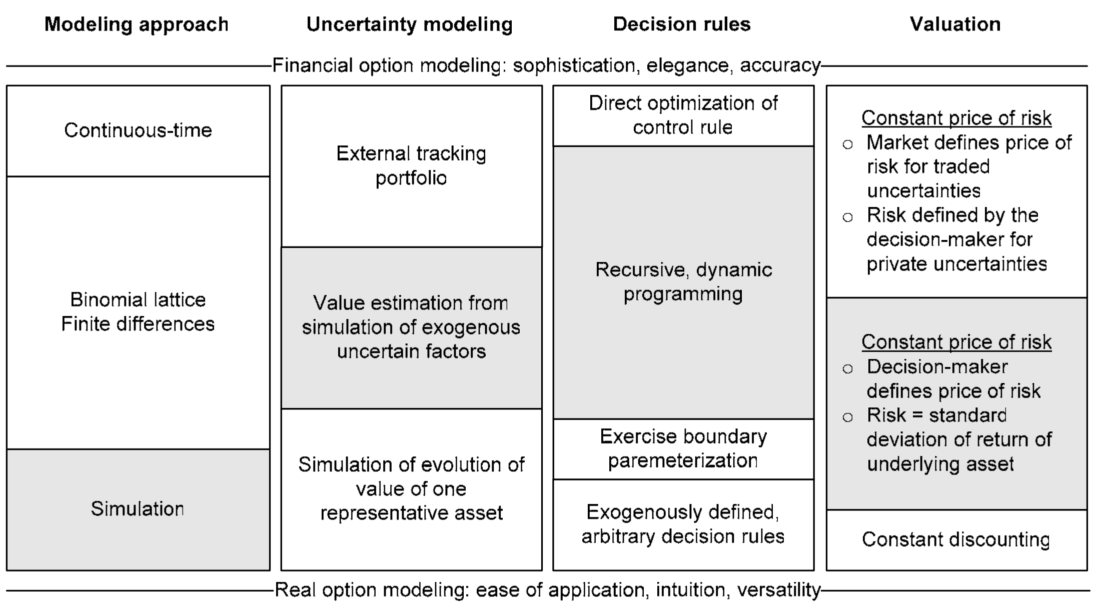

Figure 1: "Placement" of the proposed option evaluation methodology

We believe this valuation method to be a step towards bridging the "communication" gap between real option theorists and engineering practitioners, that has prevented the options method from reaching practical application. We believe that this method can enable a big improvement in engineering practice with a minimum of departure from a typical organization's evaluation culture and practices, and with a minimum of compromise from a theoretically correct solution.

The paper proceeds as follows: the next section describes in more detail the process for estimating the relative value of alternative system designs and the algorithm for flexibility valuation. The efficiency of the algorithm is compared to the Barraquand & Martineau (B&M) method and a binomial lattice solution with Richardson extrapolation for a plain vanilla call option. Next, we discuss the theoretical compromises the method entails compared to a rigorous options analysis regarding the value of a flexible system. The final section demonstrates the application of the method for a real estate development case, and compares the difference in design recommendation between a rigorous analysis and this method.

Options evaluation process

The method proposed for comparing alternative designs in terms of the managerial flexibility they enable is shown in steps below.

Step I — Define one or more operational states, i.e., alternative designs of the system. An operational state may include reversible flexibility. Estimate free cash flows (i.e., cash flows from operations, less the expense necessary to sustain these operations and their expected growth) as a function of time, the design variables and the uncertainties for each operational state. The entire organization can only be in a single operational state at a time.

Step II — Uncertainty simulation

- Simulate all exogenous uncertainties where learning cannot be affected by actions of the organization, market-traded or not.
- Express the simulation parameters of the endogenous uncertainties (i.e., those for which learning can be affected by the organization) as functions of the system design in every operational state.
- The uncertainties are simulated using their real stochastic processes (not risk-neutral). This enables the user of the process to use objective and subjective estimates for public (market-traded) and private uncertainties alike.

Step III — Value one "representative" operational state using established valuation processes for each path of simulated uncertainties. The simulation algorithm explained next can be used for this.

Step IV — Value all other operational states using a discount rate adjusted for the relative risk between each operational state and the representative operational state.

Step V Value all timing and choice options to transition between operational states the same way. The value of flexibility in each operational state is the sum of the intrinsic value of that state and the value of its options.

The process is general and can be applied to a variety of technical systems, even though it is more suitable for certain types of situations. As described in steps IV and V, the process requires the valuation of all operational states (alternative designs) as well as the possible transitions between them. Consequently, its usefulness is limited computationally by the number of operational states. Also, the process assumes that there is no strong path-dependency in the model, i.e., the cost and time lag for transition between operational states and the cash flows of operational states themselves are only functions of time, design variables and the uncertain factors.

The next section describes the valuation method for steps IV and V in the process above.

Evaluation algorithm

Steps IV and V in the process above involve the valuation of operational states (alternative designs) and options based on the value of a representative state. This is achieved with the simulation algorithm described below. The algorithm is based on the original stratified state aggregation method Barraquand & Martineau (B&M, 1995) and the Generalized Multi-Period Option Pricing model (see Arnold and Crack 2004). It is explained here for the option to transition to an operational state of known value.

Suppose that the designer knows the value of a system (operational state) A for each time on each simulated path of uncertain factors using steps I to III (or by directly observing this price in the market, if possible). The question is to value the option to obtain this state. Table 1 explains the simulation algorithm for this.

If the simulation in step 1 in Table 1 uses the risk-neutral process for the underlying asset then the algorithm is essentially identical to the Barraquand & Martineau (1995) simulation method (B&M). Then, the discounting of the expected future value of keeping the option alive (steps 5.1 and 5.2 in the table) may be done using the risk-free rate. Here however, the underlying asset (i.e., the operational state) is simulated using the real probabilities. This is because we assumed that this value is calculated from the simulated value of all the important uncertainties in the performance of the system, as estimated by the engineering organization.[1] Therefore, the risk-free rate cannot be used for discounting the future expected value of the option. Assuming the organization is generally risk-averse, it should also not use the discount rate it used for valuing operational state A. To determine the correct discount rate for each bin of paths

[1] The added complication may seem redundant to an economist, however, modeling risk-neutral dynamics is one of the most important factors for the limited penetration of real options methods in practice.

(steps 5.1 and 5.2), we use an approximation of the exact Generalized Multi-Period Option Pricing model (Arnold and Crack 2004).

Table 1: Option valuation algorithm

#	ACTION
1	Simulate K paths for the underlying asset.
2	For terminal time T ($j=0$) find the value of the option for each path
3	Go one time period δt backwards ($j=j+1$), and find the value of immediate exercise of the option for each path.
4	Divide the K paths into M groups (bins) (if $j=T/\delta t$ then all paths meet at their current value and $M=1$)
5	For each group of paths (bin): 1. Calculate the expected value of the option one period in the future (at time $T-(j-1)\delta t$) as the average for the paths in the current bin. 2. Discount this expected value appropriately to get the average present value of keeping the option alive for this group of paths 3. Compare this with intrinsic option value for each path/group of paths (from step 3), and keep greater of the two.
6	If $j<T/\delta t$ then go to step 3, else endAt time 0, the option value (all paths) is the maximum of immediate exercise or the present value of the value of waiting.

Suppose the organization uses an annualized, continuously-compounded discount rate R_A for computing the value of operational state A, V_A. This rate is usually imposed to the designer, and may represent the weighted average cost of capital for the entire organization if operational state A is representative of the risks to which the organization is exposed. This means that the organization already knows point α in Figure 2. The standard deviation of value σ_A for this representative state is also known for each bin. Assume the following:

1. The organization can borrow and lend money at an approximately risk-free rate r_f (annualized, continuously-compounded).
2. The developer requires compensation for taking risk, and perceives risk as standard deviation of returns.
3. The relationship between risk and risk premium is linear so that the price of risk is constant and equal to $\lambda=(R_A-r_f)/\sigma_A$, where σ_A is the standard deviation of returns for operational state A in each bin.

With these assumptions the value of the option for each bin can be estimated using the following formula (proof in the appendix):

$$V^t_{OPTonA} = \frac{E[V^{t+\delta t}_{OPTonA}] - (e^{R_A \delta t} - e^{R_f \delta t})V^t_A \dfrac{\sigma_{OPTonA}}{\sigma_A}}{e^{R_f \delta t}}$$

The formula is a theoretically correct approximation when applied between an option and the underlying asset, as they are perfectly correlated between them (therefore from a CAPM perspective, the ratio of betas for the operational state and the option is equal to the ratio of standard deviation of returns). Generally, it is theoretically correct when applied to infer the value of one operational state given the known value of another, if the two are perfectly correlated. It is theoretically incorrect when the values of the two operational states are not perfectly correlated.

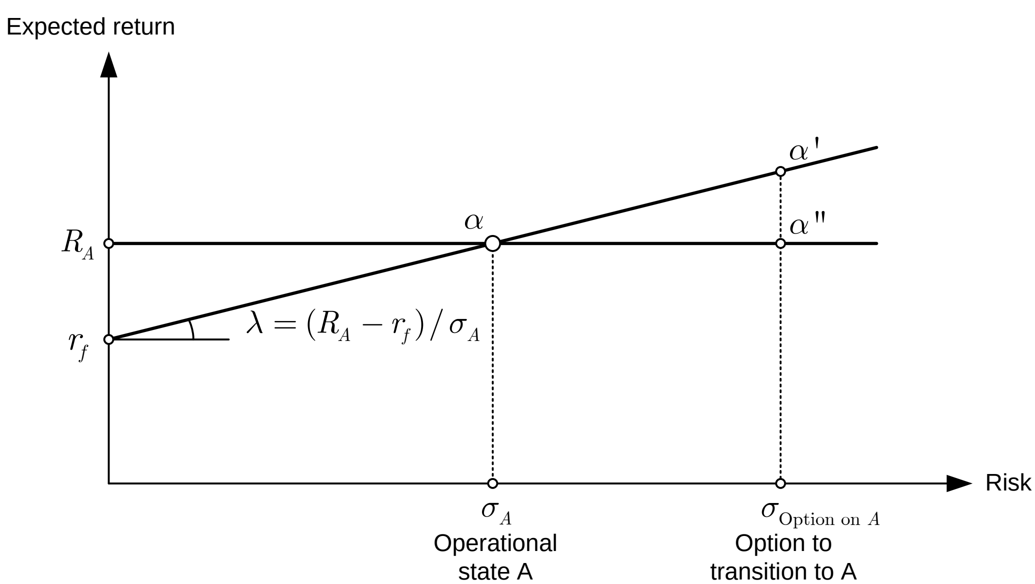

Figure 2: Constant price of risk

Discussion (incomplete)

The compromise

- Method lumps market-traded uncertainties with private ones for all operational states (assets). This is contrary to Smith and Nau (1995), but agrees with Copeland and Antikarov (2000) and current practice.
- Problem with getting exercise boundaries
- Theoretical problems (see Borison 2003)

- For perfectly correlated assets (e.g,. underlying and option), where the underlying is market-traded, the method will produce correct and consistent valuation (i.e., the actual value of the option, if it was traded).
- For perfectly correlated assets (e.g,. underlying and option), where the underlying is not market-traded, will produce consistent but incorrect valuation (i.e., the value of the option if the underlying had been traded). This is justified by the “Market Asset Disclaimer” argument, see Copeland and Antikarov (2000).
- For imperfectly correlated assets, where the reference asset is market-traded, the method will produce an inconsistent and incorrect valuation.

The benefits

- Communication: uses simulation exactly as it is used in engineering sensitivity analyses.
- The simulation method is very versatile to include development time lags and weak path-dependence.
- In terms of valuation, the process is an improvement over the original “Market Asset Disclaimer” argument because
  - Process allows ad-hoc modeling of operational flexibility within each operational state (design)
  - Process does not infer log-normal price process for the system’s value, and is not limited to mean-variance valuation
  - Process is not limited to a single price of risk for all states of the world
- Since the engineering organization is restricted by the model to lie in a single operational state at each time, and since any correlation between uncertainties is reflected in the valuation of the operational states, it follows that the proposed process will give a mean-variance optimal order of optimal designs, according to the implied risk aversion of the organization.

Application: mixed-use development (incomplete)

The methodology is applied to the design decisions of mixed-use real estate development. The goal of this case study is to demonstrate the use of the methodology and to examine if it leads to the same design decisions as a rigorous analysis. The emphasis is not in the actual valuation of the alternative options available to a developer, but the development decisions.

The scenario is as follows: the land-owner holds the right, but not the obligation to develop the land as either commercial or residential use (COM and RES) respectively within a time horizon of $T = 2$ years. The total density of development is constant (e.g., 100 units), but the developer can choose any number of commercial or residential units for the development. The rent process y_{RES}, y_{COM} for each use is considered as the underlying uncertainty, following a geometric Brownian motion; the instantaneous income return on the properties is assumed to follow a bivariate normal distribution with volatilities $\sigma_{RES} = \sigma_{COM} = 0.2$ and correlation $\rho = 0.2$. The total rent exhibits diminishing returns, as $Y = q_{RES}^{a_{RES}} y_{RES} + q_{COM}^{a_{COM}} y_{COM} = q_{RES}^{0.91} y_{RES} + q_{COM} y_{COM}$, where q_{RES}, q_{COM} are number of residential and commercial units in each development plan, respectively: they are the developer's design variables (figure 3). The initial levels for residential and commercial rents are $y_{RES} = 1.48, y_{COM} = 1.00$.

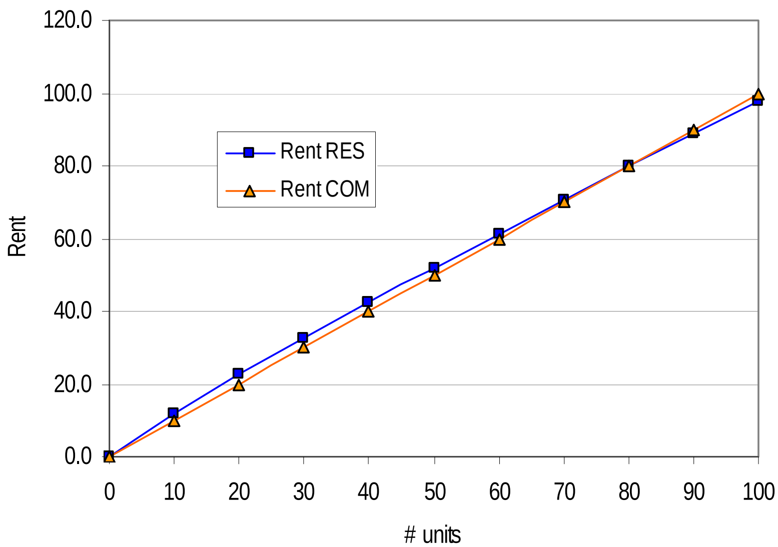

Figure 3: Diminishing returns in rent

The example is based on Geltner, Riddiough and Stojanovic (1996) and Childs, Riddiough and Triantis (1996). Geltner et al. use the value process of the alternative uses as the underlying assets, and do not consider the option to redevelop between uses; Childs et al model the rent process and consider this option. Moreover, Childs et al. show that if the development is instantaneous, then development happens in a single phase. Since we do not consider construction time lags in this example, we do not need to consider multiple development phases either. We also do not consider the re-development option.

Classic solution

[to be completed]

Solution with proposed methodology

[to be completed]

Conclusions

This paper presents a simulation methodology for the evaluation of flexibility for the purpose of design. Given that the objective of this methodology is not to provide an accurate value of the option from the viewpoint of the diversified investor, but an ordinal, relative valuation of alternative systems for design purposes, we make some simplifying compromises in the theoretical rigor of our approach. The paper demonstrates the numerical efficiency and accuracy of the methodology by benchmarking it against comparable option valuation methods. We also commented on the underlying assumptions and the conditions under which the method yields correct results. Finally, we demonstrated the benefits of the methodology in a case study, where we also show the implications for design decisions of using this approximate method rather than the rigorous analysis.

We hope that these simplifications will improve the engineering community's access to the field of real options and how it may be applied to real systems. At the same time, these simplifications may contribute to the real options literature as a means to penetrate current practice.

References

Andersen, L. A simple approach to the pricing of Bermudan swaptions in the multi-factor LIBOR market model. Journal of Computational Finance 3 (1) 2000, pp. 5-32

Arnold, T. and Crack, T.F. Option pricing in the real world: a generalized binomial model with applications to real options, 2003 Real Options Conference, Montreal, CA

Barraquand, J. and Martineau, D. Numerical Valuation of High Dimensional Multi-Variate American Securities. The Journal of Financial and Quantitative Analysis, 30,3 September 1995.

Borison, A. Real Options Analysis: Where are the Emperor's Clothes? 7th Annual International Conference on Real Options, July 2003

Chang, C. C., Chung, S. L. and Stapleton, R. C. Richardson Extrapolation Techniques for the Pricing of American-style Options, Draft, March 2002.

Childs, P. D. and Riddiough, T. J. and Triantis, A. J. Mixed Uses and the Redevelopment Option, Real Estate Economics 24, 1996

Copeland, T. and Antikarov, V. Real Options, a Practitioner's Guide, TEXERE LLC, 2001, New York, NY 100

de Neufville, R. Real Options: Dealing with Uncertainty in Systems planning and Design. Forthcoming in Special Issue, Integrated Assessment Journal, 2001

de Neufville, R., Scholtes, S. and Wang, T. Valuing Options by Spreadsheet: Parking Garage Case Example, ASCE Journal of Infrastructure Systems, in press, 2005

de Weck, O. and Chaize, M. and de Neufville, R. Enhancing the Economics of Communications Satellites via Orbital reconfigurations and staged Deployment. American Institute of Aeronautics and Astronautics, September 2003

Geltner, D. and Miller, N. Commercial Real Estate Analysis and Investments, South-Western Publishing, ISBN 0-324-13676-5, 2001, 5101 Madison Road, Cincinnati, Ohio 4522

Geltner, D. and Riddiough, T. and Stojanovic, S. Insights on the effect of land use choice: the perpetual

option and the best of two underlying assets, Journal of Urban Economics 39, 1996, pp. 20-50
Grenadier, S. R. Flexibility and tenant mix in real estate projects Journal of Urban Economics 38, 1995 pp. 357-378
Kalligeros K.C. and de Weck O.L.: Flexible Design of Commercial Systems under Market Uncertainty: Framework and Application, 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA-2004-4646, Albany, New York, 2004
Longstaff, F.A. and Schwartz, E.S. Valuing American Options by Simulation: A Simple Least-Squares Approcah. The Review of Financial Studies14 (1) 2001, pp. 113-147
Luenberger, D.G. Investment Science, Oxford University Press, 1998, New York, New York 10016
Markish, J. and Willcox, K. Value-Based Multidisciplinary Techniques for Commercial Aircraft System Design, AIAA Journal 41 (10) 2003-2012
Neely, III, J. E. and de Neufville, R. (2001) Hybrid Real Options Valuation of Risky Product Development Projects. International Journal of Technology, Policy and Management, 1(1) pp. 29-46
Smith, J. E. and McCardle, K.F. Options in the real world: lessons learned in evaluating oil and gas investments Operations Research 47 (1), January-February 1999
Smith, J. E. and Nau, R. F. Valuing Risky Projects: Option Pricing Theory and Decision Analysis Management Science 41 (5), May 1995
Trigeorgis, L. Managerial Flexibility and Strategy in Resource Allocation, MIT Press, 1996 Cambridge, MA
Wang, T. and de Neufville, R. Real options “in” projects. Draft prepared for the 9th Real Options Annual International Conference in Paris, France, June 2005 Paris, FR
Zhao T. and Tseng C.L. Valuing Flexibility in Infrastructure Expansion. Journal of Infrastructure Systems 9 (3), September 2003

Appendix A: algorithm comparison

This appendix shows the accuracy and efficiency of the algorithm from a computational perspective. For demonstration purposes, we value a plain vanilla call option with the following characteristics:

Table 2: Test call option characteristics

$K = 10000$	number of paths
$S_0 = 10$	initial stock price
$R_s = 8\%$	required expected return from stock
$r_f = 3\%$	risk free rate
$\sigma = 30\%$	annual volatility of the stock
$dt = 1$	time increment
$T = 10$	time horizon
$N_B = 200$	number of bins

We use the proposed algorithm which we label “Kalligeros & de Neufville 2006,” and compare it to Barraquand and Martineau (1995). Essentially, the Barraquand and Martineau solution is a special case of the algorithm we present, implemented by imposing the risk-adjusted expected return for the underlying asset to be equal to the risk-free rate. The results from the two methods are compared to a binomial solution. In all runs, we find the binomial solution for two values of the time increment and extrapolate the continuous-time value of the option (e.g., see Chang et al., 2002).

For the parameters in Table 2, we drew 10,000 sample paths for the underlying asset using antithetic variables. The same set of random numbers was used in both Monte-Carlo methods. The results were as follows.

Algorithm	American call option
Kalligeros & de Neufville 2006	4.7046
Barraquand and Martineau 1995	4.9879
Binomial (Richardson extrapolation)	4.6147

Next, we performed 5 analyses to explore the sensitivity of these results. The first two involved varying the volatility of returns for the underlying asset; the third involved varying the expected return of the underlying asset; the fourth examined the sensitivity of the two simulation algorithms to the number of simulated paths and bins; the final analysis involved varying the number of time steps. The results are shown in the following figures.

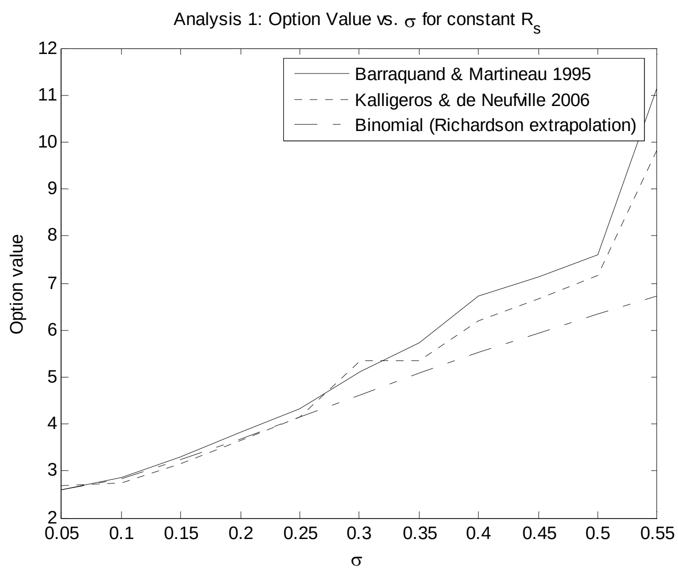

Figure 4: Option value versus volatility (constant expected return)

Figure 4 shows the behavior of both simulation algorithms (again, over the same set of random values) over a range of volatilities. Since the expected return on the underlying asset was kept constant, varying the asset's volatility essentially means we are implementing the algorithm in different economies. The figure shows that the discrepancy and its error increases with increasing volatility for both simulation methods.

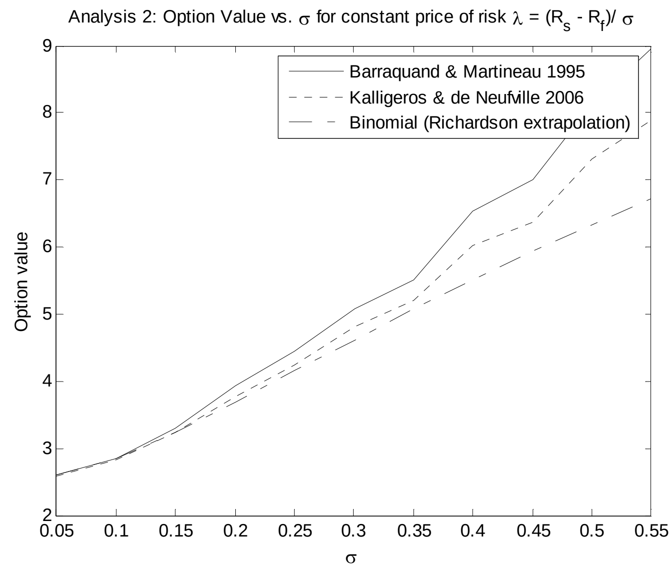

Figure 5: Option value versus volatility (constant price of risk)

Figure 5 repeats the previous analysis, but holds the price of risk constant (at the value implied by the base parameters, $\lambda = 0.1667$. So, a different expected return corresponds to every value of σ. Essentially, Figure 5 shows the performance of the three algorithms for a continuum of assets in the same economy. As before, both simulation methods deviate from the continuous time solution as volatility increases.

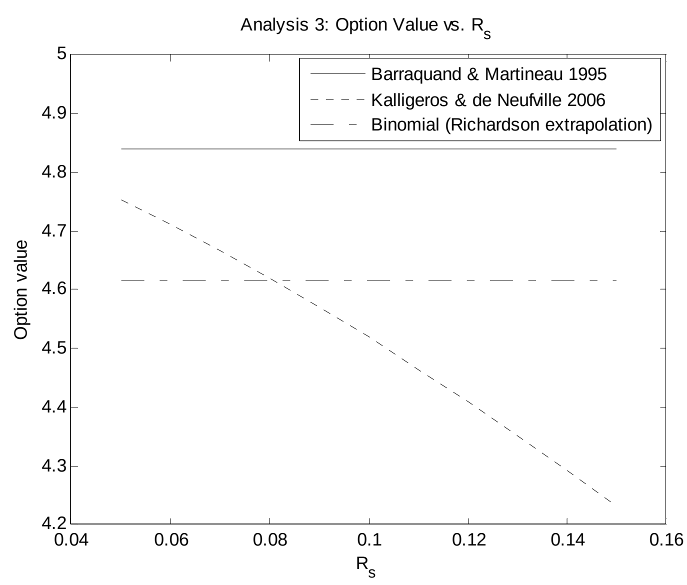

Figure 6: Option value versus expected return

Figure 6 shows the performance as the expected return of the underlying asset varies. In principle, the value of an option is indifferent to the expected return of the underlying asset, so the Figure should show the value of all three methods as flat lines. Since the expected return of the asset does not affect the calculation of either the binomial method or the Barraquand and Martineau algorithm, the corresponding results are indeed indifferent. The values of our algorithm seem to decrease as the expected return increases; this is a numerical drawback of the method.

Figure 7 shows a comparison of the option value for different number of simulated paths and bins at each time period. In both plots, the vertical axis gives the percent error compared to the binomial solution. It is seen that overall, the proposed algorithm gives more accurate results than the Barraquand and Martineau method, that are more robust to changes in the number of bins and simulated paths.

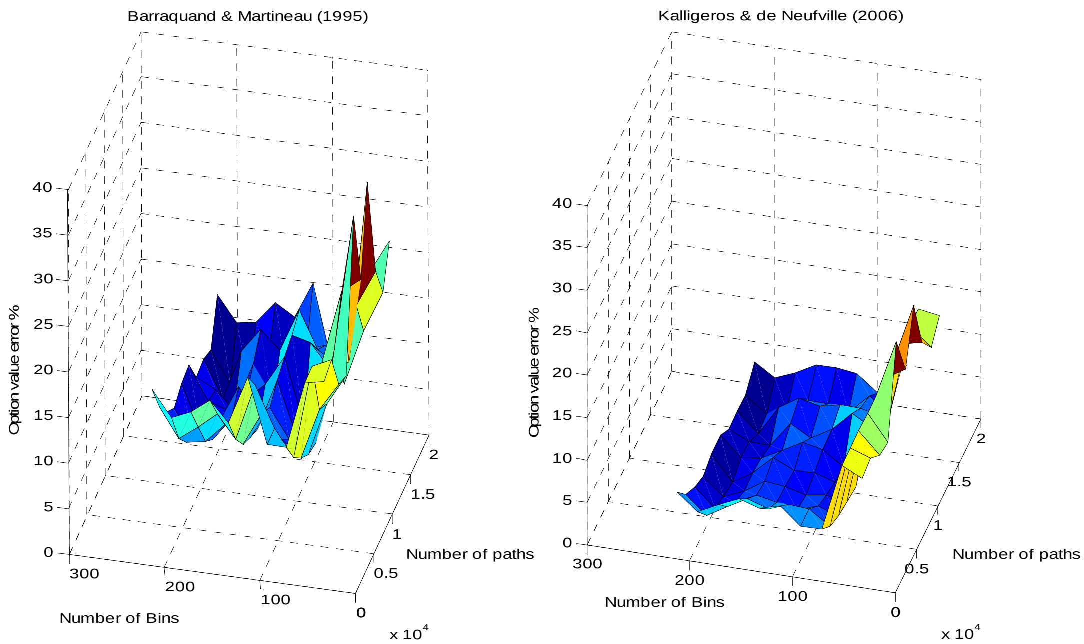

Figure 7: Option value versus number of bins and paths

Finally, Figure 8 shows how value varies with the number of steps for the same time horizon. (The binomial result is extrapolated to $dt = 0$ so it is not shown to vary.) Both simulation methods show the same pattern of oscillation around the solution that is observed with binomial methods.

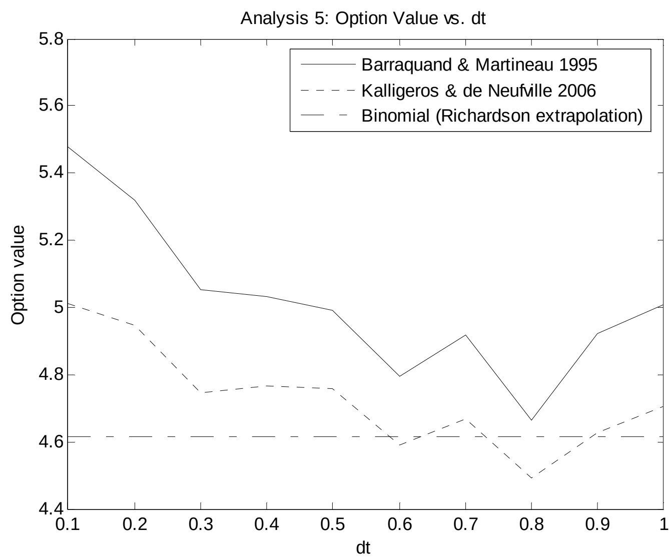

Figure 8: Option value versus time increment