

APPENDICES

TABLE OF CONTENTS

Appendix A	Flaw of Averages	2
Appendix B	Discounted Cash Flow Analysis	10
Appendix C	Economics of Phasing (to be supplied)	
Appendix D	Monte Carlo Simulation	22
Appendix E	Dynamic Forecasting	55
Appendix F	Financial Real Options Models (to be supplied)	

APPENDIX A

FLAW OF AVERAGES

The term “Flaw of Averages” refers to the widespread -- but mistaken -- assumption that evaluating a project around average conditions gives a correct result. This way of thinking is wrong except in the few, exceptional cases when all the relevant relationships are linear. Following Sam Savage’s suggestion¹, this error is called the “Flaw” of Averages to contrast with the phrase referring to a “law” of averages.

The Flaw of Averages can be the source of significant loss of potential value for the design for any project. The rationale for this fact is straightforward:

- A focus on an “average” or most probable situation inevitably implies the neglect of the extreme conditions, the real risks and opportunities associated with a project.
- Therefore, a design based on average possibilities inherently neglects to build in any insurance against the possible losses in value, and fails to enable the possibility of taking advantage of good situations.

Designs based on the Flaw of Averages are systematically vulnerable to losses that designers could have avoided, and miss out on gains they could have achieved. The Flaw of Averages is an obstacle to maximizing project value.

The Flaw of Averages is a significant source of loss of potential value in the development of engineering systems in general. This is because the standard processes base the design of major engineering projects and infrastructure on some form of base-case assumption. For example, top management in the mining and petroleum industries have routinely instructed design teams to base their projects on some fixed estimate of future prices for the product (in the early 2000s, this was about \$50 per barrel of oil). Likewise, the designers of new military systems normally must follow “requirements” that committees of generals, admirals and their staff have specified. The usual process of conceiving, planning and designing for technological systems fixes on some specific design parameters -- in short is based on the Flaw of Averages -- and thus is unable to maximize the potential value of the projects.

We definitely need to avoid the Flaw of Averages. Because it is deeply ingrained in the standard process for the planning, design and choice of major projects, this task requires special efforts. The rewards are great, however. The organizations that manage to think and act outside of the box of standard practice will have great competitive advantages over those that do not recognize and avoid the Flaw of Averages.

To appreciate the difficulties of getting rid of the Flaw of Averages, it is useful to understand both why this problem has been so ingrained in the design of technological projects, and how it arises. The rest of Appendix A deals with these issues.

Current Design Process for Major Projects Focuses on Fixed Parameters

The standard practice for the planning and delivery of major projects focuses on designing around average estimates of major parameters. For example, although oil companies know that the price of a barrel of oil fluctuates enormously (between 1990 and 2008, it ranged from about 15 to 150 dollars per barrel), they regularly design and evaluate their projects worldwide based on a steady, long-term price. Similarly, designers of automobile plants, highways and airports, hospitals and schools, space missions and other systems routinely design around single forecasts for planned developments.

Complementarily, the standard guidelines² for system design instruct practitioners to identify future “requirements”. Superficially, this makes sense: it is clearly important to know what one is trying to design. However, this directive is deeply flawed: it presumes that future requirements will be the same as those originally assumed. As the experience with GPS demonstrates, requirements can change dramatically (see Box A.1). Designers need to get away from fixed requirements. They need to identify possible future scenarios, the ranges of possible demands on and for their systems.

Box A.1 about here

The practice of designing to fixed parameters has a historical rationale. Creating a complex design for any single set of parameters is a most demanding, time-consuming activity. Before cheap high-speed computers were readily available, it was not realistic to think of repeating this task for hundreds if not thousands of possible combinations of variations of design parameters. Originally necessary, the practice of designing to fixed parameters is now ingrained in practice.

Management pressures often reinforce the pattern. In large organizations, top management and financial overseers regularly instruct all the business units to use identical parameters – for example regarding the price of oil. They do this to establish a common basis for comparing the many projects that will be proposed for corporate approval and funding. A fixed baseline of requirements makes it easier for top management to select projects. However, when the conditions imposed fixed are unrealistic – as so often is the case – these prevent designers from developing systems that could maximize value.³

The conventional paradigm of engineering further reinforces the tendency to accept fixed parameters for the design. Engineering schools commonly train engineers to focus on the purely technical aspects of design.⁴ A widespread underlying professional view is that economic and social factors are not part of engineering; that although these may have a major effect on the value of a project, they are not suitable topics for engineering curricula. The consequence in practice is the tendency for designers to accept uncritically the economic and social parameters given to them, for example the forecasts of demands for services or anticipations of the legal or regulatory rules.

In short, the practice of designing to fixed parameters is deeply entrenched in the overall process for developing technological projects. Even though the focus on requirements, on most likely futures, or on fixed estimates, leads to demonstrably incorrect estimates of value, entrenched habits are not likely to change easily. Current and future leaders of the development of technological systems need to make determined efforts to make sure that the Flaw of Averages does not stop them from extracting the best value from their systems.

The Significant Errors

The Flaw of Averages is associated with a very simple mathematical proposition. This is that:

The Average of all the possible outcomes associated with uncertain parameters
is generally not equal to the
Value obtained from using the average value of the parameters.

Formally, this can be expressed as:

$$E[f(x)] \neq f[E(x)], \text{ except when } f(x) \text{ is linear}$$

This expression is sometimes called Jensen's Law.⁵ In this formula, $f(x)$ is the function that defines the value of a system for any set of circumstances, x . It links the input parameters with the value of the system. In practice, $f(x)$ for a system is not a simple algebraic expression. It is typically some kind of computer model. It may be a business spreadsheet, a set of engineering relationships, or a set of complex, interlinked computer models. It may give results in money or in any other measure of value, for example the lives saved by new medical technology. $E[f(x)]$ indicates the Expected Value of the system, and $E(x)$ indicates the Expected Value of the parameters x .

Put in simple language, the mathematical proposition means that the answer you get from a realistic description of the effect of uncertain parameters generally differs – often greatly -- from the answer you get from using estimates of the average of uncertain parameters.

This proposition may seem counter-intuitive. A natural reasoning might be that:

- if one uses an average value of an uncertain parameter,
- then the effects of its upside value will counter-balance the effects of the downside value.

False!

The difficulty is that the effects of the upside and downside values of the parameter generally do not cancel each other out. Mathematically speaking, this is because our models of the system, $f(x)$, are non-linear. In plain English, the upside and downside effects do not cancel out because actual systems are complex and distort inputs asymmetrically.

The systems behave asymmetrically when their upside and the downside effects are not equal. This occurs in three different ways:

- the system response to changes is non-linear;

- the system response involves some discontinuity; or
- management rationally imposes a discontinuity.

The following examples illustrate these conditions.

The system response is non-linear. The cost of achieving any outcome for a system (the cars produced, the number of messages carried, etc) generally varies with its level or quantity. Typically, systems have both initial costs and then production costs. The cost of producing a unit of service therefore is the sum of its production cost and its share of the fixed costs. This means that the costs/unit are typically high for small levels of output, and lower for higher levels – at least at the beginning. High levels of production may require special efforts – such as higher wages for overtime or the use of less productive materials. Put another way, the system may show economies of scale over some range, and increased marginal costs and diseconomies of scale elsewhere. The overall picture is thus a supply curve similar to Figure A.1.

When the costs – or indeed any other major aspect of the project – do not vary linearly, the Flaw of Averages applies. Box A.2 illustrates the effect.

The system response involves some discontinuity: A discontinuity is a special form of non-linearity. It represents a sharp change in the response of a system. Discontinuities arise for many reasons, for example:

- The expansion of a project might only occur in large increments. Airports, for example, can increase capacity by adding runways. When they do so, the quality of service – in terms of expected congestion delays – should jump dramatically.
- A system may be capacity constrained and may impose a limit on performance. For instance, revenues from a parking garage with a limited number of spaces will increase as the demand for spaces increases, but will then stop increasing once all the spaces are sold. Box 1.3 illustrates this case.

Management rationally imposes a discontinuity: Discontinuities often arise from management actions, from outside the properties of the physical project. This happens whenever the system operators decide to make some major decision about the project – to enlarge it or change its function, to close it or otherwise realign its activities. Box A.3 gives an example of how this can happen.

Take-Away

Do not be a victim of the Flaw of Averages. Do not value projects or make decisions based on average forecasts. Consider, as best you practically can, the entire range of possible events and examine the entire distribution of consequences.

Box A.1

Changing Requirements: GPS

The designers of the original satellite-based Global Positioning System (GPS) worked with purely military requirements. They focused on military objectives, such as guiding intercontinental missiles. They were enormously successful in meeting these specifications.

GPS has also become a major public success. Satellite navigation now replaces radar in many forms of air traffic control. It is also a common consumer convenience embedded in cell phones, car navigation systems and many other popular devices. It provides enormous consumer value and could be very profitable.

However, the original specifications failed to recognize possible commercial requirements. The design thus did not enable a way to charge for services. Thus, the system could not take advantage of the worldwide commercial market, and could not benefit from these opportunities. The developers of the system thus lost out on major value that would have been available if they had recognized possible changes in requirements.

Box A.2

Consider a regional utility whose production comes from a mix of low-cost hydropower and expensive oil-fired thermal plants. Its average cost per kilowatt-hour (kWh) will be lower if Green policies reduce demand, and higher if economic growth drives up consumption, as in Table A.1. What is the margin of profitability for the utility when it is obliged to sell power to consumers at a fixed price of \$0.06/ kWh?

If we focus on the average forecast, that is a consumption of 1000 Megawatt-hours, then the utility has a margin of \$0.01/ kWh (= 0.06 - 0.05) for an overall profit of \$10,000. However, if we consider the actual range of possibilities, then we can see that the high costs incurred when the demand is high, boosted by the high volumes of demand, lead to losses that are not compensated on average by the possible higher profitability if demand and prices are low. In this example, the average value is thus actually \$1,600 compared to the \$10,000 estimate as Table A.2 shows. A focus on average conditions thus leads to an entirely misleading assessment of performance. This is an example of the Flaw of Averages.

Tables A.1 and A.2 about here

Box A.3

Valuation of an oil field

Consider the problem of valuing a 1M barrel oil field that we could acquire in 6 months. We know its extraction will cost \$75/bbl, and the average estimate of the price in 6 months is \$80/bbl. However, the price is equally likely to remain at \$80/bbl, drop to \$70/bbl or rise to \$90/bbl.

If we focus on the average future price, the value of the field is \$5 M = 1M (80 – 75). What is the value of the field if we recognize the price uncertainty? An instinctive reaction is that the value must be lower because the project is more risky. However, when you do the calculation scenario-by-scenario, you find that intuition to be entirely wrong!

If the price is \$10/bbl higher, the value increases by \$10M to \$15M. However, if the price is \$10/bbl lower, the value does not drop by \$10M to – \$5M as implied by the loss of \$5/bbl when production costs exceed the market price. This is because management has the *flexibility* not to pump, and to avoid the loss. Thus the value of the field is actually 0 when the price is low. The net result is that the actual value of the field would be \$6.67M, higher than the \$5M estimate based on the average oil price of \$80/bbl.

The field is worth more on average, not less. This manifestation of the Flaw of Averages illustrates why it worthwhile to consider flexibility in design. If management is not contractually committed to pumping, it can avoid the downside if the low oil price occurs while still maintaining the upsides. This flexibility increases the average value of the project compared to an inflexible alternative.

Table A.1. Cost of supplying levels of electricity

Level of Use Megawatt-hours	Probability	Average Cost \$ / kilowatt-hour
1200	0.3	0.08
1000	0.4	0.05
800	0.3	0.04

Table A.2. Actual profitability under uncertainty

Level of Use Megawatt-hrs	Probability	Cost \$ / kwh	Margin \$ / kwh	Overall Profit \$	Expected Profit \$
1200	0.3	0.08	- 0.02	- 24,000	- 7,200
1000	0.4	0.05	+ 0.01	10,000	4,000
800	0.3	0.04	+ 0.02	16,000	4,800
Total					1,600

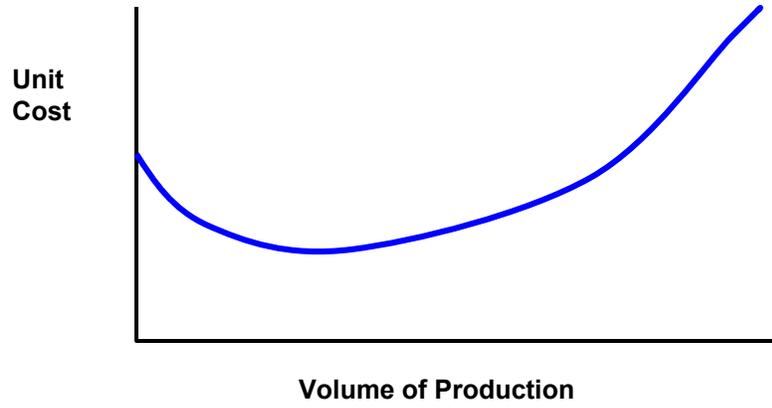


Figure A.1. Typical Supply Curve for the Output of a System

APPENDIX B

DISCOUNTED CASH FLOW ANALYSIS

This text refers throughout to discounted cash flow (DCF) analysis, the most common methodology for an economic appraisal and comparison of alternative system designs. Although widely used, DCF has significant limitations in dealing with uncertainty and flexibility. In response, some academics call for its wholesale replacement by a different and “better” method.⁶ This is not our approach. We wish to build upon the widespread use of the discounted cash flow methodology. We thus advocate a pragmatic incremental improvement of the method to alleviate its limitations in dealing with uncertainty and flexibility.

To improve DCF sensibly, it is important to understand its elements and procedures. This appreciation supports the use of Monte Carlo simulation to deal with uncertainty, as Appendix D indicates.

The purpose of this Appendix is to remind readers of the basic principles of DCF:

- Its main assumptions and associated limitations;
- The mechanics of discounted cash flows;
- The calculation of a net present value and an internal rate of return; and importantly,
- The rationale for the choice of a suitable discount rate.

The issue

Every system requires cash inflow (investments and expenses) and generates cash outflow (revenues) over time. These cash flows are the basis for the economic valuations of system designs. Indeed, economists tend to regard projects or system designs merely as a series of cash flows. They are oblivious to the engineering behind their generation. The question is: how should we value these revenues and expenses over time?

The underlying economic principle is that money (or more generally, assets) has value over time. If we have it now we can invest it productively and obtain more in the future. Conversely, money obtained in the future has less value than the same amount today.

Algebraically:

X money now $\rightarrow (1+d)^t X = Y$ at a future date

Y at a future date $\rightarrow Y/(1+d)^t = X$ money now,

where $d > 0$ is the rate of return per dollar we could achieve if we invested money over the respective period. The rate of return captures the time value of money. This means that cash flows Y in the future should have less value, that is, be “discounted”, when compared to investments X now. This is the rationale behind **discounted** cash flow analysis

The most fundamental assumption behind a DCF analysis is that it is possible to project the stream of net cash flow, inflow minus outflow, with a degree of confidence over the lifetime of

a project or system. To facilitate the analysis, it is usual to aggregate cash flows temporally, typically on an annual basis. Table B.1 shows illustrative cash flows of two system designs.

Table B.1 about here

Which design would you prefer? Design A requires a lower initial investment but annual cash investments of \$100M for three more years before it is completed and sold off for \$625M. Design B requires a substantially larger initial investment but delivers positive cash flows from year 1 onwards. However, Design B has to be decommissioned in year 5, at a cost of \$100M. Note that if you do not discount income and expenses, Design A is a winner (it nets 155M = 625 - 470) and Design B is a loser (it shows a net loss of 20M = 550 - 570). This is the perspective of tax authorities, but it does not constitute a proper economic analysis because it does not account for the time value of money.

A proper economic valuation has to answer two questions:

- Is a particular design worth implementing? Does it have positive economic value?
- If there are several design alternatives, which design is preferable?

The answers to these questions depend on the time value of money, specifically on the availability and cost of the capital needed to finance a project. This is the return the organization will have to pay to its financiers: to its banks in the form of interests on loans, and to its owners in the form of dividends. Discounted cash flow (DCF) analysis offers a way to answer these questions.

DCF principle

DCF is based on a rather simplistic view of an organization's finances: all its capital sits in a single virtual bank account. The group finances all its investments from this account. It receives any surplus that the project generates and pays out any shortfall or investment needed. In reality, firms finance their investments from a variety of sources with different conditions attached. They also place capital in different investments with different risks and returns. However, the assumption of a single bank account simplifies comparisons between projects considerably and, importantly, makes these comparisons conceptually independent of the specific financing arrangements. For government agencies, the sources of income and expenses are much more complicated, but the DCF principle applies equally to them and to government projects. In practice, the main difference between government and corporate uses of DCF analysis lies in the choice of discount rate.

The virtual bank account has an interest rate attached, like any bank account. A second simplifying assumption of a traditional DCF analysis is that this interest rate is fixed and the same for both deposits (positive cash flowing back from projects) and loans (cash injections required to run the projects). This is obviously different from real banks. The interest rate of the company's virtual bank account is called its **discount rate**. We discuss this discount rate in more detail later.

Once we accept the financial view of the company or government agency as a virtual bank account, it is natural to value the cash flow stream that a project generates over time as its contribution to the bank account. The analysis considers that the bank account finances any shortfalls, that is, negative cash flows; and receives and earns interest on any positive cash flows; all at a fixed interest rate. Any project is worthwhile if it generates a net positive return. Design A is preferable to Design B if A's net return is larger than B's.

Box B.1 about here

DCF mechanics

One way to quantify the value of a project is to calculate its net contribution to the company bank account at the end of its life. Table B.2 shows this calculation in detail for the two designs in Table B.1. (We show the details to explain the process. In practice, analysts use standard spreadsheet functions to get the results without the detail.) The end-of-life contribution of Design A is \$12,003 after 4 years; the corresponding contribution of Design B is \$6,287 after 5 years. Note, from the perspective that recognizes interest, that Design B appears worthwhile -- as it does not if we fail to take into account the time value of money.

Table B.2 about here

If two projects have the same duration, then their end-of-life contribution is a sensible way of comparing their economic value. However, we generally have to compare projects with differing durations, as in the case of Designs A and B in Table B.2. In this situation, it is not fair to compare the end of life results (for example \$12,003 for Design A versus \$6,287 for Design B) because the same amounts in different years are not equivalent. To make the comparisons fair, to compare apples with apples as it were, it is customary to relate the end-of-life contributions of projects to a sum of money at a common time. This common time is normally the present. We thus refer to the **present value** of projects. More specifically, to indicate that we account for the difference between the revenues and expenses, we focus on the **net present value** (NPV).

The net present value of a project is independent of its duration. It is natural to think of today's equivalent of the positive end-of-life contribution of a project as the sum of money we can borrow from the bank account today against the net contribution of the project, its NPV. In other words, the NPV is the amount we can withdraw from the company bank account today (and spend in other ways), so that the end-of-life contribution of the project will allow us to pay back the accrued debt of our withdrawal at the end of the project.

An amount X borrowed from the company account at annual interest rate r will accrue to $X(1+r)^T$ at the end of T years. This sum has to be equated to the project's end-of-life contribution. So, $NPV = (\text{end-of-life contribution})/(1+r)^T$. Using a discount rate of 10% ($=0.1$), the NPV of Design A is therefore $\$12,003/(1.1)^4 = \$8,198$. Likewise, the NPV of Design B is $\$6,287/(1.1)^5 = \$3,094$.

If the end-of-life contribution of a project is negative, then its NPV is the amount we need to deposit in the company bank account today to cover the project's shortfall at the end of its lifetime. Financial theory and common sense suggest that companies should not invest in projects with a negative NPV, unless there are significant externalities that offer value beyond the project cash flow, such as access to the client's future more profitable projects.

Short-cut calculation: It is common practice to calculate the NPV using a short cut based on discounted annual cash flows. We do this by calculating for each annual cash flow the sum that could be borrowed from or needs to be deposited in the account today to balance out the cash flow in the respective year. These sums are of course smaller than the actual cash flow, which is why they are called "discounted cash flows". The NPV for the entire project is then the sum of these discounted annual cash flows.

For example, Design A's cash flow in year 3 is -\$100,000. If X is the amount deposited today to balance out this shortfall, then $X \cdot 1.1^3$ is the amount available after 3 years. Therefore the amount X today that balances out the negative cash flow of -\$100,000 in year 3 is $-\$100,000/1.1^3 = -75,131$. Table B.3 illustrates the short form of the NPV mechanics for the Designs in Table B.1.

Table B.3 about here

Choosing the discount rate

The discount rate is at the heart of the DCF principle of a virtual bank account. Which discount rate should one choose? As a practical matter, most system designers have to use the discount rate set by higher authority. In most companies, the board of directors or the chief finance officer sets a discount rate employees should apply to all projects across the company. The process is similar for government agencies: somebody is responsible for establishing the applicable rate.⁷ Designers may thus not have a choice about the discount rate. However, it is still useful to understand the rationale for its selection.

The discount rate should capture the firm's cost of capital: The firm should be able to raise funds at the discount rate and should be able to invest these funds so that the return on these investments equals the discount rate. Suppose a firm can borrow from a bank at 6% interest. Does that mean that its discount rate should be 6%? The answer is no. In fact, a significant proportion of the firm's capital will come from the owners of the firm, rather than from banks. The owners are more vulnerable to the risk of default because debtors have preferential access to liquidation proceeds in the case of bankruptcy. Indeed, the bank will only grant the loan because there are owners who are primarily liable with their invested capital (their "equity") in the case of bankruptcy. The owners therefore take more risk and will demand a higher return from the company than the bank's 6%. The cost of the owners' equity, that is, their return expectations, is higher than the cost of debt, which is the interest on loans. To determine a sensible discount

rate, we need to derive the average cost of capital by appropriately weighting the cost of equity and debt.

Weighted average cost of capital (WACC): This quantity provides a reasonable estimate of the discount rate. It represents the average cost return expected by the owners and banks that finance a project. A simple example illustrates the process for its calculation. If the company's total market value (number of shares times share price) amounts to 75% of the company's total invested capital and the shareholders expect an annual dividend of 10%, and the remaining 25% of the total invested capital is financed through debt at an average interest rate of 6%, then the company's WACC is $75\% \cdot 10\% + 25\% \cdot 6\% = 9\%$. In practice, the details of the calculation are more complex and depend on the specific context of the company⁸. Conceptually, the important point to retain is that the proper discount rate is higher than the interest rate paid for loans to finance a project.⁹

Making decisions with NPV

Recall the fundamental questions: is a project economical? If projects compete, which one should we implement? The so-called NPV rule stated in Box B.2 gives the answers.

Box B.2 here

A project may be worthwhile even if its NPV is not hugely positive. All the entire costs of financing the project, that is, the dividend payments and the interests on various loans, are already factored into our set-up of the virtual bank account -- via the discount rate. Any positive NPV indicates that the project delivers more than the normal, threshold rate of return. If a company continuously produces highly NPV positive projects, the market will realize this and value the company higher, which will lead to higher dividend expectations. This in turn will lead to an adjustment of the discount rate, which will reduce the NPVs of typical projects. A project is desirable if its NPV is positive, even slightly.

Box B.3 about here

Dependence of NPV on discount rate

The NPV depends on the discount rate in a rather complex way. Figure B.3 illustrates this by exhibiting the NPVs of Designs A and B as the discount rate changes. Design A is much more sensitive to discount rate changes, even though it is of shorter duration. Its NPV decreases as the discount rate increases. This makes intuitive sense. Insofar as the discount rate is the company's cost of capital, the more the company has to pay the banks and owners, the less the company keeps. Design B is more stable but its NPV has an interesting property: its dependence on the discount rate is non-monotonic. For very low discount rates, its NPV is negative. Rising discount rates then lift the NPV above zero (by decreasing the importance of the large closure cost at the

end of the project). As discount rates rise further, they adversely affect NPV (by diminishing the value of the positive returns) and it turns negative again for very large discount rates.

Figure B.3 about here

Why does NPV depend so much on the discount rate and project life? To answer this question, consider the effects on positive and negative cash flows separately. The contribution of a positive cash flow in year T to the NPV equals the sum of money we can withdraw from the company account today, so that the accrued debt in the account in year T will be covered by the positive project cash flow in year T . If the discount rate on the virtual account increases, the hole in the account grows faster. The amount we have available in year T , however, will remain the same size. Therefore, we have to make the hole smaller, that is, withdraw less money. This deteriorating effect of higher discount rates on the value of positive cash flows is larger the further in the future the cash flow lies. This results from the compounding of interest that we pay on our withdrawal. Figure B.4 shows the effect: Positive cash flow further in the future deteriorates faster as the discount rate increases. In summary, positive cash flows contribute less to the NPV as discount rates increase and this deterioration is greater for late cash flows.

Figures B.4 and B.5 about here

The effect on negative cash flow is the opposite. As the interest rate on the account increases, the amount we need to deposit today to fill the hole due to the future negative cash flow becomes smaller because the deposit will grow faster. A negative cash flow of \$100 in 2 years at a 10% discount rate is equivalent to borrowing $\$100/1.1^2 = \82.64 , from the bank account today. The same cash flow in 4 years is equivalent to withdrawing $\$100/1.1^4 = \68.30 today. Therefore, as the discount rate increases, it dampens the negative contribution of negative cash flows to the NPV and this positive effect is greater for late cash flows (see Figure B.5).

Box B.4 about here

In the example in Table B.1, Design A has a large late positive cash flow, whose contribution is very vulnerable to changing discount rates. This causes the sensitivity of this design's NPV, see Fig. B.3. On the other hand, Design B has most of its positive cash flow early on. The negative effect of increased discount rates on these cash flows will be relatively mild because they occur early. At the same time, the project has a large negative cash flow at the end. The increasing discount rate dampens the negative effect of this late payment, so has an overall positive effect on the NPV. These moderate negative effects on the early positive cash flows and the positive effect on the late negative cash flow balance out and lead to the relative robustness of design B to changes in the discount rate.

Box B.1

Assumptions of discounted cash flow analysis

1. We can confidently estimate annual net cash flows of projects over the project lifetime.
 2. Cash flow shortfalls in any one year come from a “virtual bank account”, net positive cash flow will be deposited in this account.
 3. The project sponsor, a company or government agency, can borrow arbitrary amounts from the virtual account.
 4. The interest rate for borrowing is fixed and the same as for deposits. This interest rate is called the discount rate.
 5. Projects and design alternatives should be compared by their net contribution to the virtual bank account over their lifetime.
-

Box B.2

NPV rule

Assumption: The cash flow profiles of projects fully capture their economic value.

- A) A project is economical if its NPV is non-negative.
 - B) If two projects are mutually exclusive, then the project with the higher NPV is preferable.
-

Box B.3**Internal rate of return (IRR)**

A frequently used alternative to the NPV method is the calculation of the internal rate of return (IRR). As with NPV, one starts by calculating the net contribution at the end of the project. This net contribution depends on the interest rate charged on the virtual bank account, the discount rate. How high an interest rate can be charged before the end-of-life contribution of a project turns negative? Alternatively, equivalently, what interest rate would lead to a zero NPV? This interest rate is called the internal rate of return (IRR).

If the discount rate equals the IRR, then the net contribution of the project to the bank account is zero; cash inflows and outflows balance out, accounting associated interest payments.

Accordingly, we should only invest in projects with an internal rate of return at or above their discount rate. If we thus use IRR as a method of appraisal, we often refer to the discount rate as the hurdle rate. The IRR has to exceed the hurdle rate for a design to be worth implementing.

One way of calculating the IRR is to plot the NPV for various discount rates and find the value where the NPV is zero. The result for Designs A and B from Table B.1 appears in Figures B.1 and B.2.

Figures B.1 and B.2 about here

The IRR of design A is 11%, slightly higher than the hurdle rate of 10%. Project B is an interesting case. It has two IRRs, one at 7%, the other at about 23%. In fact, the NPV will be positive for any interest rate between 7% and 23%. Since the hurdle rate of 10% is in that range, the project is economically viable. Multiple IRRs occur only if there is more than one sign change in the cash flow. Put another way, if all negative cash flows or investments happen before the first positive cash flow, then the IRR is unique.

Box B.4

The higher the discount rate, the lower the positive contribution of late positive cash flows, and the lower the negative contribution of late negative cash flows to the overall NPV.

Need to make more obvious that data are all in millions...

Table B 1. Cash flow profiles of 2 designs

CASH FLOW PROJECTIONS						
YEAR	Upfront Investment	1	2	3	4	5
DESIGN A (\$M)	-\$170.000	-\$100.000	-\$100.000	-\$100.000	\$625.000	
DESIGN B (\$M)	-\$260.000	\$200.000	\$200.000	\$100.000	\$50.000	-\$310.000

Table B 2. Lifetime contribution of two designs

LIFETIME CONTRIBUTION						
		Discount rate 10%				
YEAR	Upfront Investment	1	2	3	4	5
DESIGN A CASH FLOW (\$M)	-\$170.000	-\$100.000	-\$100.000	-\$100.000	\$625.000	
Starting balance		-\$170.000	-\$287.000	-\$415.700	-\$557.270	
Interest		-\$17.000	-\$28.700	-\$41.570	-\$55.727	
Inflow / outflow of capital at year-end		-\$100.000	-\$100.000	-\$100.000	\$625.000	
Year-end balance		-\$287.000	-\$415.700	-\$557.270	\$12.003	
DESIGN B CASH FLOW (\$M)	-\$260.000	\$200.000	\$200.000	\$100.000	\$50.000	-\$310.000
Starting balance		-\$260.000	-\$86.000	\$105.400	\$215.940	\$287.534
Interest		-\$26.000	-\$8.600	\$10.540	\$21.594	\$28.753
Inflow / outflow of capital at year-end		\$200.000	\$200.000	\$100.000	\$50.000	-\$310.000
Year-end balance		-\$86.000	\$105.400	\$215.940	\$287.534	\$6.287

Table B 3. Net present values (NPV) of two designs

NET PRESENT VALUE - SHORT VERSION						
		Discount rate 10%				
YEAR	Upfront Investment	1	2	3	4	5
DESIGN A (\$M)	-\$170.000	-\$100.000	-\$100.000	-\$100.000	\$625.000	
Effect today (discounted cash flow)	-\$170.000	-\$90.909	-\$82.645	-\$75.131	\$426.883	
Sum of discounted cash flows (NPV)	\$8.198					
DESIGN B (\$M)	-\$260.000	\$200.000	\$200.000	\$100.000	\$50.000	-\$310.000
Effect today (discounted cash flow)	-\$260.000	\$181.818	\$165.289	\$75.131	\$34.151	-\$192.486
Sum of discounted cash flows (NPV)	\$3.904					

Note: Need to format Tables consistently and, in particular, to provide complete borders.

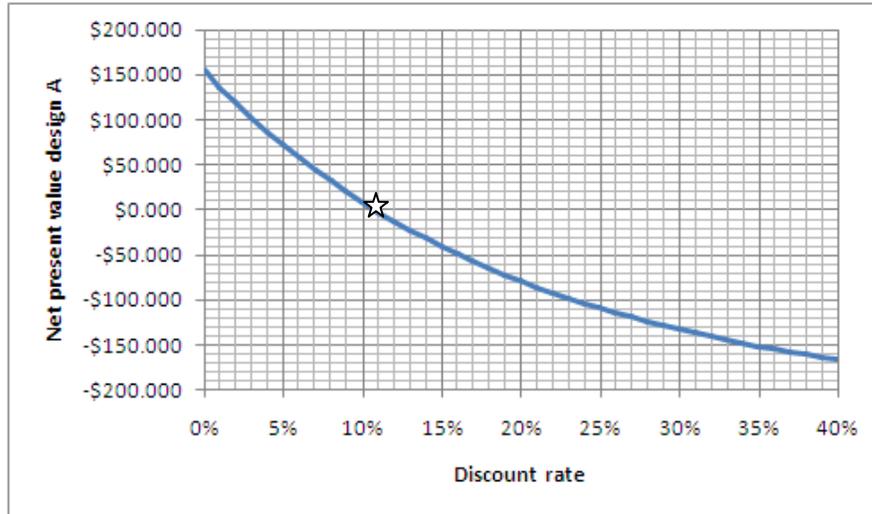


Figure B.1. Net present value for Design A as a function of discount rate. IRR is 11%.

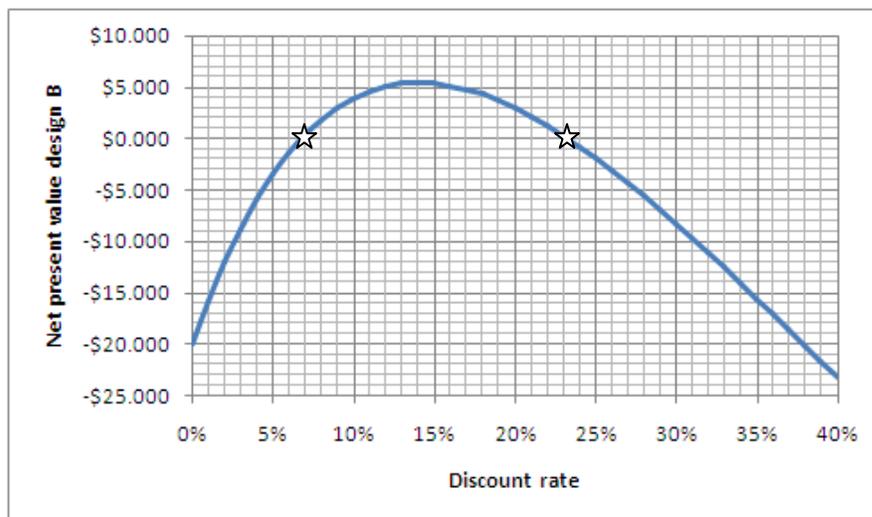


Figure B.2. Net present value for Design B as a function of discount rate. IRR is either 7 or 23%.

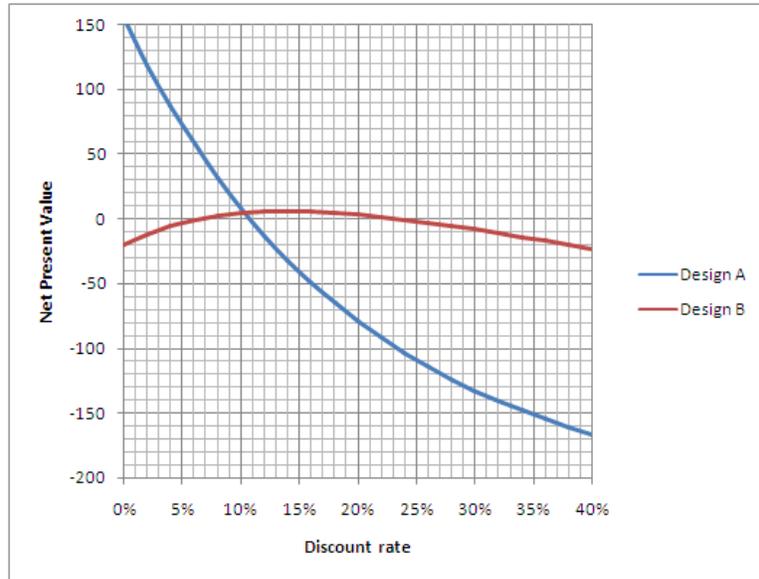


Figure B.3: NPV of two designs for varying discount rates



Figure B.4. Present value of \$1 occurring in 2 or 4 years, for a range of discount rates.



Figure B.5. Dependence of negative cash flows on project life.

APPENDIX D

MONTE CARLO SIMULATION

Converting a standard valuation into a Monte Carlo simulation with flexibility exercise

Chapter 3 introduces Monte Carlo simulation as a preferred methodology for the study of the value of flexibility. Part 2 refers to this methodology throughout. Monte Carlo simulation models efficiently generate thousands of futures, run all these futures simultaneously through a model of system performance, and summarize the distribution of possible performance consequences graphically.

The purpose of this Appendix is to provide a roadmap to the conversion of traditional static system performance models into Monte Carlo models. The aim is to clarify these models, which can be daunting for designers, engineers or clients who are not familiar with the technique, and who rely on simpler static models such as standard NPV spreadsheets. Along the way, we introduce a variety of techniques to model uncertainty in spreadsheets.

Specifically, this Appendix takes you through the conceptual steps that lead from a static valuation model to a Monte Carlo model that can be used to articulate the value of flexibility:

- Step 1: Produce a standard valuation model
- Step 2: Perform a standard sensitivity analysis: Change one variable at a time
- Step 3: Perform a probabilistic sensitivity analysis: Change all variables simultaneously
- Step 4: Introduce distributional shapes for uncertain numbers
- Step 5: Introduce dependence between uncertain numbers
- Step 6: Introduce dynamically changing uncertain numbers
- Step 7: Model flexibility via rules for exercising flexibility.

To illustrate these steps we use the parking garage illustration of Chapter 3 as a working example.

A note on software

As in the rest of the book, the focus is on spreadsheet models because of their ubiquity. The Appendix assumes the reader is familiar with standard formulas and commands in Microsoft Excel[®]. Monte Carlo simulations can be performed in that program. Without additional software however, Monte Carlo modeling in Excel can become cumbersome and simulations of larger models tend to take a long time. Commercial software packages facilitate the modeling effort and speed up execution. Examples of commercial packages include @Risk[®], Crystal Ball[®], XLSim[®], and RiskSolver[®].

All Monte Carlo software packages have three main components:

- A collection of built-in random number generators, which are special spreadsheet formulas that allow sampling from pre-specified distributions. For example, the formula “=gen_normal(0,1)” in XLSim draws a sample from a normal distribution with mean 0 and

standard deviation 1 and puts it in the cell that contains the formula. Every time the spreadsheet is recalculated (e.g. when a cell has been modified and the enter key is hit; or when the recalculate key, the F9 function key in Excel, is hit), another sample is drawn from this distribution and appears in the cell.

- A convenient and fast way of executing thousands of scenarios and storing the results.
- An interface that facilitates the calculation of summary statistics and the generation of charts to visualize and communicate the results of simulations.

Unfortunately, different software packages use different formulas for their random number generators and their spreadsheet models are not portable.

Since we do not wish to limit our readers to a specific package, this Appendix uses standard Excel commands to generate distributions for uncertain numbers. These commands are slower and more cumbersome than the specialized formulas of commercial packages, but the resulting spreadsheets work with all packages. Readers who use Monte Carlo simulations frequently should invest in a commercial package. For convenience, we used the XLSim[®] software package to produce some of the graphs in the Appendix. This package is relatively inexpensive and produces standard Excel graphics.¹⁰

Step 1: Produce a standard valuation spreadsheet

Valuation models for system designs are input-output models. They capture how a system design:

- Converts system inputs, such as capital, labor, material, energy, demand,
- Within constraints, e.g. of a physical, regulatory, or legal nature,
- Into outputs.

Some system outputs will be desirable, such as profit, demand satisfaction, better health; others will be undesirable, such as congestion or pollution. Roughly speaking, the inputs and constraints describe the world the system faces, the outputs describe the difference that the system makes. The system itself is described by (i) a set of design parameters, such as capacity, productivity, reliability, etc, and (ii) a set of formulas that relate inputs and system parameters to outputs. Figure D.1 illustrates a valuation spreadsheet for the parking garage example.

Figure D.1 and Box D.1 about here

Modeling tips: A few tips on spreadsheet modeling are in order at this point. The advantage of a spreadsheet is that we can hard-wire calculation steps into it via formulas. This is very useful when valuation assumptions change. Suppose you have calculated an NPV for a complex project using a 10% discount rate when your boss tells you that the plan needs to be recalculated using a 12% rate. If you were using an electronic calculator to compute the NPV, you would have to perform the same calculations all over again using this 12% rate. However, if you use a well-

programmed spreadsheet, you simply change the number in the cell that contains the discount rate from 10% to 12% and the spreadsheet performs the update instantly.

To exploit this advantage of the spreadsheet, you have to dedicate a single cell to contain the discount rate, and refer to it whenever a calculation uses the discount rate. If you place the discount rate in cell A1, you should always refer to it and never use the numerical value of the discount rate. To discount the value in cell A2 by 10%, you should write “=A2/(1+A1)”. If you use “=A2/1.1” instead, your spreadsheet will not update correctly when you increase A1 from 10% to 12%. If you have calculated your NPV in a spreadsheet with the numerical discount rate mixed into formulas, as in “A2/1.1”, you will have to find all the cells that contain this rate and change them manually – a process that may well take more time than repeating the calculation with an electronic calculator. You have given away the main advantage of the spreadsheet – an instantaneous update when inputs change.

The most important rule in spreadsheet modeling is therefore that every numerical input gets one and only one dedicated cell. Then whenever you need to use that input in a calculation, you reference that respective cell. Never mix formulas and numbers. It is very useful to dedicate a range of the spreadsheet, or indeed a separate worksheet, as the input range that contains all the numerical values. The remaining space is the model range; all its cells contain formulas or cell references only. If the layout of the spreadsheet requires inputs at various places, then an alternative to a dedicated input range is suitable color-coding of all cells that contain a numerical value.

A second important point in spreadsheet modeling is that it is easy to make mistakes, to commit “slip of the keyboard” errors. Therefore models need to be very carefully validated. Read every formula several times. If you use a complex formula, input the same formula again into an adjacent cell to verify that it produces the same result. Once you include a formula, check that it behaves as expected (e.g. see that it does not produce negative values when it should not). Do this by changing the inputs the formula uses to values that may be unrealistic but for which you know what the formula result should be, e.g. the revenue for zero demand should be zero. Validate as you build the model and then again when the model is finished. Change the inputs to see if the outputs behave in the expected way. Finally, if you have the resources, the best way to validate a model is to ask a second modeler to produce an independent model of the system. While this will not avoid systematic mistakes due to a shared misunderstanding of the system behavior, it will greatly reduce the chance of a slip of the keyboard error.

A third important point is to document the model well. Use text boxes and cell comment boxes. Use more of these boxes rather than less. Regard the model as a written piece and make it easy for a “reader” to follow the line of reasoning in the model.

Step 2: Perform a standard sensitivity analysis: change one variable at a time

Every model is based on assumptions that can be called into question. The parking garage model invites questions such as “what if demand is lower than expected?”, “what if the average annual revenue per space used is lower than projected because there is more use of discounted long-time parking?”, “what if operating costs exceed projections?”, “what if the maximal capacity utilization due to variation in demand is lower than anticipated?”

A first step towards the recognition of uncertainty is to understand the effect of deviations of individual inputs from their base-line assumptions. This process is called sensitivity analysis and should be a routine part of practical valuation procedures. To carry out this analysis, you keep all inputs except one at their base values and alter the free input to track corresponding changes of performance measures. The data table command performs this analysis efficiently. In fact, it is arguably the most powerful command in Excel. Everyone who works with valuation spreadsheets should familiarize themselves with data tables.

Figure D.2 about here

Figure D.2 shows a specific sensitivity chart for the parking garage case. The original model assumed initial demand of 750 spaces, additional demand by year 10 of 750 spaces, and then a further additional demand beyond year 10 of 250 spaces. The data table command is used to analyze the effect of deviations from these assumptions by +/- 50%.

Note that the graph shows an important asymmetry. Low demand has a more pronounced effect on the NPV than high demand. This is a consequence of the fact that we cannot capture high demand when the garage is at full capacity. As Chapter 3 and Appendix A explain, such asymmetries cause the Flaw of Averages. That is, the NPV based on base case conditions, is not the average of the NPVs obtained as the condition varies around the base value. Therefore, whenever you generate sensitivity graphs that are not straight lines, this is a sign that there is Flaw of Averages in the system. You should also be suspicious when you generate straight-line sensitivity graphs and check that you have not missed a constraint in the system.

Tornado diagram: The tornado diagram shown in Figure D.3 is a useful tool for sensitivity analysis. This graph summarizes the relative effects of variations of several inputs over their ranges, under the assumption that the other variables remain at their base values. Several freeware spreadsheet add-ins are available to facilitate the production of tornado diagrams.¹¹

Figure D.3 about here

The tornado diagram illustrates which uncertain inputs most affect the relevant performance metric. Each variable has an associated bar representing its impact on the performance metric as it varies over its prescribed range. The diagram sorts the bars from top to bottom by their length, going from longer ones at the top to shorter ones at the bottom, so that the result looks like a funnel -- hence the name tornado diagram. The graph is particularly useful if there are a large number of uncertain inputs in the valuation model and it is impractical to analyze

them all carefully. The tornado diagram provides a way to prioritize and choose the most important uncertainties to consider: the variables associated with the longest bars.

Tornado diagrams have a second important advantage. Just as sensitivity graphs, they allow the detection of asymmetries and therefore potential Flaws of Averages. In our example, the tornado diagram in Figure D.3 shows that when demand varies symmetrically by +/- 50% its effect is skewed: the bar to the right is shorter than the bar to the left of the base value. Equal changes in this input lead to unequal changes in performance, the indication that the Flaw of Averages is at work.

Box D.2 about here

Step 3: Perform a probabilistic sensitivity analysis: change all variables simultaneously

A considerable drawback of traditional sensitivity analysis is that it inspects the effect of changes only one input at a time, holding the other inputs constant.¹² In mathematical terms, standard sensitivity analysis is akin to estimating the partial derivatives of performance measures with respect to the input variables. This variable-by-variable information only provides a good approximation of performance sensitivity close to the base values of the inputs.

Simultaneous effects when two or more variables differ jointly from their projected values can be very important, in particular when their ranges are large. For example, consider net revenues as the product of margin and sales volume. If the projected margin is very small, then additional sales do not greatly increase net revenues. In the extreme case of a zero margin, additional sales have no impact on net revenues. So traditional sensitivity analysis may conclude that sales volume uncertainty has little effect on performance. However, if margin and sales volumes grow simultaneously, then this can have a significant effect – one that standard what-if-analysis would overlook.

Probabilistic sensitivity analysis alleviates the problem of one-dimensional sensitivity analysis. This technique sits between standard sensitivity analysis and a full Monte Carlo simulation. Just as standard sensitivity analysis, probabilistic sensitivity analysis works only with ranges and assumes that all input values within their range are equally likely. However, in contrast to standard sensitivity analysis, it simultaneously randomly samples inputs from their respective ranges. A particular trial of a probabilistic sensitivity analysis therefore consists of a choice for each input from its range. The process records the inputs and their associated outputs, and repeats the sampling many times. The advantage of probabilistic sensitivity analysis is that it explores the effect of joint changes in the inputs. It is thus more realistic than the traditional one-at-a-time sensitivity analysis.

Box D.3 about here

Creating the probabilistic analysis: To build a model for probabilistic sensitivity analysis, it is necessary to develop a way to sample the ranges of variation. To do this we can use Excel's

random number generator, the RAND() function. A cell that contains “=RAND()” will contain a number between 0 and 1. Moreover, this number updates when the spreadsheet re-calculates. A Microsoft Excel spreadsheet re-calculates every time a cell is changed. It also re-calculates when the function key F9 is hit. Hitting F9 repeatedly is like rolling dice for the cell with the “=RAND()” formula. If you recalculate the spreadsheet many times, say several thousand, you will find that the numbers in the cell spread uniformly over the interval 0 to 1.

The RAND() function can be used to sample an uncertain input from its range. Suppose the lower bound of an input is in cell A1 and its upper bound in cell B1. If cell C1 contains the command “=A1+RAND()*(B1-A1)”, It will contain a number between A1 and B1. Every time the function key F9 is hit, a different number will occur. The RAND() function ensures that the numbers between A1 and B1 have, for all practical purposes, the same chance of being sampled into cell C1.

When model inputs change, the performance of the system changes as well. Monte Carlo simulation refers to a computer program, or spreadsheet, that allows automatic sampling and recording of input values and associated output values. Figure D.4 shows the first 20 trials of a Monte Carlo simulation for the parking garage example. Each row refers to one hit of the F9 key, i.e., one combination of inputs, sampled randomly from their prescribed ranges. The results are recorded in columns with sampled inputs first, followed by generated outputs (NPV in this case). With Monte Carlo software add-ins it takes literally seconds to generate these data. An alternative, using standard Excel, is to employ the data table command.¹³

Figure D.4 about here

Correlation analysis: Once the Monte Carlo data is generated, it can be analyzed statistically to provide additional understanding of the sensitivity of the relationships between inputs and outputs. Correlation analysis is a way of studying the dependence of outputs on the various inputs. As the tornado diagram, it helps prioritize uncertain inputs, identifying those that contribute most to the uncertainty of the NPV. Figure D.5 shows the correlation between the NPV and the uncertain inputs for the garage case, and confirms that demand uncertainty is the input that affects its performance most significantly.

Figure D.5 about here

Correlation analysis has the advantage of being more realistic than a tornado analysis. It allows all uncertainties to vary simultaneously over their ranges, and not one-by-one with the others kept fixed at their base values, as in the alternative tornado diagram.

Correlation has the disadvantage of being an unintuitive concept. It measures the degree to which two uncertain numbers are linearly related. However, uncertain numbers with a low correlation can still be closely related nonlinearly. This is an issue because nonlinearities between system inputs and outputs occur quite often, for example because of system constraints. Take for

example an inventory of 2000 parts. Suppose demand ranges between 1000 and 3000 parts with all values equally likely. If demand is lower than 2000, then there is waste which costs \$1,000 per part; if demand is higher than 2000, then there are lost sales which are again costly, say also \$1,000 per part. Therefore cost = $|\text{demand} - 2000 \text{ units}| * \$1000/\text{unit}$. Because of the absolute value function $|\cdot|$ the output (cost) is a nonlinear function of the input (demand). There is a clear relationship between demand and cost, as Figure D.6 shows, yet the correlation coefficient of the generated pairs of demands and costs is about 0.03, statistically indistinguishable from zero.¹⁴

Figure D.6 about here

Scatter plot: A scatter plot matrix is a better tool for the analysis of relationships between input and output variables. It simply replaces each correlation coefficient in a correlation matrix by a scatter plot of the two relevant variables.¹⁵ Scatter plots are more informative and more intuitive than correlation coefficients.

Figure D.7 shows two scatter plots generated from Monte Carlo output for the parking garage, with *all* inputs sampled simultaneously from the ranges given in Figure D.3. Notice that the range of the performance metric on the vertical axis, the NPV, is the same for both plots, while the horizontal axis covers the interval specified for the input variable. The first observation is that neither of these inputs has a dominating effect on the overall uncertainty. The residual NPV uncertainty, driven by the remaining input uncertainties and depicted by the vertical spread of the points cloud, is substantial for both scatter plots. Nevertheless, we recognize that demand uncertainty has a more pronounced effect. The effect of the uncertainty in capacity utilization is in almost entirely wiped out by the uncertainty in the remaining inputs.

Figure D.7 about here

Figure D.7 shows that the relationship between demand and NPV is non-linear, roughly following the curve in Figure D.2. As indicated earlier, this is because capacity constraints cut off the benefits of high demands. The scatter plot also shows that the spread of points to the right is larger than to the left, that is, the expected NPV and its residual uncertainty, driven by the remaining uncertain inputs, increase when demands are larger than expected.

Optimization in the context of uncertainty: Optimization, finding the “best” set of system parameters, is an important step in the design process. In our illustrative garage case the number of levels is the only design parameter. Going back to our fixed projection model in Figure D.1, we can find the optimum by varying the number of levels, as in Figure D.8. On this basis it seems sensible to build 6 levels, albeit the difference to the 5-level garage is not large.

Figure D.8 about here

An alternative optimization technique is to simply sample the number of levels, together with the other uncertain inputs from their intervals, say between 0% and 5%. The

“RANDbetween” function in Microsoft Excel allows you to sample integers between any two given integers, giving each the same probability. “=RANDbetween(1,8)” for example gives every integer between 1 and 8 the same probability of being selected. Figure D.9 shows the result for the parking garage. It makes clear that the choice between levels 4, 5, 6 or 7 could be dominated by the uncertainties in outcome, especially since the possible losses are 10 times the expected NPVs.

Figure D.9 about here

In general, designs “optimized” for deterministic cases are often not best when we recognize uncertainties and understand their effects. A second example reinforces this message. This case is taken from a service industry and concerns the staffing level, which can vary continuously. See Figure D.10. Whereas the static base case NPV model seemingly provides clear optimization guidance about the optimum level, the more realistic probabilistic model shows that there is no point in thinking too much about precise optimization, relative to the uncertainty in the environment. The base case NPV model gives the false impression that there is a clear-cut optimal solution – and this is not right.

Figure D.10 about here

Step 4: Introduce distributional shapes for uncertain numbers

Step 3 extends traditional variable-by-variable sensitivity analysis to a simultaneous sensitivity analysis that samples thousands of input combinations from their respective ranges. That analysis assumes that all input realizations within a given range are equally likely. However, we may have good reasons to believe that this is not the case. For example, realizations in the centre of the range may be regarded as more likely than those at the extremes. The accommodation of differential likelihoods of inputs over their range is the essence of step 4 -- in fact, of Monte Carlo simulation.

Histogram: A histogram is the most common means to display differential likelihoods. It is a bar chart of the desired distribution. To obtain it, we first divide the range over which a variable can vary into a number of regions of equal width, called bins. To each bin we then allocate the fraction of realizations of the uncertain number that we would expect to find in it if we sampled many times. We then scale the height of the bar to represent the frequency.

Figure D.11 compares the theoretical histogram for the “RAND()” function in Microsoft Excel with an experimental histogram based on 10,000 samples from “RAND()”. These histograms are quite although not perfectly similar. They do not exhibit differential likelihoods. They are equal for practical purposes.

Figures D.11 and D.12 about here

The histogram in Figure D.12 assigns different likelihoods to different regions of the range over which the uncertainty varies. Values in the middle are more frequent than values at either end. This shape is called a triangular distribution.¹⁶ Symmetric triangular distributions are easily implemented in Excel as the sum of two uniform distributions.¹⁷ In some cases it may make sense to use a non-symmetric triangular distribution, with the peak not in the middle of the range. Generating such distributions with the RAND() function, although possible in theory, is more cumbersome as Box D.4 indicates. This is where Monte Carlo software packages add value; they make it easy for users to generate uncertain numbers from many different shapes of distribution.

Software catalogs of histogram shapes provide a versatile tool to express uncertainty in inputs. However, the greater the choice, the more difficult it is to choose between the shapes. In this connection one should keep in mind that working with some shape of the uncertainty is better than assuming that there is no uncertainty, as we do when working with base case, deterministic valuation models. A deterministic input is equivalent to an extreme case of Monte Carlo simulation, where the histograms of all uncertainties consist of a single bar at the base case value. Even a slightly spread-out distribution will be more realistic than the single bar.

Although a wrong distribution is better than no distribution, it is good practice to perform sensitivity analysis on the shapes. We can explore the effect of several distributional assumptions by running Monte Carlo using each of them. Figure D.13 for example shows the difference between using a uniform uncertainty, i.e. a simple RAND() function, and a triangular uncertainty for the uncertain inputs in the parking garage example. Different shapes of uncertainties do lead to differences in system performance. Extreme outcomes are less likely in the case of a triangular distribution, both along the vertical axis leading to a narrower cluster of clouds, and along the horizontal axis, with more points clustered in the middle and fewer points at either end. However, the overall effect of the different distributions is, qualitatively, relatively mild. The uncertainty in system performance is significant, whether one uses a uniform or a triangular distribution.

Figure D.13 and Box D.4 about here

Generating output shapes: It is natural to think about the shape of the distribution of the uncertain performance. Calculating these output shapes from given input shapes is the essence of Monte Carlo simulation. Traditional spreadsheet models are “numbers-in, numbers-out” models; Monte Carlo spreadsheets are “shapes-in, shapes-out” models.

The histogram of the outputs can be generated directly from the recorded results of the Monte Carlo simulation. Monte Carlo simulation add-ins can be very useful in the construction of these graphical outputs. The facilities in Excel to create histograms are somewhat cumbersome.

A target curve, on the other hand, is easy to build in Excel. This can be done by sorting the output of interest with values in ascending order and plotting the values against their respective percentiles, calculated from 0 to 1 in ascending steps of $1/n$ if there are n sorted output

values. Alternatively, we can use the PERCENTRANK function to calculate the percentage rank for each output within complete set, and then scatter plotting the output against its percentage rank, as Figure D.14 illustrates.

Figure D.14 about here

Step 5: Introduce dependence between uncertain numbers

This step recognizes the relationships between uncertain inputs. This is the most difficult step. While it is relatively easy to acknowledge the existence of relationships, quantifying them in a model is hard, not least because intuition about the nature of the relationships between uncertainties is often relatively poor.

In the parking garage example, one can assume that demand and annual revenue per space are related. If demand is high, the operator of the parking garage may be able to charge more per space, which will increase the annual revenue per space. This results in a positive relationship between revenue per space and demand. There are two ways to model this relationship. We can either model it directly, or model the driving mechanism for the relationship. For example, in the case of the parking garage, we could capture the driving mechanism with a pricing model that inputs demand over time, calculates appropriate prices we would charge, and thereby produces revenues per space as a function of demand.

A direct model of the relationship could be of the form

$$\begin{aligned} \text{Average revenue/ space} = & \text{sampled annual revenue/ space} \\ & + b * \text{demand deviation from projection,} \end{aligned}$$

where b is a parameter that needs to be determined sensibly. If we choose $b = 0$ then the average revenue per used space is sampled independently of demand, as in Step 4. The second factor corrects this by taking demand deviation from projection into account. So what should b be? First, b should be positive since intuitively the relationship between the two inputs is positive, i.e. the higher demand, the higher revenue per used space. Second, parameter b reflects the price elasticity of demand, which is itself an uncertain input in the model. It should therefore have a distribution rather than a single value. If there are data on price elasticity of demand for parking garages, these may be useful in estimating this distribution.

Notice that the above model does not change the average revenue per space as long as the price elasticity parameter b itself is independent of the demand deviation from projection. This is because the average demand deviation from projection is zero¹⁸. However, relationships between the uncertain variables can significantly affect the shape of the distribution of the output, as Box D.5 indicates.

Box D.5 about here

Common undercurrents cause relationships between uncertain inputs: Relationships between uncertain inputs are often a consequence of common undercurrents, of global drivers further up the causal chain that are not directly included in the model. For instance, the state of the economy, characterized by metrics such as GDP growth, is a common undercurrent that affects the performance of most systems. If the economy thrives, then both the demand for our parking garage and its operating costs may go up. This induces a positive relationship between costs and demand.

It is useful to explore such common causes. First, make a list of the key undercurrents that might simultaneously affect the uncertain inputs. Then capture the qualitative nature of this effect in a matrix, as Figure D.15 illustrates. The columns correspond to undercurrents; the signs in the matrix denote the anticipated effect on the uncertain input as “higher than expected” (+) or “lower than expected” (-), respectively.

Figure D.15 about here

An undercurrent matrix is a useful tool to start a discussion about relationships between uncertainties. Its development is a pragmatic rather than scientific exercise. Although data can and should be used as much as possible, the critical challenge is prioritization, the determination of a manageable list of key undercurrents from a vast number of possible variables that may affect the uncertain inputs. This requires expert judgment and context knowledge.

Once a list of key undercurrents is determined, we can use them as additional, hidden inputs in our model and determine the uncertain inputs y_i via equations of the form

$$y_i = a_{i0} + a_{i1}x_1 + \dots + a_{in}x_n + e_i,$$

where x_1, \dots, x_n are the undercurrents and e_i is a residual uncertainty that accounts for uncertainty in the variable y_i on top of the undercurrents. Such models can easily be implemented in a Monte Carlo spreadsheet, provided we have estimated coefficients $a_{i0}, a_{i1}, \dots, a_{in}$, and agreed on distributions for the inputs x_1, \dots, x_n and e_1, \dots, e_m . The variables y_i are then calculated using the above formula. It is appropriate to seek statistical advice at this point.

A word of caution is in order. It would be wrong to get the impression that relationships between undercurrents and uncertain inputs can be straight-forwardly established. This can be a lengthy debate, hopefully informed by research and data. The relationship may well not be easily captured by a + or – sign, let alone a linear equation as assumed above. However, neglecting relationships can be worse than getting their magnitudes slightly wrong.

Step 6: Introduce dynamically changing uncertain numbers

For the modeling of flexibility it is important to acknowledge explicitly that many uncertain inputs to the system model will evolve over time and that we can use the flexibility as the uncertainties unfold. We therefore need to spend some time on dynamic models of uncertainty.

As Appendix E discusses, the classical dynamic model of uncertainty is a random walk. The simplest example of a random walk is a repeated coin flip, for example where you pay \$1 when the coin shows head and gain \$1 when it shows tail. The random walk keeps track of your profit and loss over time. In mathematical terms, a simple random walk is a sequence of uncertain numbers where $X(t)$ evolves from $X(t - 1)$ by adding an uncertain shock $\epsilon(t)$

$$X(t) = X(t - 1) + \epsilon(t),$$

where $X(0)=x_0$ is known and $\epsilon(t)$ are independent uncertain numbers, typically with the same distribution. In the example of the coin flip, $x_0=0$ and $\epsilon(t) = 1$ with probability $\frac{1}{2}$ and $\epsilon(t) = -1$ with probability $\frac{1}{2}$.

Random walk models provide a more realistic view of future scenarios. Consider uncertainty in demand growth for the parking garage example. In Step 3 we modeled its uncertainty by sampling a random deviation within +/- 50% of the base case demand and then calculating the corresponding growth curve. The scenarios were thus all within a limited band around the base case, which implied a rather smooth growth projection. It is likely that the actual growth will be much more uneven. To simulate the more realistic situation, we can use a random walk model. The initially sampled growth curve will only give us the trend and the random walk will modulate around this trend.

Random walks are very easy to implement in spreadsheets. We begin by sampling the smooth growth curves, giving expected demands d_1, d_2, \dots, d_{15} over the 15 year planning horizon. Then we start the random walk process with $X(1) = d_1$ and use the recursion

$$X(t) = X(t - 1) + \epsilon(t), \quad t = 2, 3, \dots, 15,$$

where $\epsilon(t)$ is a random shock with an expected value of $d_t - d_{t-1}$.¹⁹ This guarantees that the expected value of $X(t)$ is d_t . However, $X(t)$ fluctuates around this expected value, $X(t)$ fluctuates around this expected value. Figure D.16 shows 10 realizations of demand paths, with normal random shocks, and illustrates how the realizations modulate around the average growth curve. It also shows that the demand realizations spread out over time, that is, uncertainty in demand grows with increased time. This is sensible. It is more difficult to predict the more distance future. Notice that this was not the case in our original model, which only modeled the average growth curves. The models can of course be combined by first generating an average growth curve and then a random walk modulation around this curve. Appendix E provides more details on the specifications of dynamic input distributions.

Figure D.16 about here

Step 7: Using rules for exercising flexibility

Steps 1 – 4 turned a fixed-number model with base case inputs into a Monte Carlo simulation model that allows us to calculate with shapes. It converts input shapes into output shapes and allows us to explore the effects of uncertainty, of the variation around the base values as well as

of the dependence between variables. Such probabilistic models are necessary for a systematic valuation of flexibility because flexibility, by its very nature, is only exercised in certain scenarios. If we do not have a way to simulate the scenarios, the value of flexibility is invisible.

To value flexibility within a probabilistic model, we have to tell the model when to use this flexibility. This is the last part of the modeling exercise. In the parking garage example we need to tell the model when to expand and by how much. A simple way of doing this is to use the Excel IF functions. For example, we could stipulate that we should add an extra level to the garage if it ran at its effective capacity for the past two years. To do this, we simply need to keep track of capacity utilization, year on year, and the IF function would trigger the addition when demand met the stated conditions.

Figure D.17 shows a modification of the NPV spreadsheet in Figure D.1 that includes such a rule for exercising flexibility. Row D6:P6 contains the IF function statements. We have not included expansion in year 1 or year 15, on the grounds that we could not meet the condition in the first year, and would not want to expand in the last year. For example, cell E6 contains the formula

`"=IF(AND(D4<MAX_CAP,MIN(D4,D5)+MIN(E4,E5)=D5+E5),"expand","")"`.

This statement has two conditions. The first is that the number of levels is not yet at its maximum; MAX_CAP refers to a cell that states the maximum number of floors. The second condition guarantees that demand during the past two years was above the garage's effective capacity. When the conditions are met, building cost is incurred (E11) and the capacity becomes available the following year (F5). In this case, the rule for exercising flexibility stipulates that only one level at a time would be built. Because demand grows rapidly in the scenario in Figure D.17, it leads to expansions in both years 3 and 4. At the end of the planning horizon in year 15, the garage has been expanded to 9 levels.

Figure D.17 about here

With a rule for exercising flexibility in place, the performance of the garage depends not only on the design parameters, such as the initial number of levels, but also on our choice of the rule. The optimization of today's actions is complemented by an optimization of our anticipated future actions. Notice that it is very important that the rule is only based on information available at the time of the decision. A rule that would exercise expansion in year 4 on the basis of demands in years 5 and 6 is not allowed.

Final comments

Expectation consistency: When we turn a static base case model into a probabilistic model for Monte Carlo analysis, we have to make sure that we remain expectation consistent, i.e. that the expected values for our inputs stay fixed at the base case values. Otherwise we start comparing apples with pears. If, for example, we implemented a demand model that would lead to higher

demands, on average, over the time horizon of the parking garage, then it would not be surprising to see that the parking garage is worth more, on average, on the basis of the probabilistic model than on the basis of the static model.

To check expectation consistency, analysts should record both the model outputs of a Monte Carlo simulation and the associated inputs. They should then compare the input averages with the corresponding base case values. Expectation consistency requires these values to be close to one-another.

Trials needed for a robust Monte Carlo simulation: The goal of a Monte Carlo simulation is to perform a shape-in, shape-out calculation; to approximate the distributions of uncertain outputs, given distributions for uncertain inputs. It is important to understand that Monte Carlo simulations only provide approximate results. They could only provide precision after an infinite number of trials. The more trials, the more accurate the approximate result – but it will never be precise. Therefore, a critical question is: how many trials do we need for practical purposes?

To get a first idea of the accuracy of the Monte Carlo process, we can calculate the accuracy of our estimate of the mean of the distribution of the output. Clearly, a good approximation of the mean is a necessary condition for a good approximation of the whole distribution. This approximation of the mean is the average of the generated output trials. An important statistic for its accuracy is its standard error, which is the standard deviation of the generated output trials divided by the square root of the number of trials. Elementary statistics tells us that the following formula defines the 95% confidence interval (CI) for the mean

$$95\% \text{ CI} = \text{Average} + 2 * \text{Standard Error}.$$

What does a 95% confidence interval signify? It refers to the fact that we can expect 19 out of 20 Monte Carlo simulations to produce 95% confidence intervals that contain the actual (unknown) mean of the distribution of the uncertain output. Indeed, repetitions of the Monte Carlo process sample different inputs, and produce different values for the approximation of the output distribution, and therefore different approximations of the mean and different confidence intervals for the mean. Thus, a minimum requirement for accuracy is that the 95% confidence interval is small, that is, that the standard error of the mean is small relative to the mean. If this is not the case, we need to run more trials.

Note that an accurate estimate of the mean is not a sufficient criterion for an accurate estimate of the distribution. To estimate the probability that the output value falls into any specific region within its range, for example into a specific bin of its histogram or below a certain target value, we can again use the standard error of the mean. This probability is the average of a “counting” random variable obtained directly from the output, assigned the value 1 if the sampled output value falls in the specified region and the value 0 if not. We estimate a 95% confidence interval for this probability using the standard error of the counting variable. If P is the proportion

of trials that fall into the region, then the standard error of the counting variable is given by the square-root of $P^*(1-P)/n$, where n is the number of trials. Figures D.18 and D.19 illustrate the calculation of confidence bounds for a target curve.

Figures D.18 and D.19 about here.

Box D.1

Modeling vocabulary

Input: Numbers that describe the environment that the system will face

System parameters: Numbers that describe the system design

Outputs: Numbers that describe the performance of the system

Box D.2

Sensitivity analysis vocabulary

Sensitivity analysis: replaces fixed assumptions on inputs by ranges on inputs and produces graphs that show how system performance changes as individual inputs change over their range, with all other inputs fixed at their base case values.

Tornado Diagram: A bar chart that summarizes the effects of the changes of variables across specified ranges .

Box D.3

Monte Carlo simulation vocabulary

Uncertain inputs: The cells in the spreadsheet whose uncertainty significantly affects system performance.

Input distribution: Distribution of the uncertain inputs, including their relationships and dynamic evolution. The determination of a suitable distribution for the uncertain inputs is the result of a dynamic forecasting exercise, as Appendix E explains.

Output distribution: Distribution of performance metrics as a function of the input distribution.

Trial: A single run of a Monte Carlo simulation model, sampling one input combination from the defined distribution and recording the associated values of all relevant output cells, calculated by the valuation model.

Monte Carlo simulation: A list of many sampled input combinations and associated calculated output metrics, ready for statistical analysis and graphical display.

Box D.4**Generating distributions in Microsoft Excel**

Microsoft Excel has two random number generators: RAND() for a continuous variable between 0 and 1, and RANDBETWEEN for a discrete variable between any two specified integers. These functions allow us to generate a variety of other distributions. There are three main ways to do this.

Inverse Transform Method: Microsoft Excel has several inverse cumulative distribution functions (ICDF). For example, the NORMINV function is the ICDF of the normal distribution, GAMMAINV is the inverse of the so-called Gamma distribution, etc. We can use such functions to generate samples from the distribution, using the RAND() function. In fact, if FINV is the ICDF of a random variable X with cumulative distribution function F, then “=FINV(RAND())” samples from the distribution of X.²⁰ In that sense, the RAND() function is the mother of all distributions. For example the formula “=NORMINV(10,2,RAND())” generates a normal distribution with mean 10 and standard deviation 2. Likewise, the formula “=-A1*LN(RAND())” generates an exponential distribution with mean in cell A1.²¹

Sampling from a user-defined distribution: This approach uses a histogram, that is, a list of numbers (the mid-points of the histogram bins), and associated frequencies. The following example defines values in C3:C6 and associated probabilities in D3:D6.

	A	B	C	D	E
1					
2		Vlookup values	Value	Probability	
3		0%	10	12%	
4		12%	20	17%	
5		29%	50	37%	
6		66%	120	34%	
7				100%	Total
8	Sample		120		

We can combine the VLOOKUP and RAND() functions to sample from these values with the associated probabilities. To do this, set up a column with cumulative probabilities to the left of the value column C. This is done in B3:B6. It is important to start with 0% against the lowest value, i.e. the cumulative distribution values have a lag of 1, the 29% probability that the value is less than 20 is not put against the value 20 but against the next number up, the 50 in this case. Cell C8 then contains the formula “=VLOOKUP(RAND(),B3:C6,2)” and samples from the specified distribution.

Sampling from historical data: We can similarly sample directly from data of past occurrences of the uncertain input. For example, an important uncertain input for hospital operations is the length of stay. Suppose you have this data for the past 1000 patients. Input it into a spreadsheet, labeling each record consecutively as in columns B and C below. You can then sample from the data with the function “=VLOOKUP(RANDBETWEEN(1,1000),B2:C1001,2)”, which is the formula in cell F2.

	A	B	C	D	E	F
1						
2		Patient	Length of Stay		Sample	5
3		1	5			
4		2	3			
5		3	4			
6		4	3			
7		5	3			
8		6	1			
9		7	2			
10		8	2			
11		9	6			

Sampling from historical data is only sensible if the process is reasonably stationary, that is, if the past is a good predictor for the future. For example if you plot length of stay over time and observe that it tends to reduce, then you should modify the historical data to capture this trend before you can properly use it to sample future length of stay.

Box D.5

Why relationships between uncertain inputs matter

Relationships between input variables can significantly affect the shape of the output of models. To illustrate this, consider the simplest of all models, the sum of uncertain inputs. Note that the expectation of the sum of uncertain summands is the sum of their expectations, regardless of the relationships between them. However, whilst average of the histogram of the model output remains fixed when relationships are introduced, its shape can change dramatically.

When we add up unrelated uncertain numbers, the shape of the sum will be more peaked in the middle than the shape of the summands.²² To get an extremely high or low sum, we have to be lucky enough to sample only high or only low summands. Sums in the middle are more likely because they can be a mix of ups and downs of individual summands. The histogram of a sum peaks in the middle.

This peaking effect is affected when the summands are related to each other. To see this, consider the sum of two spreadsheet cells, A1 containing the formula “=RAND()”, and A2. Thus cell B1 contains the formula “A1+A2”. The content of A2 is a random number more or less related to A1..

- Suppose first that A2 also contains the formula “=RAND()” and is independent of A1. When you do a Monte Carlo simulation you will find that the shapes of the uncertain numbers in A1 and A2 are both flat but that the shape in B1 is triangular, that is, it peaks in the middle as expected.
- Now suppose that A2 “=A1”. Both A1 and A2 contain uncertain numbers and their shapes are both flat. However, the shape of their sum in B1 is now also flat, a stark difference from the previous triangular shape. The reason is that the uncertain numbers in A1 and A2 are now perfectly positively related. The mixing of high and low realizations of A1 and A2 to get a result in the middle no longer happens. If A1 is high, so is A2. In this extreme case, the peaking completely disappears. In less extreme cases, when there is a less than a perfect positive relationship, the peaking is weakened.
- Finally, suppose A2 “=1-A1”. A1 and A2 are flat shapes as before. However the shape of B1 is now extremely peaked. B1 now always is 1, no matter what. The uncertain numbers in A1 and A2 have a perfect negative relationship. Whenever we sample a low A1, it is balanced out by a correspondingly high A2 which leads to the sum of 1. In the extreme case of a perfectly negative relationship, the peaking becomes maximal. In less extreme negative relationships, the sum exhibits a stronger peaking than in the case of independent summands.

In summary, when summands are unrelated, the shape of the sum peaks more than the shape of the summands. Negative relationships between the summands amplify this peaking effect, positive relationships weaken it.

Relationships between input variables can also affect the average output, thus inducing a Flaw of Averages. For example, the expression: net revenue = margin * sales volume. If margin and sales volume are unrelated, then the average revenue equals the product of average margin and average sales volume. However, if margin and sales are related, as one would expect because demand drives them both, then this identity fails. If demand for a patent protected product is higher than expected then the company can charge higher prices, and obtain both higher margins and higher sales volume -- a positive relationship. However, if demand for an unprotected product is higher than expected, this can lead to more competitors in the market, fiercer price competition, and depressed margins. Meanwhile, the consolidated marketing effort of all competitors may increase market size and higher than expected sales volume. This dynamic induces a negative relationship between sales volume and margin.

When margin and sales volume are positively related, the expected net revenue will be larger than the product of expected margin and expected sales volume. If the relationship is negative, then the inequality is reversed, expected net revenue is smaller than projected. To illustrate this effect in a spreadsheet, take the distribution of margin in cell A1 as “=RAND()” and set sales volume in cell A2 as “=A1+RAND()” for a positive relationship, in cell B2 as “=2-margin-RAND()” for a negative relationship, and in cell C2 as “=rand()+rand()” for no relationship. In all three cases, the individual distributions of margin and sales volume are the same; margin is uniformly distributed between 0 and 1 and sales volume has a triangular distribution between 0 and 2. However, both the shape and the average of the distribution of net revenue, the product of the respective cells, are quite different.

INPUT TABLE					
DEMAND PROJECTION					
Demand in year 1	750	spaces			
Additional demand by year 10	750	spaces			
Additional demand after year 10	250	spaces			
REVENUE					
Average annual revenue	10,000	per space used			
COST					
Average operating costs	3,000	per space available			
Land lease and other fixed costs	3,330,000	annual			
Capacity cost	17,000	per space			
	10%	growth/level > 2			
DISCOUNT RATE					
	10%				
SYSTEM PARAMETERS					
Capacity per level	200	cars			
Number of levels	6	levels [DESIGN PARAMETERS]			
PERFORMANCE CALCULATION					
Year	0	1	2	...	15
Demand		750	893	...	1,634
Capacity		1,200	1,200	...	1,200
Revenue		7.5	8.9	...	12.0
Operating costs	0.0	3.6	3.6	...	3.6
Land leasing and fixed costs	3.3	3.3	3.3	...	3.3
Cashflow	-3.3	0.6	2.0	...	5.1
Discounted cashflow	-3.3	0.5	1.7	...	1.2
Present value of cashflow	26.7				
Capacity cost for up to two levels	6.8				
Capacity costs for levels above 2	17.4				
Net present value	2.5				

Figure D.1. Inputs, system parameters, and performance calculation (NPV) for parking garage with 6 levels.

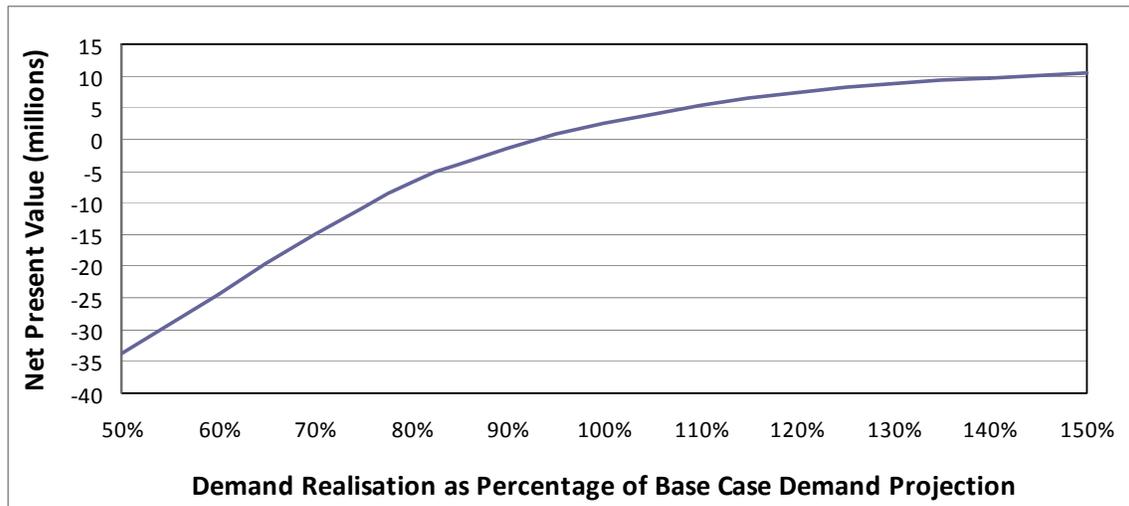


Figure D.2. Sensitivity chart for parking garage.

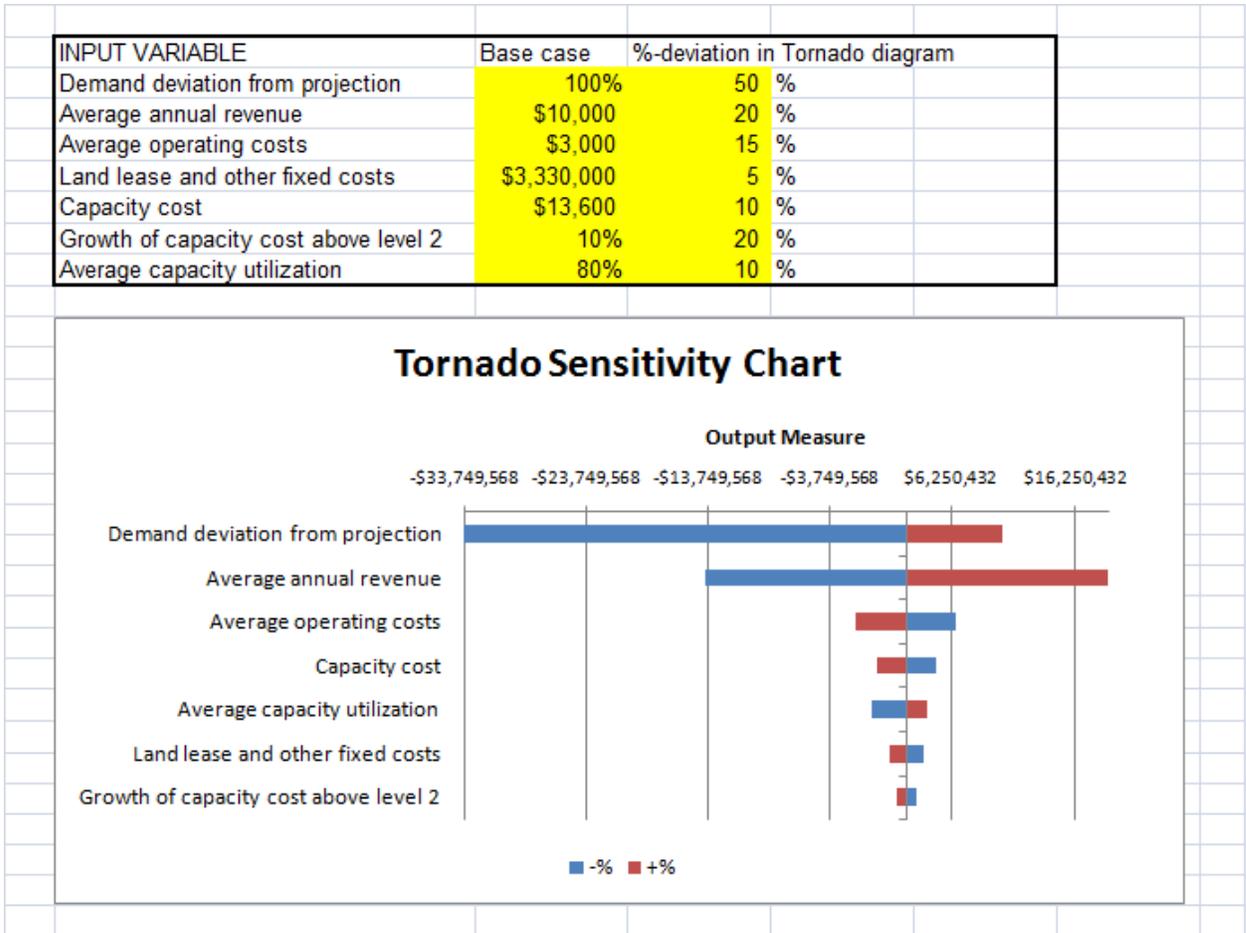


Figure D.3. Tornado chart for parking garage.

Trial	Demand deviation	Average annual revenue	Average operating costs	Land lease and other fixed costs	Capacity cost	Growth of cap cost above 2 levels	Average capacity utilization	NPV
1	134%	\$11,016	\$3,140	\$3,495,435	\$12,399	9.88%	79.69%	\$17,625,663
2	141%	\$11,852	\$3,260	\$3,304,375	\$13,795	9.84%	79.06%	\$23,221,598
3	86%	\$10,193	\$3,312	\$3,430,938	\$13,068	10.43%	79.50%	-\$4,718,694
4	83%	\$9,055	\$3,283	\$3,354,325	\$14,036	9.90%	76.09%	-\$15,456,972
5	83%	\$10,114	\$3,047	\$3,316,798	\$12,685	11.52%	87.40%	-\$3,934,426
6	62%	\$8,111	\$3,283	\$3,206,261	\$13,100	11.02%	84.27%	-\$36,456,858
7	76%	\$10,515	\$2,647	\$3,386,862	\$14,260	11.81%	83.77%	-\$6,799,900
8	68%	\$8,343	\$3,440	\$3,269,874	\$14,830	9.71%	78.82%	-\$32,839,236
9	114%	\$10,272	\$3,291	\$3,410,315	\$14,852	10.11%	85.44%	\$5,120,553
10	68%	\$10,151	\$3,413	\$3,437,877	\$14,110	10.78%	85.28%	-\$24,320,813
11	68%	\$9,664	\$3,153	\$3,163,520	\$13,841	11.51%	82.62%	-\$21,314,749
12	129%	\$11,189	\$2,636	\$3,430,040	\$14,713	10.59%	75.51%	\$16,118,317
13	132%	\$11,832	\$2,695	\$3,179,077	\$14,469	11.09%	84.10%	\$30,373,102
14	114%	\$10,250	\$2,586	\$3,189,684	\$14,311	9.34%	75.61%	\$9,916,978
15	105%	\$9,754	\$2,641	\$3,218,399	\$14,147	9.14%	72.32%	\$2,261,573
16	63%	\$9,595	\$3,089	\$3,240,407	\$13,222	9.92%	77.66%	-\$22,983,479
17	147%	\$10,881	\$2,702	\$3,374,710	\$13,140	10.25%	79.15%	\$20,546,268
18	91%	\$8,783	\$3,275	\$3,222,284	\$12,731	9.97%	86.09%	-\$10,350,864
19	67%	\$9,320	\$3,182	\$3,477,948	\$14,446	9.57%	81.29%	-\$26,886,841
20	67%	\$8,961	\$3,302	\$3,206,613	\$14,770	9.87%	75.37%	-\$26,433,140

Figure D.4. Data table with first 20 trials of a simulation run for the parking garage example (ranges as in Figure D.3).

Correlation coefficients	Demand deviation	Average annual revenue	Average operating costs	Land lease and other fixed costs	Capacity cost	Growth of cap cost above 2 levels	Average capacity utilization
NPV	0.734039	0.597362	-0.18708	-0.063773723	-0.07688	-0.037357715	0.071097

Figure D.5. Correlation between NPV and uncertain inputs for the parking garage example (ranges as in Figure D.3).

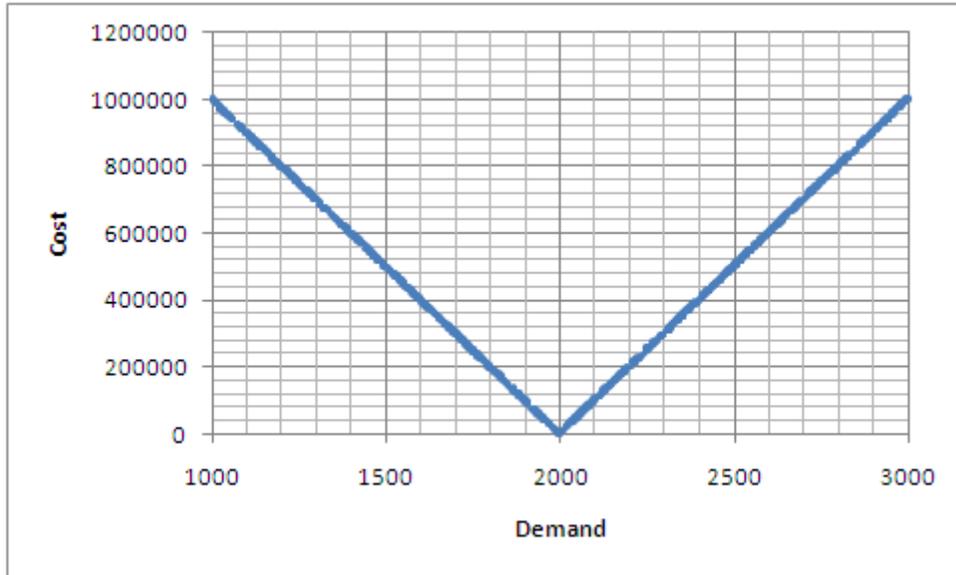


Figure D.6. A nonlinear relationship between inventory cost and demand with a vanishing correlation.

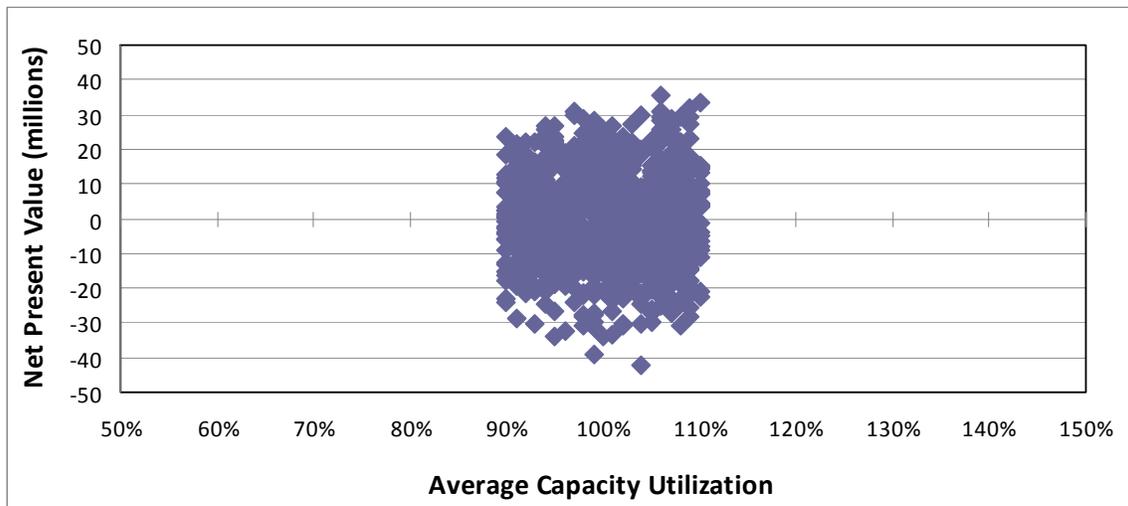
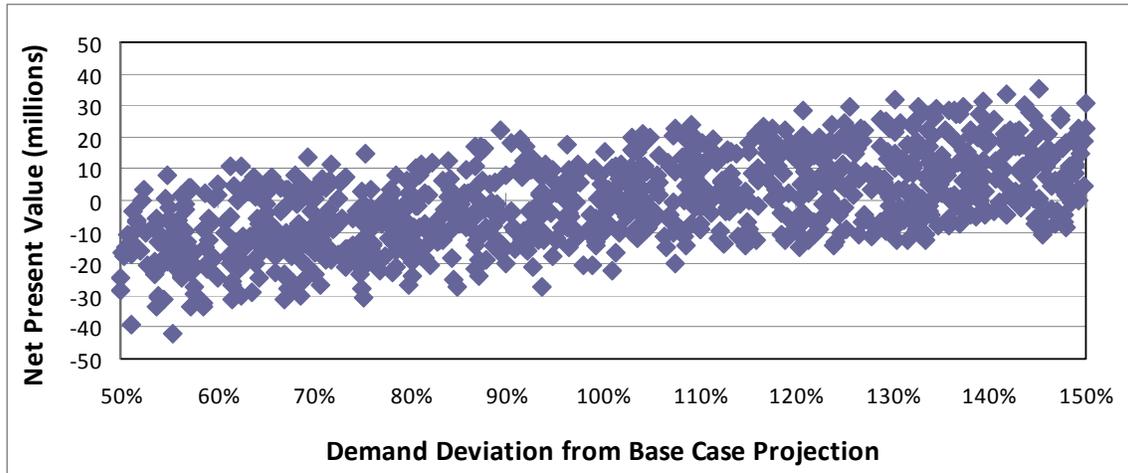


Figure D.7 Scatter plots of NPV versus uncertain input for 1000 Monte Carlo trials. The vertical variation parallel to the y-axis is due to variability of the other uncertain inputs.

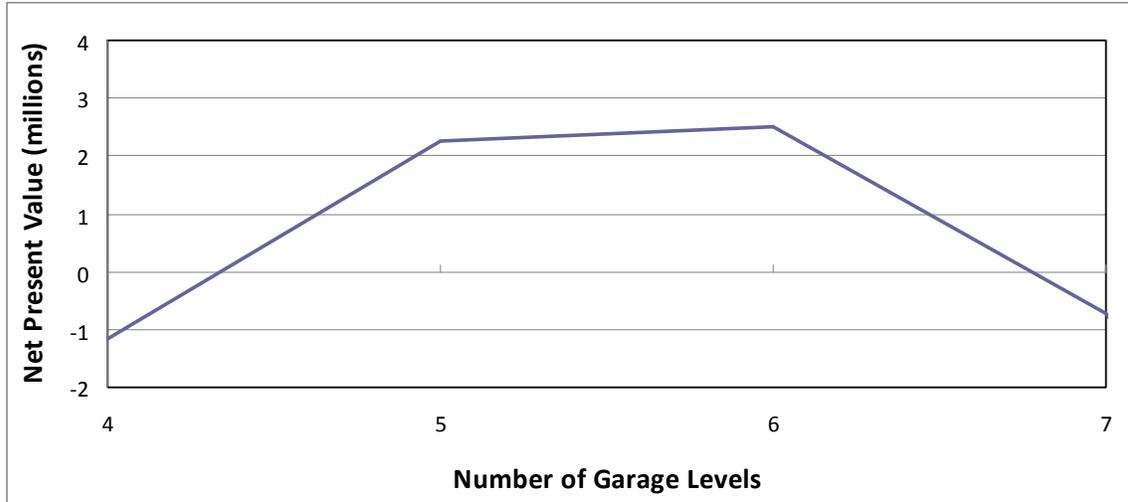


Figure D.8. Optimizing the number of levels of the parking garage.

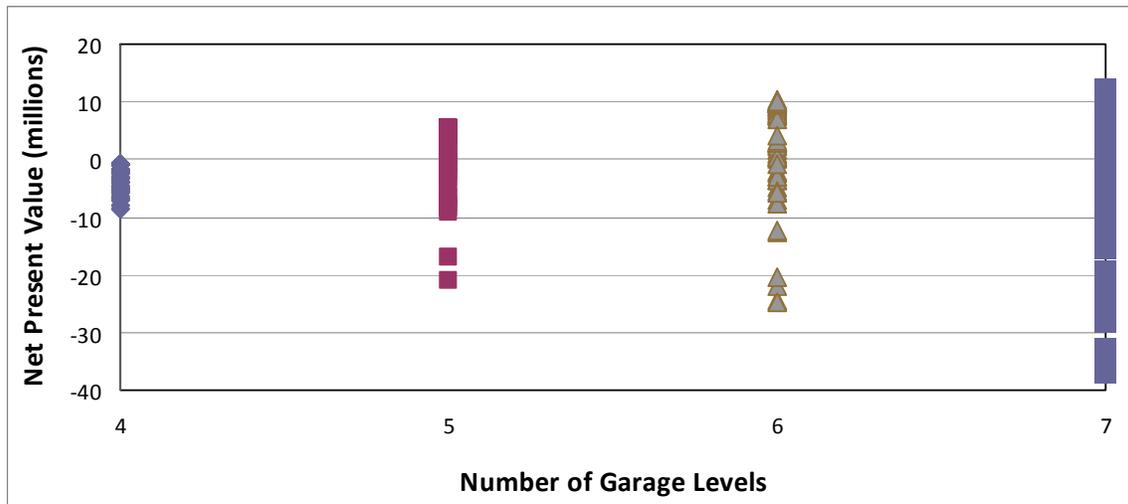


Figure D.9. Optimization under uncertainty. Notice how Figures E.8 and E.9 differ in scale.

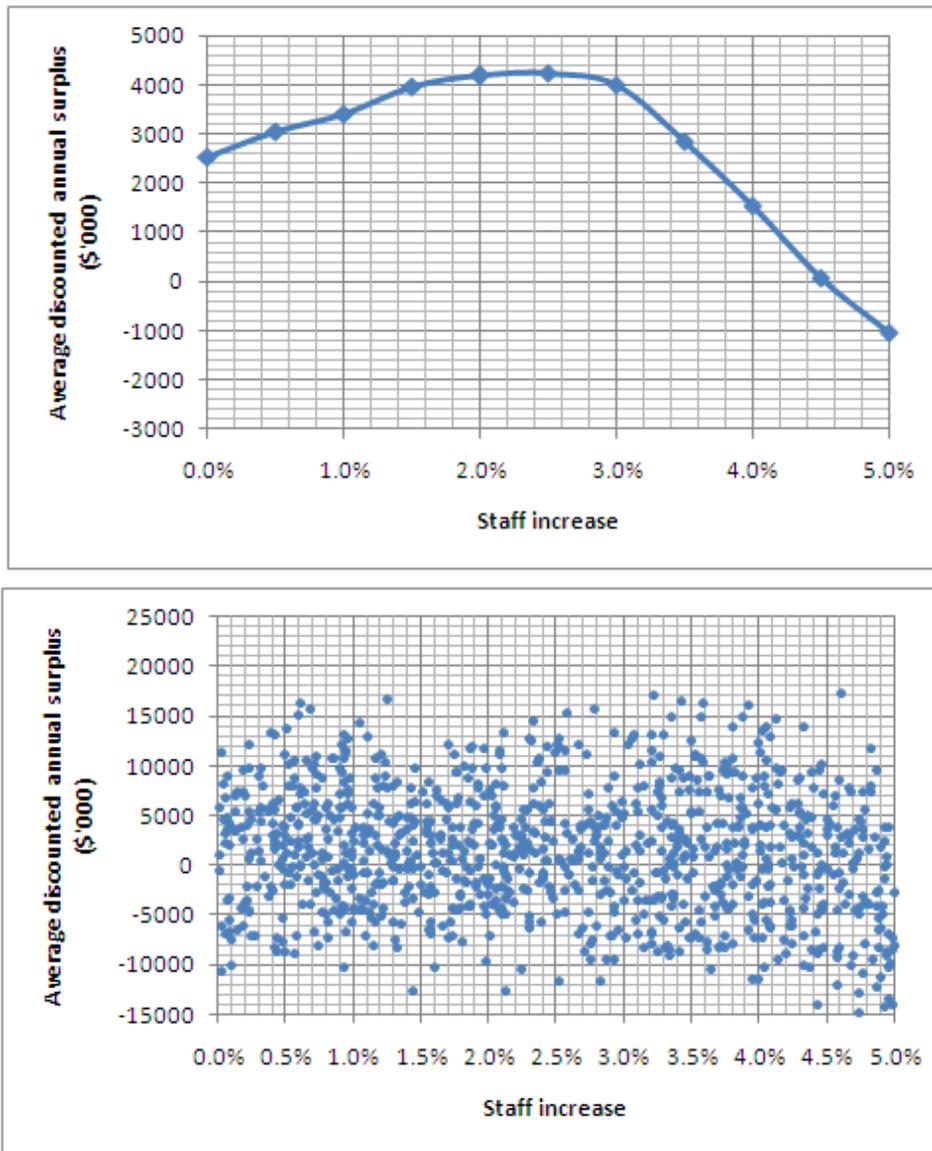


Figure D.10. Example of optimization of staffing level using base case inputs (left) versus probabilistic inputs (right).

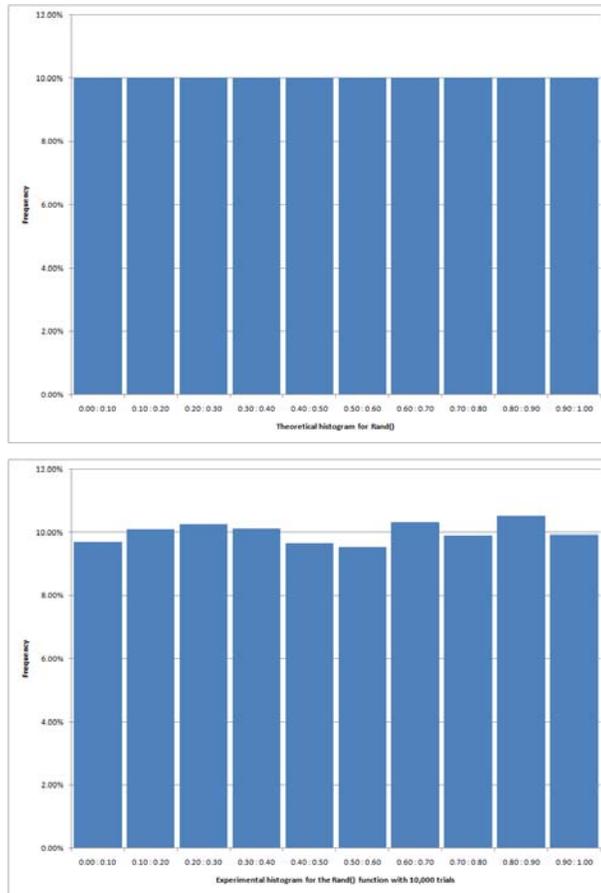


Figure D.11. Theoretical and Experimental Histogram for the RAND() function.

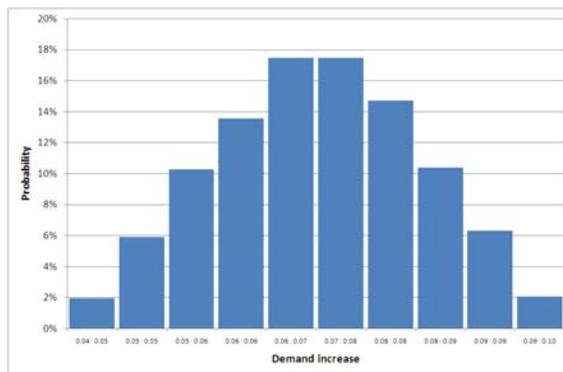


Figure D.12. Histogram of 10,000 realizations of a triangular distribution.

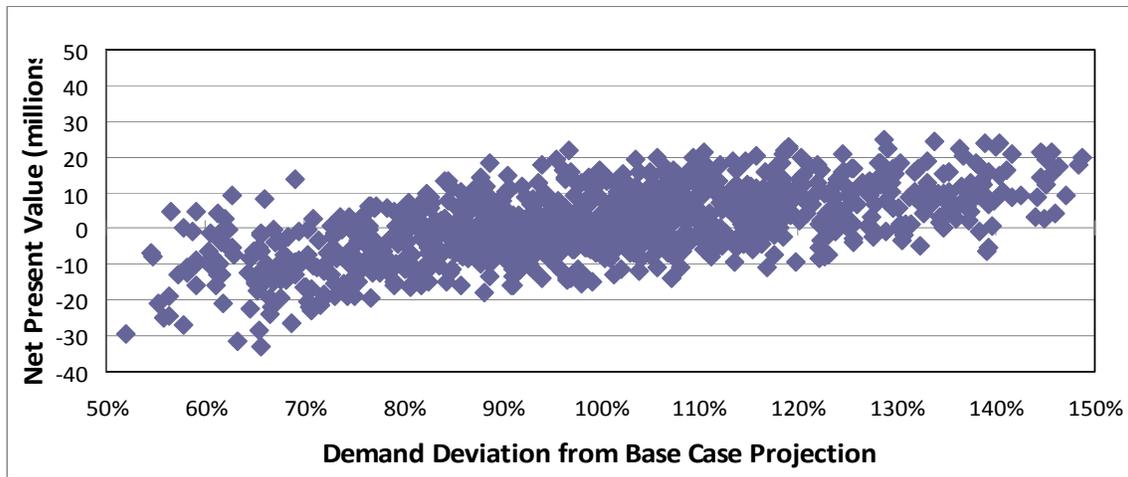
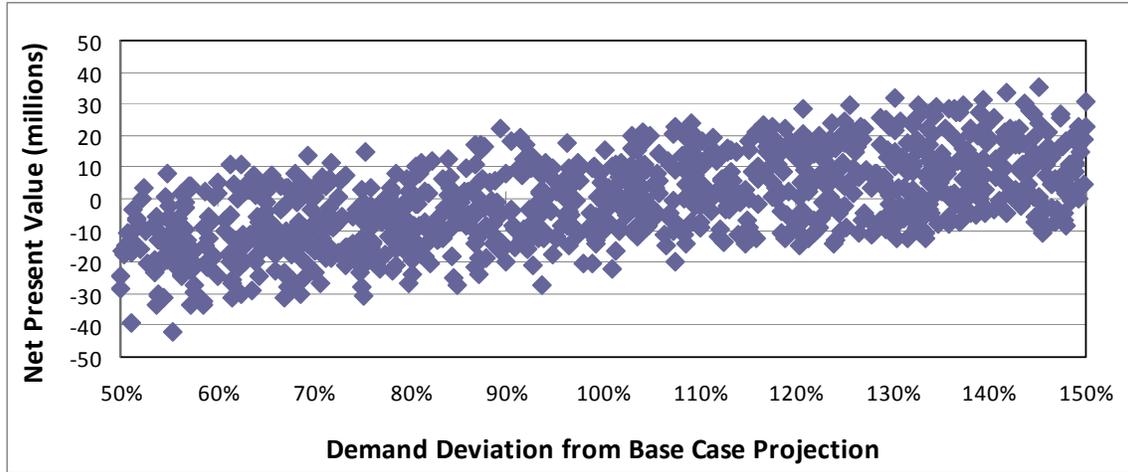


Figure D.13. NPV distribution as a function of deviation from demand projection for uniform uncertainties (left) and triangular uncertainties (right).

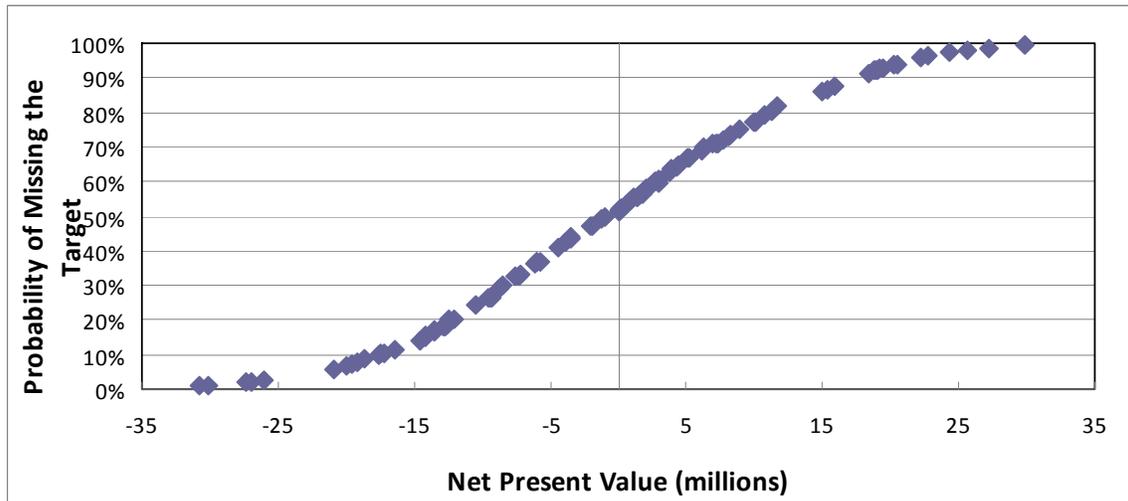


Figure D.14. Target curve generation via scatter plot of output against its percentage rank.

	Economy grows faster than expected	Inflation increases	Tighter immigration rules are introduced
Demand increase	+	-	-
Unit revenues increase	+	+	
Operating costs increase	+	+	+

Figure D.15. Undercurrent matrix for uncertain inputs for the parking garage model.

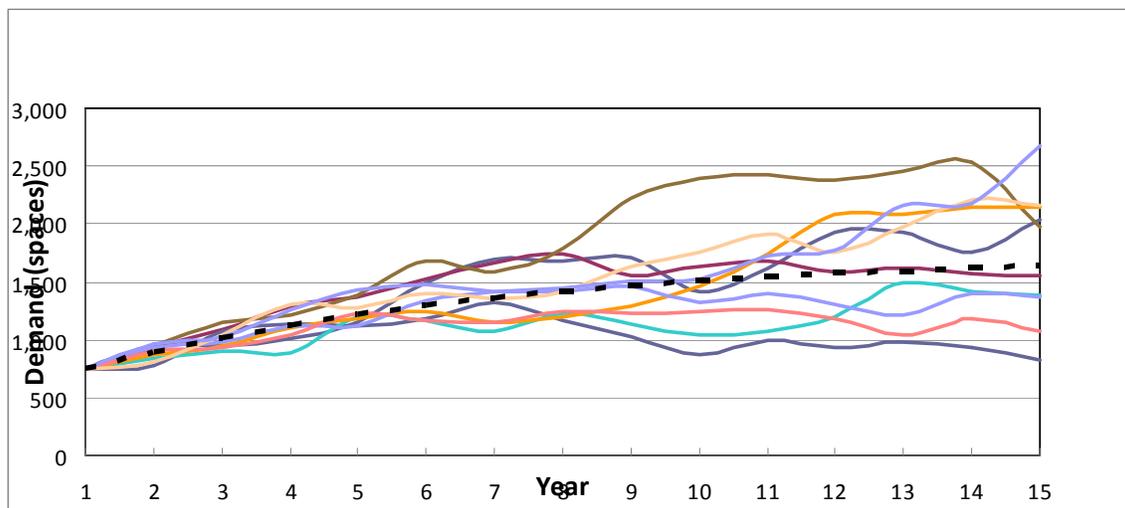


Figure D.16. Realizations of a random walk model for demand.

PERFORMANCE CALCULATION							
Year	0	1	2	3	4	...	15
Levels		4	4	4	5		8
Realised demand		750	893	1015	1120	...	1634
Capacity		800	800	800	1000	...	1600
Expansion?				expand	expand	...	
Build extra capacity		0	0	200	200	...	0
Revenue		7.5	8.0	8.0	10.0	...	16.0
Operating costs		2.4	2.4	2.4	3.0	...	4.8
Land leasing costs	3.3	3.3	3.3	3.3	3.3	...	3.3
Expansion cost	0.0	0.0	0.0	4.5	5.0	...	0.0
Cashflow	-3.3	1.8	2.3	-2.3	-1.3	...	7.9
DCF	-3.3	1.6	1.9	-1.7	-0.9	...	1.9
Present value of cashflow	20.4						
Capacity cost for up to two levels	8.8						
Capacity costs for levels above 2	7.9						
Net present value	3.7						

Figure D.17. Spreadsheet with decision rule for expansion.

	A	B	C	D	E
1	-2.49992	0.01	0	0.0299	
2	-2.22714	0.02	0	0.048	
3	-2.16875	0.03	0	0.064117	
4	-1.64289	0.04	0.000808	0.079192	
5	-1.61701	0.05	0.006411	0.093589	
6	-1.53192	0.06	0.012503	0.107497	
7	-1.32943	0.07	0.018971	0.121029	
8	-1.29886	0.08	0.025741	0.134259	
9	-1.25556	0.09	0.032764	0.147236	
10	-1.25383	0.1	0.04	0.16	
11	-1.18777	0.11	0.047422	0.172578	
12	-1.02597	0.12	0.055008	0.184992	
13	-1.00145	0.13	0.062739	0.197261	

Figure D.18: Calculation of confidence bounds on target curve. Column A: First 13 of 100 observations sampled from a standard normal variable, sorted in ascending order. Column B: Target curve value P associated with sorted observations, ascending in steps of 1/100. Column C: Calculation of lower confidence bound, e.g. C1: “=max(0,B1-2*SQRT(B1*(1-B1)/100)” Column D: Calculation of upper confidence bound, D1=”min(1,B1+2*SQRT(B1*(1-B1)/100)”

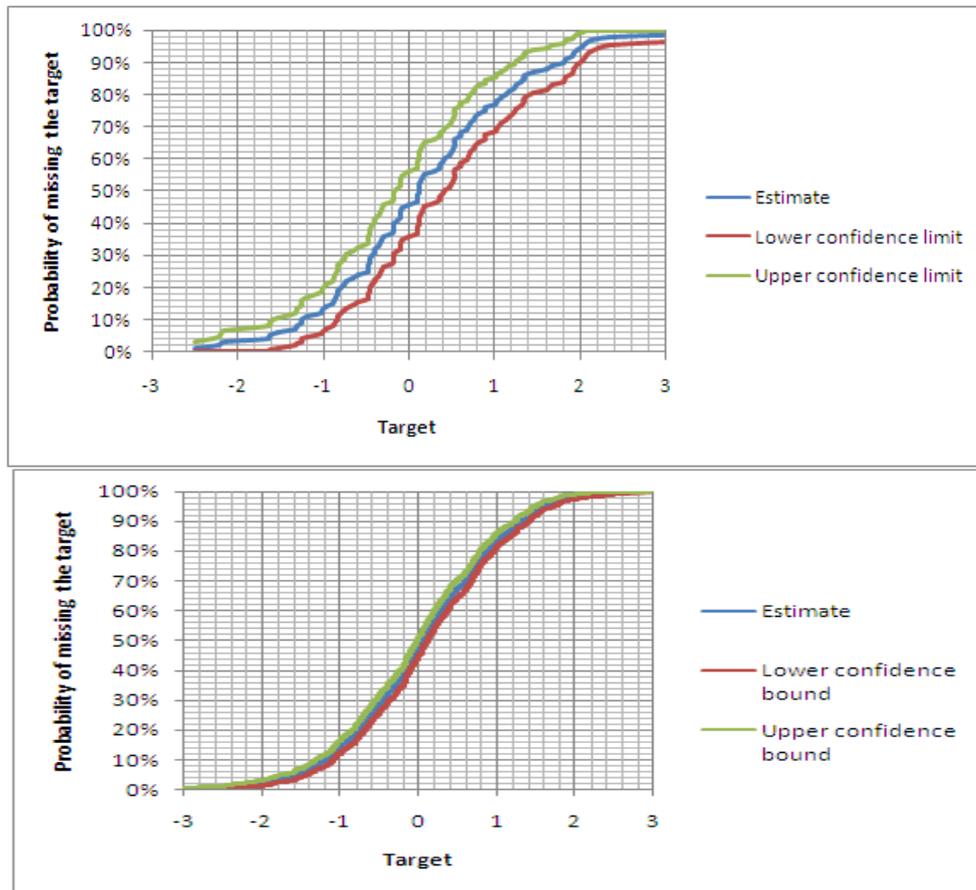


Figure D.19: 95% confidence bounds on target curve sampled from a standard normal distribution (mean 0, standard deviation 1) with 100 trials (top) and 1000 trials (bottom)

APPENDIX E: DYNAMIC FORECASTING

Current forecasting practice produces single-number (or “point”) projections for the future. Such projections, suggesting that it is possible to pinpoint the future, are unrealistic. They do not indicate the level of uncertainty appropriate for the forecast. Thus, they are certainly inadequate for a systematic appraisal of flexibility.

Instead of point predictions, we need a practical way to present the uncertainties around forecasts. Our suggestion is that we should use dynamic forecasts. These are spreadsheet modules of the uncertain environment in which the system operates. This context is characterized by the evolution of uncertain variables, such as demand, costs, prices, and productivity. Dynamic forecasting spreadsheet modules implement the joint distribution of these variables, including dependencies and their evolution over time, where appropriate. We link these modules to our valuation models to drive the uncertain input variables.

An example of a dynamic forecast is the demand module that drives the parking garage case. The purpose of this Appendix is to introduce the main models that are used for such dynamic forecasts and to illustrate how to

- Build forecasting modules in spreadsheets; and
- Calibrate the modules to historical data.

Random walk models

The simplest model of an uncertain variable that evolves over time is a series of coin flips, resulting in a sequence of the form HTHHTTHTTHT. The model is particularly simple because the variable can only take on two values but also, and more importantly, because the variable in period t does not depend on the variables in earlier periods $t-1$, $t-2$,... This is an unusual situation. For most parameters of interest -- such as demand, price, productivity, etc. -- their level at any time will depend on the past.

A stock-flow model is the simplest model of a variable that depends on the past. It is a direct extension of the coin flip model. The variable of interest is the stock, that is, the amount of some variable aggregated over time. In the financial flow version, we focus on the amount of money gained. In that case, the coin flip leads to either a specified gain or a loss, for example with a head leading to a gain of \$1 and a tail leading to a loss of \$1. This model is of the form

$$X(t) = X(t-1) + \varepsilon(t),$$

where $X(0)=0$ and $\varepsilon(t)$ is the financial flow driven by consecutive independent coin flips, i.e. $\varepsilon(t)=-1$ with probability $\frac{1}{2}$ and $\varepsilon(t)=1$ with probability $\frac{1}{2}$. This model is simple to implement in a spreadsheet as Figure E.1 illustrates.

Figure E.1 about here

The stock flow model is versatile: its random component $\varepsilon(t)$ can have any distribution. For example, if we choose $\varepsilon(t)$ to be a normal random variable then we obtain the so-called arithmetic Brownian Motion (also known as a “random walk”) model. We may also wish to model the fact that some factor of $X(t)$ carries over to the next period, which leads to a model of the form

$$X(t) = a \cdot X(t-1) + \varepsilon(t),$$

where “a” the expression of a non-random rate of appreciation or depreciation²³. For example, future demand might be growing at the rate of $r = 5\%$ (thus $a = 1.05$) with some variation around this trend.

Multiplicative models are an alternative to additive models. In this case the error $\varepsilon(t)$ is added to a scaled past period level, and the model is of the form

$$X(t) = \varepsilon(t) X(t-1)^a.$$

When the exponent $a = 1$, this model is called a multiplicative random walk. It can be thought of as a generic model of a bank account with a randomly fluctuating interest rate $\varepsilon(t)$. A multiplicative model reflects random growth, proportional to the existing amount in stock. Because it is multiplicative, if we start with a positive initial level $X(0)$, then future amounts $X(t)$ will never be negative. This is a useful feature of many kinds of variables, such as the demand for or price of some asset, that realistically are never negative.

Note that the multiplicative model turns into an additive model after a log-transformation:

$$\log(X(t)) = a \cdot \log(X(t-1)) + \log(\varepsilon(t)).$$

A typical assumption for the distribution of the error term $\varepsilon(t)$ of a multiplicative model is that it follows a log-normal distribution, i.e. $\log(\varepsilon(t))$ is normally distributed.

If the random error term $\varepsilon(t)$ has a normal distribution with mean μ and variance σ^2 , then the additive model $X(t) = a \cdot X(t-1) + \varepsilon(t)$ has a normal distribution with expectation $E[X(t)] = a^t X(0) + t \cdot \mu$ and variance $V[X(t)] = t \cdot \sigma^2$. If in the multiplicative model the error term $\varepsilon(t)$ is log-normal and the transformed variable $\log(\varepsilon(t))$ has mean μ and variance σ^2 , then $X(t)$ is log-normal with $E[\log(X(t))] = a^t \log(X(0)) + t \cdot \mu$ and $V[\log(X(t))] = t \cdot \sigma^2$.

The simplest models are of the form $X(t) = X(t-1) + \varepsilon(t)$ or $X(t) = \varepsilon(t)X(t-1)$, where $\varepsilon(t)$ has only two values, that is, $\varepsilon(t) = u$ with probability p and $\varepsilon(t) = d < u$ with complementary probability $1-p$. These models are called (additive or multiplicative) lattice models because the values change in a discrete fashion as if moving from vertex to one of the neighboring vertices in a lattice. If we start from some value $X(0) = x_0$ in the additive model, $X(1)$ can only move to either $x_0 + u$ or $x_0 + d$. $X(2)$ can therefore only achieve one of 3 values, $x_0 + 2u$ (twice up), $x_0 + u + d$ (once up, once down), or $x_0 + 2d$ (twice down). These models are thus said to be “recombinant” in that possible outcomes combine in each stage, as in the case of “once up, once down” = “once down, once up”. They thus have the great advantage that the number of possible end points after N stages is $(N + 1)$ rather than N^2 . Generally, $X(t)$ can have $t+1$ values $x_0 + su + (t-s)d$, where s ranges from 0 to t , and assumes a binomial distribution over these values, with success

probability p . The same principle applies to the multiplicative model, where $X(1)$ would have values x_0u or x_0d , $X(2)$ can have values x_0u^2 , x_0du , or x_0d^2 and so forth. These simple lattice models are very useful in the calculation of real options, as Appendix F indicates.

Calibration of random walk models

We can easily implement random walk models in spreadsheets, analogously to Figure E.1. Before doing so, we need to make sure we are doing the right thing. We need to address two questions:

- Are the models appropriate? And
- What values should we choose for the unknown parameters?

Historical data and statistics help with these questions. A detailed exposition goes beyond the scope of this book but an example illustrates the main principles.

At this point, it is important to be clear about what statistics can and cannot do for us. Statistics alone cannot help you determine an appropriate model; they can only be used to test whether a given model is consistent with the observed data. If the model is consistent with the data, then that does not mean that the model is appropriate, it only means that the data does not allow us to refute the model. In general, many models, often with contradictory implications, can fit a given set of data. This is particularly the case when we build models to capture the relationship between variables based on sets of data over time. This is because many phenomena exhibit monotone growth, either upwards or downwards. Such steadily growing and declining variables are naturally correlated, even if there is no causal relationship at all between them. Statistics by themselves are not sufficient to determine an appropriate structural model that captures the relevant causal relationships. Its development is a task for a domain expert who understands the relationships between different factors.

To illustrate how we would go about determining an appropriate model, consider the data in Figure E.2 that represent an uncertain variable, say demand for some service or product, over a period of 35 months. These data do not come from a real-world process but have been simulated from a pre-determined process. The reason we did this is to make the point that we often put too much meaning into data, when the actual cause of the data is random variation.

Figure E.2 about here

It is quite easy in many practical circumstances to develop a convincing story that “explains” some set of data. Suppose in this case that the data are orders of a new product that a company introduced 35 months ago. The “explanation” of the pattern of sales could be the following sequence:

- Initial success due to significant pre-launch marketing;
- Sales fluctuation over the next 10 months (“our competitor started reacting to the new product with mixed success”);

- A more stable period with somewhat lower sales (“our competitor and the market adjusted to the new product”);
- A decline to a low in month 25 (“our competitor launched their own new product”); and
- Significant recovery (“we hired a new chief technology officer who turned things around”).

As convincing as such a story might be, much historical data can equally well be explained as a consequence of pure chance fluctuations (as in this case).

Simplicity of structural models is very important in forecasting, for two reasons. The obvious reason is that simple models are easier to explain to a wide audience. Complex black box models are much less appealing. A secondary, less appreciated, rationale for simple models is that they will typically depend on fewer parameters. When you have many parameters to estimate, and your forecast depends on the forecast of parameters in the model, then these complex models tend to produce worse results than simpler, less “accurate” models.²⁴

In this case, consider the possibility that chance generated the data, according to the additive random walk model described above: $X(t) = a \cdot X(t-1) + \varepsilon(t)$. We can fit this model by scatter-plotting $X(t-1)$ against $X(t)$ and performing a regression to calculate a value for the parameter a . Figure E.3 shows the result. The scatter plot is not convincing, though. It shows somewhat irregular behavior. Many points are clustered close to the lower end of the line. The variation parallel to the y-axis is not the same at the upper and lower end (the upper end is larger). This phenomenon is called heteroscedasticity (a fancy word meaning “different variation”). It renders linear regression problematic because it implies that the error term is not from the same normal distribution, which is a fundamental assumption of regression analysis. In practical terms, it means that the large variation points have an undue weight on the position of the line.

Figures E.3 and E.4 about here

One way to correct for heteroscedasticity is to work with a different model, one that represents the greater effect of higher values. The multiplicative model discussed above is such a model. To obtain it, we perform a log-transformation of the variables. We can then estimate the unknown parameter by regression analysis as before. Figure E.4 shows the result. This model is a better fit to the data. It is specified as²⁵

$$\text{Log}(X(t)) = 0.87 \cdot \text{log}(X(t-1)) + 0.26 + \varepsilon(t),$$

or, equivalently,

$$X(t) = X(t-1)^{0.87} \cdot \exp(0.26 + \varepsilon(t)).$$

What would be a sensible distribution for $\varepsilon(t)$? This distribution is best approximated by the historical errors. The distribution of the errors has a mean of zero (see Figure E.5), and its standard deviation is calculated as 0.22²⁶. We have got a relatively small set of errors, so it is difficult to ascertain that the errors are normally distributed just by looking at the histogram.

However, the comparison of the empirical target curve of the data and the associated normal distribution with mean 0 and standard deviation 0.22 does not give reason to dismiss the normality assumption as illustrated in Figure E.6.²⁷

Figures E.5 and E.6 about here

We therefore conclude that the model $X(t) = X(t-1)^{0.87} \exp(0.26 + \varepsilon(t))$ with independent normal errors $\varepsilon(t)$ with mean zero and standard deviation 0.22 is a reasonable model for the data in Figure E.1, provided that the multiplicative random walk model was sensible in the first place. The actual model that generated the data was $X(t) = X(t-1)^{0.9} \exp(0.1 + \varepsilon(t))$ where $\varepsilon(t)$ are independent shock terms with a normal distribution with mean zero and standard deviation 0.2.

Note that we have calibrated the model on the same data that we use to test it. This is common practice but it is statistically flawed. In the ideal situation is one where you would use one set of data to estimate the parameters, and another set to test model consistency. In practice this is often not done because of the lack of sufficient data. However, if you have sufficient data, it is recommended that you split the data into one set for model calibration and another for model testing.

Seasonality

Random walks of the form $X(t) = aX(t-1) + \varepsilon(t)$ with independent and identically distributed error terms $\varepsilon(t)$ have a fixed trend in the form of the expected error. This fixed trend is sometimes made explicit by writing $X(t) = aX(t-1) + b + \varepsilon(t)$, where the associated error has a mean of zero (as in the above calibration for the example data). Such models can be extended to models of the form $X(t) = aX(t-1) + b(t) + \varepsilon(t)$, where $b(t)$ captures some known seasonal effects. For example, if seasonal effects are assumed to be quarterly, the function $b(t)$ can be chosen of the form $b(t) = b_0 + b_1 * Q_1 + b_2 * Q_2 + b_3 * Q_3$, where $Q_i = 1$ if t lies in quarter i of the year, and $Q_i = 0$ otherwise. The determination of the coefficients can be done with regression, using so-called dummy variables for the quarters.

Mean reversion processes

Random walks have the unpleasant property that they can “blow up” over long time horizons, meaning that they project unreasonably large or small values for the long term. When the model has moved to a level $X(t)$ at time t , then, given this position, it does not know about, has “forgotten” about the old mean $X(0)$. The new mean of $X(t + s)$ given $X(t)$ at time t is now $X(t)$, not $X(0)$. This forgetting about the mean, about the “natural home” of the process, can lead to seriously high or low values over time, to the extent that the generated paths look unrealistic to experts. Put another way, for $X(t) = X(t-1) + \varepsilon(t)$ where $\varepsilon(t)$ has the mean zero and variance σ^2 , then $X(t)$ is a random variable with mean $X(0)$ and variance $t * \sigma^2$. In other words, the mean does not change but the variance increases linearly. Figure E.7 shows this funnel effect.

Figure E.7 about here

A way to avoid this often unrealistic situation is to use a model that tends to center on a mean value, that is, a mean reversion model. To create this effect, we augment a model with a mean reversion term: $r^*(m - X(t - 1))$, where m is a fixed number, the mean or “natural home” of the process, and r is the fixed rate of mean reversion between 0 and 1. Mechanically, we can think of this as a damping effect: if $r = 0$, there is no damping or mean reversion; if $r = 1$, then the process resets to the mean each period.

For example, the mean reversion model corresponding to the additive random walk model has the form:

$$X(t) = X(t - 1) + r^*(m - X(t - 1)) + \varepsilon(t).$$

The mean reversion term $r^*(m - X(t - 1))$ moves $X(t - 1)$ towards m before the random shock $\varepsilon(t)$ is added. This correction towards the mean before a random error term is added implies that the process becomes less likely to move too far from m . Moreover, the further away $X(t - 1)$ is from m , the larger the correction $r^*(m - X(t - 1))$. Figure E.8 illustrates this tendency to move towards the mean in contrast to the random walk in Figure E.7.

Figure E.8 about here

To calibrate the above mean reversion process, we can calculate $X(t) - X(t - 1)$ and regress $X(t)$ on $X(t - 1)$. This will give a model $X(t) = a^*X(t - 1) + b + \varepsilon(t)$. We then match the coefficients with $X(t) = r^*m - r^*X(t - 1) + \varepsilon(t)$, i.e. $r = -a$, and $m = -b/a$.

Models with Jumps

Random walks and mean reversion models change incrementally. Large deviations from $t-1$ to t , although possible, are unlikely. It sometimes makes sense to include the possibility of large deviations, possibly due to special events such as wars, regulation changes, or other disruptions. We typically do this by overlaying a “jump process” onto a smoother process, such as a random walk.

One way to model a jump process is by assuming that the chance of more than one large deviation event is very small in any one period between $t-1$ and t , and can be neglected for practical purposes. This implies the time until the arrival of the next event is exponentially distributed with a given mean, as in Figure E.9. The mean time between two disruptive events is all we need to specify in terms of timing of these events. The exponential distribution has a sensible shape for arrival times and can be easily implemented in Excel (see Box D.4).

Figure E.9 about here

The exponential distribution has an interesting and unique property, called lack-of-memory. It means that the distribution of the chance of something happening is not affected by the length of time you may have already been waiting. In technical terms, the probability of having to wait until time $(t + s)$, given that you have already waited until time t , is the same as the

probability of having to wait until time s from the start. We can therefore at every time period $t = 1, 2, 3 \dots$ draw a sample waiting time to the next event from an exponential distribution. If the waiting time generated at time $t-1$ leads to an event before time t , then we add an additional shock corresponding to the event, otherwise we do not change the process. We can continue with this procedure at time t and so forth, because of the lack-of-memory property.

The effect of the generated jump event needs to be specified as well, and may well be a distribution itself. It may simply add a one-off extra charge to the process, without changing the position $X(t)$ that is need to calculate $X(t-1)$, or change the position of $X(t)$ itself, that is, shift the entire process. The specifics depend on the context.

Time series models of higher order

Random walk and mean reversion processes depend only on the last level of the variable. This may make sense in some situations, but in others, the path by which that last level was achieved can play a role in providing momentum. For example, if prices are dropping one might argue that there is a tendency that they continue to go down. Such path dependence can be modeled with higher order models which are of the form:

$$X(t) = a_0 + a_1 X(t-1) + \dots + a_s X(t-s) + \varepsilon(t).$$

Such autoregressive models can be fit to the data just as random walk models. This leads into the vast domain of time series analysis, coverage of which goes well beyond the scope of this book. The interested reader can refer to standard literature.²⁸

	A	B	C	D
1	Time	Flow	Stock	
2	0		0	
3	1	1	1	
4	2	1	2	
5	3	1	3	
6	4	-1	2	
7	5	1	3	
8	6	1	4	
9	7	-1	3	
10	8	-1	2	

Figure E.1. Stock-flow model implemented in a spreadsheet. Cell B3 contains the formula =if(RAND() < 0.5,-1,1), which generates 1 with probability $\frac{1}{2}$ and -1 with residual probability $\frac{1}{2}$. Cell C3 contains the formula =C2+B3, i.e. the generated flow is added to the stock of the past period. These formulas are then copied down in B4:C10.

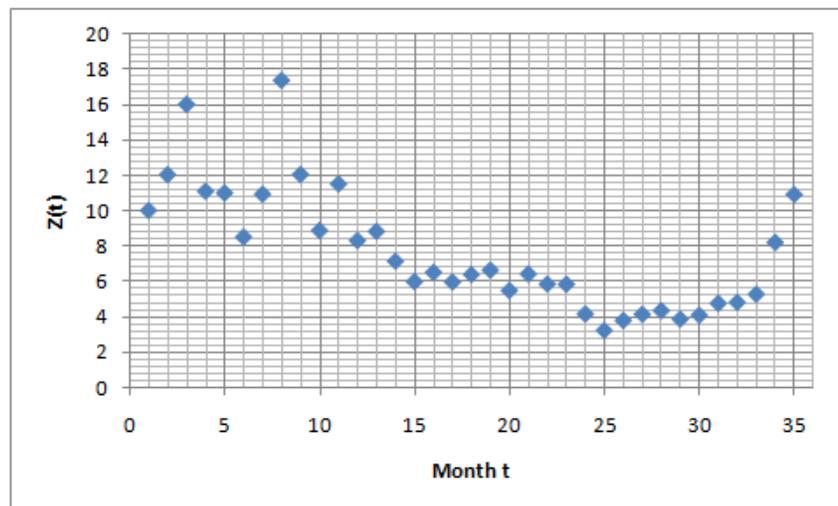


Figure E.2. Monthly data series for 35 consecutive months

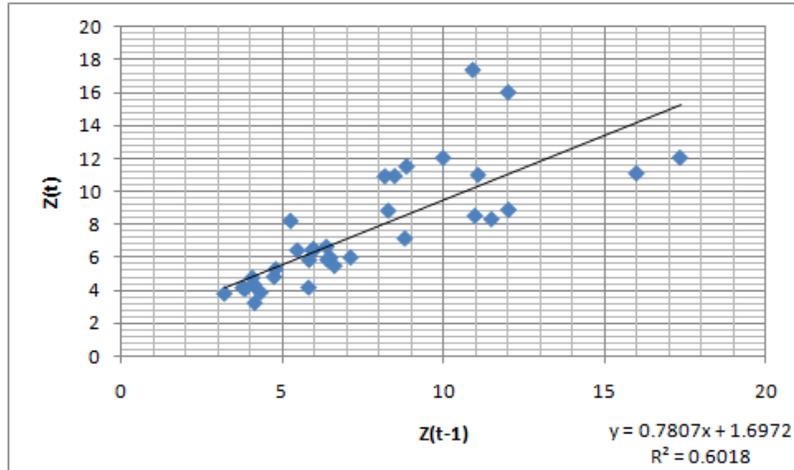


Figure E.3. Regression model for $X(t) = aX(t-1) + \epsilon(t)$

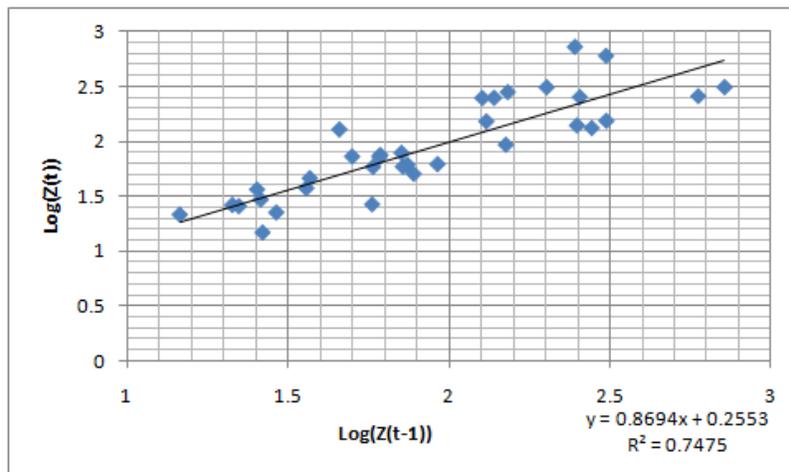


Figure E.4. Regression model for $\log(X(t)) = a \cdot \log(X(t-1)) + \epsilon(t)$

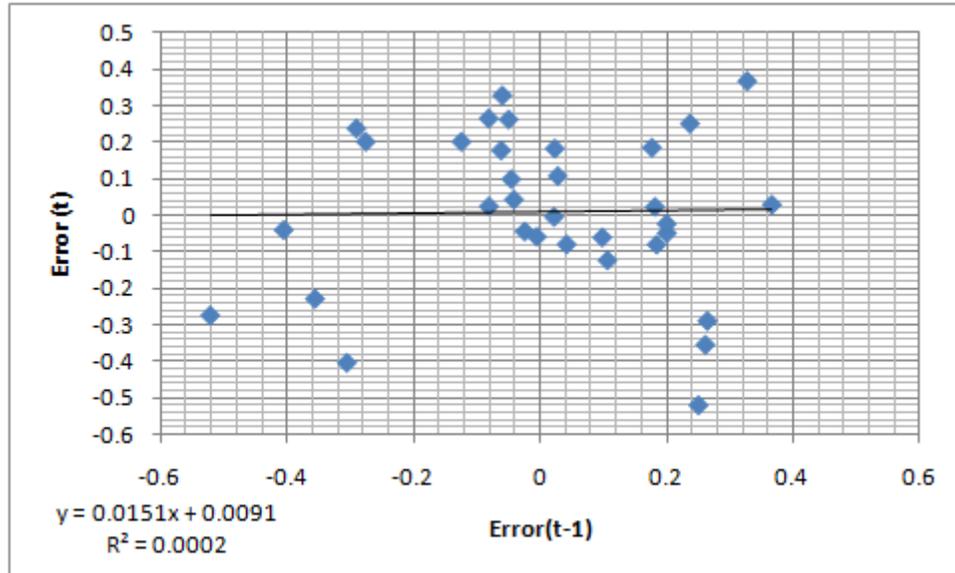


Figure E.5. Scatter plot of successive errors. The line slope is not statistically significant, the correlation, as measured by R^2 is practically zero

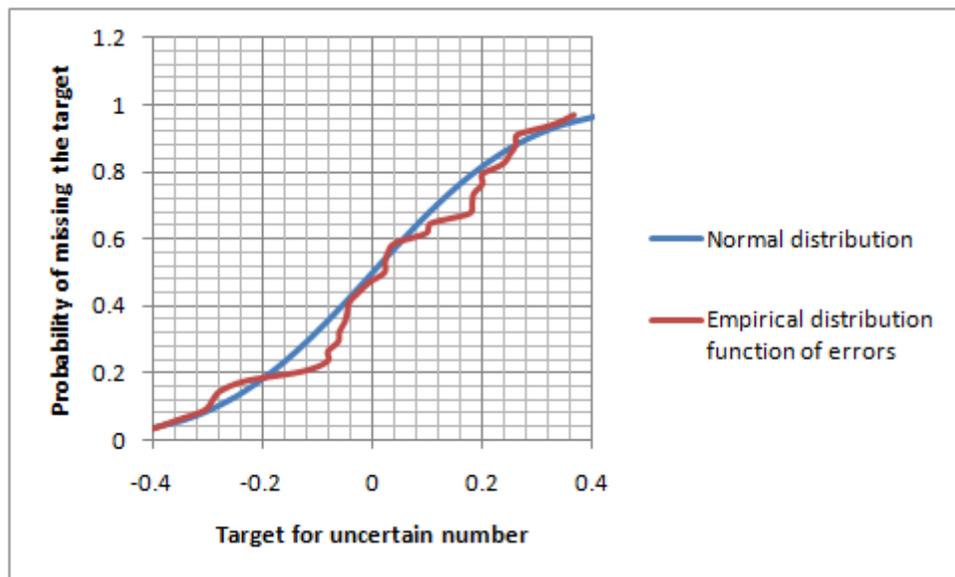


Figure E.6. Empirical cumulative distribution function of the regression errors versus normal distribution

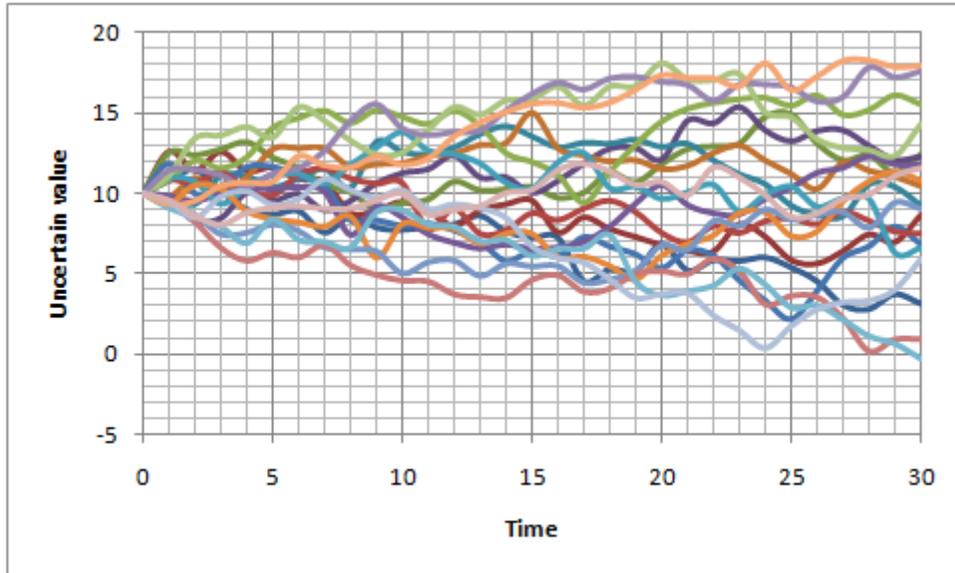


Figure E.7. 20 paths of an additive random walk $X(t) = X(t-1) + \varepsilon(t)$

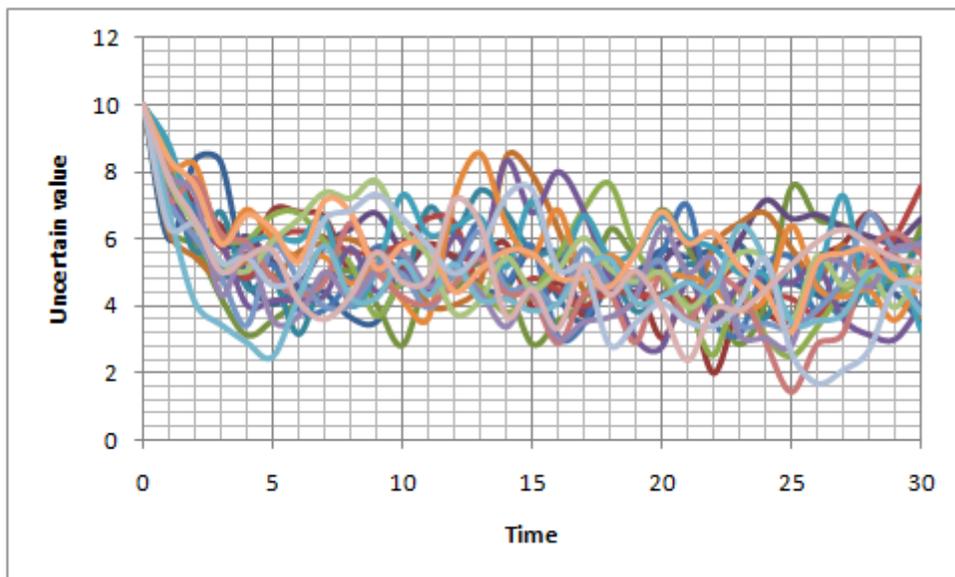


Figure E.8. 20 paths of a mean reversion process $X(t) = X(t-1) + 0.5*(5 - X(t-1)) + \varepsilon(t)$

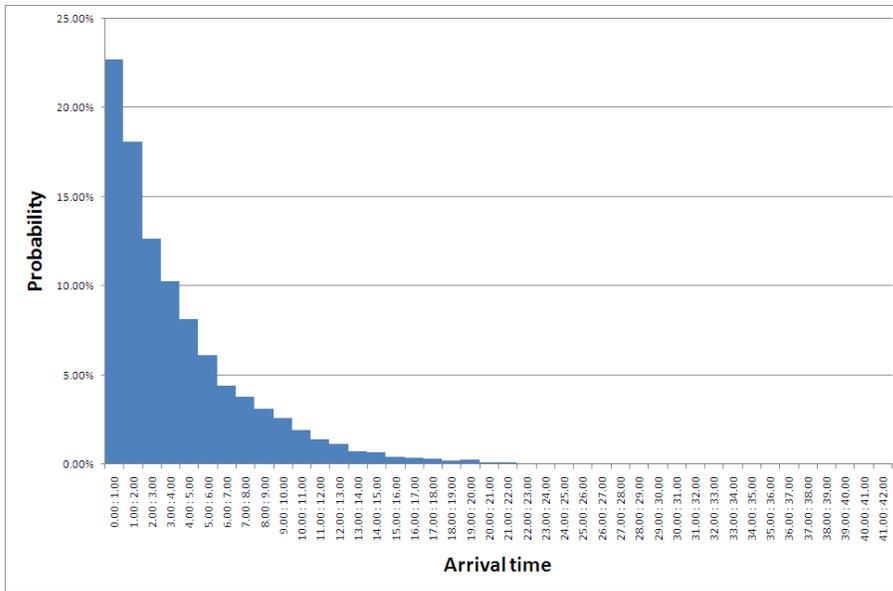


Figure E.9. Example exponential distribution

¹ See Savage (2009).

² See Shisko et al. (1995).

³ When base case assumptions are replaced by distributional assumptions, e.g. on the future price of oil, it is critically important that all projects being evaluated use the same distributions for common inputs. Technology is available to facilitate probabilistic planning across projects, see <http://www.probabilitymanagement.org> .

⁴ The engineering curricula at leading European and Asian universities commonly place little emphasis for economic or social studies. In the United States, the accreditation boards require engineering undergraduates to take a quarter of their credits in some form of liberal arts. However, the professional curriculum often does not integrate this material.

⁵ To be precise, Jensen's law is somewhat more specific. It states that $E[f(X)] \leq f(E[X])$ if $f(\cdot)$ is a convex function.

⁶ The argument against using discounted cash flow analysis when cash flows are uncertain centers on the inability either to account for optionality, such as the option to defer investments, or to vary the discount rate as risk varies as the project advances. See Trigeorgis (1996) for example. From a theoretical perspective, the discount rate should reflect the degree of risk: the greater the uncertainty, the greater the discount rate. However, greater uncertainty in the cash flows gives value to the option to wait with the investment until the uncertainty is resolved to some degree. The time value of money competes with the value of learning about the future, and the latter is not included in discounted cash flow analysis. Thus finance theory has developed what is known as "options analysis" that uses a "risk neutral analysis" that transforms the distribution of future cash flows into an alternative "risk neutral" distribution that justifies a constant discount rate, in fact a "risk free rate". Theory justifies this approach on the basis that broad markets will enable "arbitrage enforced pricing" of the assets generating the cash flows. Appendix F provides details.

Financial markets have widely implemented options analysis. The functioning of these markets most often justify the assumptions under which Merton, and Black and Scholes (1973) pioneered this approach. However, these assumptions generally do not apply to real world projects we might design. Moreover, options analysis requires special training. In any case, many practicing designers and system managers have rejected the use of real options based on finance theory. They neither understand nor trust it.

⁷ In the United States, the Office of Management and Budget sets the discount rate for agencies of the national government. Its current pronouncements appear in Circular A-94, for example US OMB (2008).

⁸ See Brealey and Myers (2002) for details.

⁹ At a deeper level, the discount rate reflects available opportunities. This is because organizations and investors will rationally select the projects with the highest returns and these will set the bar for the price of the capital invested.

¹⁰ A student version of XLSim[®] is included in Sam Savage (2004), which is an excellent tutorial to Monte Carlo Simulation.

¹¹ The tornado chart in Figure D.3 was produced with the Sensitivity Toolkit by Baker, Powell and Burnham (see <http://mba.tuck.dartmouth.edu/toolkit/index.html>)

¹² Two-way data tables are occasionally used to produce sensitivity charts that show how system performance changes when two inputs vary simultaneously. However, these 3d-graphs are somewhat difficult to interpret and are limited to a maximum of two variables.

¹³ Figure D.4 is part of a data table with 1000 trials, stored in the range T10:AB1010 within the randomized NPV spreadsheet, i.e., new inputs and an associated NPV were sampled every time F9 is hit. The data table was produced by (i) filling column T11:T1010 with consecutive numbers 1-1000, (ii) linking U11:AB11 to the cells that contain the generated inputs and outputs that we wished to track, (iii) highlighting T11:AB1010, (iv) invoking the data table command, and nominating an arbitrary unused cell, here S11, as "column input cell". Note that the data table command does not work across worksheets, so the changing inputs and outputs have to be in the same sheet as the data table.

¹⁴ The statistically trained reader will realize that the theoretical correlation between demand and cost in this example is precisely zero.

¹⁵ The open source statistics package R has a scatter plot matrix function. SAS sells a more user-friendly commercial package with many more graphic choices under the name of Jmp.

¹⁶ Graphical output for Monte Carlo simulations in this book has been produced with XLSim[®].

¹⁷ The formula " $=0.5*(RAND() + RAND())$ " generates a triangular distribution between 0 and 1. The formula " $=A1 + 0.5*(RAND() + RAND())*(B1 - A1)$ " generates a triangular distribution between a number in cell A1 and a number in cell B1 with a peak at $(B1 - A1)/2$.

¹⁸ If R is the annual revenue when independently sampled as in step 4, R' the new revenue variable, and D the demand deviation, then the formula is $R' = R + b*D$. Therefore $E[R'] = E[R + b*D] = E[R] + E[b*D] = E[R] + E[b]*E[D] = E[R]$. Here we have used the fact that $E[X*Y] = E[X]*E[Y]$ when X and Y are independent, and that the expected deviation of demand from its projection $E[D] = 0$.

¹⁹ If the shocks can be large, it may be necessary to safeguard against negative demands, for example by using a formula of the form $X(t) = \max\{0, X(t-1) + \epsilon(t)\}$. This leads to a slight violation of the expectation consistency requirement, as discussed at the end of Appendix D. However, the bias is typically small and acceptable in most practical situations.

²⁰ This is called the inverse transform method. to see why it works, recall that $\text{Prob}(\text{RAND}() \leq z) = z$ for $0 \leq z \leq 1$. Hence for every x $\text{Prob}(X \leq x) = F(x) = \text{Prob}(\text{RAND}() \leq F(x)) = \text{Prob}(F^{-1}(\text{RAND}()) \leq x)$ and therefore X and $F^{-1}(\text{RAND}())$ have the same distribution.

²¹ Exponential distributions are often used to model the time between two events. The exponential distribution with mean m has the cumulative distribution function $F(x) = 1 - \exp(-x/m)$, Hence its inverse is $F^{-1}(u) = -\ln(1-u)*m$. Note that $1-\text{RAND}()$ has the same distribution as $\text{RAND}()$ so that the use of either in the formula does not make a difference.

²² This result is known as the central limit theorem. It essentially says that the distribution of the sum of many well-behaved unrelated uncertain numbers tends towards a normal distribution, no matter what the distribution of the summands might be. See for example W. Feller (1968).

²³ Appendix E uses fixed non-random rates. In general they may well be time dependent, that is, the model may be of the form $X(t) = a(t) X(t-1) + C(t)$.

²⁴ Add [reference to forecasting literature \(see special issue International Journal of Forecasting 2006\)](#).

²⁵ The coefficients calculated by a linear regression analysis have standard errors attached to them. Here, these are 0.18 for the intercept and 0.09 for the slope coefficient. In other words, whilst the above model is the most likely model, assuming the standard assumptions for a regression analysis are satisfied, we can also sample the model itself randomly as $X(t) = X(t-1)^a \exp(b + \epsilon(t))$, where a and b are sampled from a normal distribution with mean 0.87 and 0.26, respectively, and standard deviation 0.09 and 0.18, respectively. However, if we do this, we should adjust the error term distribution.

²⁶ Microsoft Excel has functions SLOPE, INTERCEPT and STEYX to calculate the slope and intercept of the model as well as the standard deviation of the residual of a regression of one array of data onto another.

²⁷ A Kolmogorov-Smirnov test would be an appropriate way to test this statistically.

²⁸ For example [xxxxxxx some US and British standard texts -- Stefan, your favorite?](#)