Supported Intelligence

Guide to Recursive Models

With applications from the Rapid Recursive® toolbox

February 2013

Patrick L. Anderson
Executive Chairman,
Supported Intelligence, LLC

Supported Intelligence, LLC
1555 Watertower Place, Suite 300
East Lansing, Michigan 48823
Tel: (517) 333-7024

http://www.supportedintelligence.com

ISBN:0989083306

Note on Intellectual Property

The Rapid Recursive® toolbox is patent-pending. It is licensed by Supported Intelligence LLC to users under an agreement that protects the intellectual property of both the supplier of the software and the user. Licensed users have access to a wide variety of information, maintenance and support services, and contributed files.

If you are not already a licensed user, please visit the Supported Intelligence LLC website and learn about your options for licensing it.

Note on the Use and Limitations of Analytical Models

This *Guide* is intended to assist interested people in the general power, use, and features of recursive models; and to provide instructive examples using the Rapid Recursive® toolbox. It is not intended as a comprehensive report on valuation, accounting, finance or business administration; neither it is intended to convey an estimate of value for any specific business or investment. Readers are reminded that *all* models represent simplifications of the real world, and involve assumptions and forecasts about the future that cannot be expected to be correct.

Supported Intelligence LLC, as well as the editors and contributors of this *Guide*, do not possess knowledge of any specific investment or business decision contemplated by a reader, and have no financial interests in any such decisions. For these and other reasons, responsibility for any such decisions rests solely with the reader.

Table of Contents

List of Tables, Figures, and Equations

Part I. Introduction to Sequential Decision Problems

Chapter 1. Introduction

PURPOSE OF THIS GUIDE

The purpose of this *Guide* is threefold:

1. To introduce the concepts of sequential decision problems, and the recursive approach to such problems;
2. To describe a step-by-step process to organize the information necessary to compose a sequential decision problem in a form that can be solved by the Rapid Recursive® toolbox, error-check the model, solve the model, and report the results; and
3. To present a set of such models that represent common decision problems, and which have already been composed into models that can be immediately solved by the Rapid Recursive® toolbox. Some of these are *Solution Templates* that are included as part of the software licensed to toolbox users.

In addition, this *Guide* can assist readers to:

4. Identify a set of practical issues—such as the selection of discount factors and state variables—for which there is currently no standard convention.
5. Provide the reader a set of references on the theory of value functional models, real options, sequential decision problems, related valuation concepts, and the underlying mathematics.

LIMITATIONS OF THIS GUIDE

This *Guide* is a valuable resource. However, it is not intended to be provide basic instruction in the use of the toolbox, complete historical discussion of the underlying theory, or an accurate assessment of any specific practical decision problem. We suggest the following resources for these purposes:

- The *User's Guide* is a program reference for the Rapid Recursive toolbox, and should be consulted for information on the proper command syntax and usage, as well as troubleshooting, licensing, and installation.
- The Rapid Recursive toolbox, version 1, requires MATLAB®, a product of The Mathworks. User of this version of the toolbox can consult many resources that are available for learning and using MATLAB®. See Appendix B. Bibliography and References for a brief list of such resources, or consult the Mathworks website.
- The published works cited here include derivations of the theory behind sequential decision problems, recursive models, and related concepts such as control theory, MDPs, dynamic programming, finance, economics, and valuation. See "Appendix B. Bibliography and References" on page B-1.

Also, please remember that professional judgment, knowledge, and assessment of the risks and opportunities involved in any important decision should be paramount in any practical financial decision. No software, no matter how advanced or powerful, can replace such judgment.

USER GUIDE AND HELP DOCUMENTS

In addition to this *Guide,* licensed users have access to a *User's Guide* that serves as a reference for commands in the Rapid Recursive toolbox software, as well as for installation, licensing, and supported operating systems.

In addition to these resources available to all licensed users within the software or accompanying it, additional assistance, as well as other resources, is available online at the Supported Intelligence LLC website.

Finally, expert assistance is available to licensed users on an incident-by-incident basis. Please see the *User's Guide* or your license terms for more information.

SPECIAL ASSISTANCE TO SCHOLARS AND AUTHORS

We encourage scholars, researchers, and authors to make use of the Rapid Recursive® toolbox in their work. In some cases, Supported Intelligence LLC can provide special assistance to authors that are performing research that they intend to publish (such as in a book or journal). This assistance includes suggestions for data and other sources, as well as improving Rapid Recursive models.

If you are an author with a specific research project with a goal of publishing the results, you can contact us regarding this special assistance as described in Appendix A. Contacting Supported Intelligence LLC. Please include in your inquiry information about the topic and time line of your research, and the expected form of publication.

ACCREDITED CONSULTANTS; CUSTOMIZED MODELS

Supported Intelligence LLC provides training to individuals and organizations. Those completing this training may receive accreditation. Accredited consultants that regularly provide consulting services to their clients using the Rapid Recursive toolbox have access to a licensing agreement with SI LLC that provides enhanced services and a fee schedule that matches the use of the product with the licensing fee.

For organizations that wish to create, use, and periodically update customized models, these accredited consultants may be able to provide assistance on a consulting basis. Please consult the listing of accredited consultants on the SI LLC website, or contact us as described in Appendix A. Contacting Supported Intelligence LLC.

Please note that accredited consultants are independent businesses. Organizations wishing to retain their services should formalize their arrangement in an agreement that covers such topics as confidentiality of any proprietary data, quality assurance, rights to use models developed during the course of the consulting assignment, quality assurance, work plan, and fees.

RESPECTING INTELLECTUAL PROPERTY

Bringing the power of "recursive thinking" to our customers required many years of work, and major financial investments by our founders and investors. We ask that you respect this investment and the effort we have placed into developing this product by limiting the distribution of our software to licensed users.

We also respect the work that you do, and the importance you place on being able to share the results of your analyses with your co-workers, colleagues, and customers. For that reason, we acknowledge your rights under our licensing agreement, and specifically allow our licensed users to use, adapt, and report the results of Rapid Recursive® toolbox models in a manner consistent with that agreement.

While we ask that you cite the Rapid Recursive toolbox (and any solution template you adapted) in any report or document distributed to those who are not licensed users.[1] You may also excerpt short passages from this document in your work, provided you properly quote the passage and cite it by title, author, and date. However, any proprietary data you use in such an analysis, and the resulting results and report, remain your property until you distribute it to others.

OUTLINE OF THIS GUIDE

Before setting up a model in the Rapid Recursive toolbox, it is important to understand the *recursive* approach to *sequential decision problems.* Thus, this *Guide* is laid out as follows:

1. Part I introduces the software; and then discusses sequential decision problems. These are discussed in "Sequential Decision Problems" on page 6.
   It provide examples of problems in a variety of fields that can be attacked with recursive methods. We then discuss a set of possible decision rules for solving such problems, only a few of which are commonly mentioned in financial curriculum. These are discussed in "Approaches to Decision Problems" on page 10.
2. Part II provides guidance on a variety of topics important to those actually using recursive models. It includes a discussion of the serious problems with using traditional discounted cash flow ("DCF") analysis to evaluate decisions under uncertainty in "Problems with the DCF Model: Ignoring Management Flexibility" on page 17.
   Practical steps for composing a recursive model and identifying the information to use in that composition are addressed in "Using the Rapid Recursive® Toolbox" on page 24. A focused discussion of discounting, discount

1. An acceptable citation would be something like the following:

   We used the Rapid Recursive® toolbox, a software program from Supported Intelligence LLC, to compose and solve the decision problems described in this research project. The web site for Supported Intelligence LLC is: http://www.supportedintelligence.com.

and growth rates, and certain mathematical conventions is contained in the following chapters, including “Discounting and Discount Rates” on page 27.

3. Part III introduces a set of selected Solution Templates that are provide to licensed users of the Rapid Recursive toolbox. These cover various decision problems in valuation, management, finance, risk analysis, and economic policy. Excerpts from these templates as well as output from running them are excerpted here. Licensed users have access to thousands of lines of unencrypted code in these Solution Templates, that they can adapt and use professionally and in their academic research in a manner consistent with the License Agreement.

   See “Selected Solution Templates” on page 52.

4. A set of appendices, including:

- Appendix A. Contacting Supported Intelligence LLC
- Appendix B. Bibliography and References
- Appendix C. Sample Worksheets

Let us know how you use the Rapid Recursive Toolbox!

We hope you will find this *Guide* useful.

If you have suggestions for inclusion in future versions, or think you have found a mistake, mis-statement, or paragraph that needs clarification, please contact by mail or email using the addresses listed in the appendix.

Also, if you find the toolbox particularly useful, or use it for something particularly interesting, or have some of your work published, please let us know. I am sure we will be pleased, and may be able to share any published work with others.

Chapter 2. Sequential Decision Problems

SEQUENTIAL DECISION PROBLEMS

Sequential decision problems ("SDPs") are problems of the following type:

- There is a time *sequence* of events. Information on market conditions, the amount of resources the subject controls, and other relevant facts arrive as events unfold in this time sequence.
- The subject must make a *decision* at one or more points in that time sequence. These decisions are represented by actions of the subject. These actions could affect the current reward or cost, and also market and other conditions that will affect the subject in the future.
- At each point in time, the subject receives some type of reward (which might be a payoff, benefit, earnings of some type) or incurs some type of cost.
- The subject faces the *problem* of choosing the best decision, according to some decision rule, based on the information available to him or her at that time.

Such problems are also called "sequential decision models" and may be defined more formally in set notation.[1] The concept of a sequential decision problem was formalized by the statistician Abraham Wald in the 1940s.[2] Wald's statistical approach to decision problems spawned many areas of fruitful research in fields such as economics, statistics, natural sciences, finance, control theory, and operations research.

---

1. Puterman (1994, page 1) describes the model as follows:

   The key ingredients of this sequential decision model are the following:

   A set of decision epochs

   A set of system states.

   A set of available actions

   A set of state and action dependent immediate rewards or costs.

   A set of state and action dependent transition probabilities.

   The sequential decision problem is to choose, prior to the first decision epoch, a policy to maximize a function of this reward sequence.

2. In particular, Wald (1947) and Wald (1950) formalized sequential decision problems and described approaches to solving them. Many of the subject of Wald's analyses were statistical decision problems, and these are sometimes called *Statistical Decision Problems* in the literature.

   On the intellectual history of decision theory, see Brown (2000).

EXAMPLES OF SDPS

Many of the most difficult, important, and common decision problems faced by people around the world can be structured as sequential decision problems. For example, all of the following are sequential decision problems:

Investment problems

- You have a chance to invest in a company developing a promising technology, which may or may not result in significant future earnings. Alternatively, you can invest in very safe short term bonds that pay a predictable, but very low, return. Which do you do?

Entrepreneurial Problems

- You are one of the founders of a new technology firm. You and your partners own 100% of the company, which is burning through its capital at an alarming rate. A “angel” investor offers to purchase 50% of the stock in return for an infusion of funds, but also wants a liquidation preference. Do you take the badly-needed funds in return for a loss of control, and a much lower potential payout in the future?

Managerial Decisions

- You are managing a division of a very large manufacturer of consumer goods in a competitive industry. You know that significant investments in new product is necessary to maintain your current market share, but that size of investment would cost more than your entire current profits. Without the investment, however, your market share will decline. Should you make the investment?

Real Options Problems

- You are managing a natural resource company (a logging, mining, or similar enterprise) that has the rights to continue to exploit the raw materials in a certain area. However, the size of those reserves are steadily declining. You can sell the land (and lose the rights) for a predictable, and stable, price anytime during the next five years. When, if ever, do you sell?

Life Decisions

- You have saved some money that you can spend on a house (which improves your life and makes it easier to study and work) or continuing to attend graduate or technical school (which increases the wages you expect to earn in the future). Which should you do?

You can observe that all of the decision problems listed above have a time sequence, where information about relevant conditions are revealed over time, and where the manager, owner, or other subject person is capable of taking actions that affect the future course of events and future value of the business or other asset.

EXAMPLES OF OTHER DECISION PROBLEMS

Sequential decision problems are qualitatively different from many other decisions. In particular, they are different from what are known as "one-off" decisions, in which a person makes a one-time decision with no possibility of changing course in the future.[1] Throwing a coin into a fountain is one example (assuming, of course, that you are not going to climb into the fountain to retrieve it!) This is a single event: you throw the coin into the fountain, and it is gone. The sequence stops there.

Sequential decision problems are also different from whimsical actions that result in no consequences. An example is choosing between chocolate and vanilla ice cream on a spur-of-the-moment visit to an ice cream parlor, once you have already arrived at the parlor and decided to enjoy some ice cream. This is a decision, of course, but it is difficult to call it a decision *problem.*

One can conceptually broaden these examples to arrive at a sequential decision problem (for example, by visiting the fountain or ice cream parlor every day, and having financial or physical health depend on the decisions made there). However, in this book we will consider only sequential decision problems with non-trivial elements including a time sequence, set of rewards, and decision opportunities that carry with them consequences for the future.

ADVICE ON RECOGNIZING AND STRUCTURING SEQUENTIAL DECISION PROBLEMS

The examples above provide the basis for some important practical advice about using analytical tools in general, and decision analytical tools in particular.

1. Use an analytical tool only when it makes sense to use it.

   The corollary is also sound advice: if a person doesn't think they have a problem, face a decision, or have the ability to affect future conditions that are important to him; then it probably doesn't make sense to model a sequential decision problem for that person.

   It may be possible for another person, such as a skilled analyst, to identify the necessary components for a formal decision model within the person's situation. However, one should be very cautious about assuming that other people have the same motivations, information, and objectives as you have.

2. Be sure that the conditions, rewards, and costs you put in your model are recognized by the subject of the problem

   Good models capture the major incentives and beliefs of the subject of those models. Therefore, be sure that the rewards, costs, probabilities, and other mathematical elements of your model comport in broad terms with the state of the world *as seen by the subject of the problem*. If it does not, you may have a great analytical model—but you can't expect it will explain the behavior of the original subject!

---

1. The term "one-off" appears to originate in the artificial intelligence literature. See, e.g., Poole & Mackworth (2010, section 9.2).

If the purpose of the model is to analyze a decision problem, then the model and its inputs should reflect the perspective of the subject, not the analyst.[1]

3. Ensure your model captures the major motivation and opportunities of the subject of the decision problem

   Recursive models allow the analyst to capture much more complexity in incentives, information, possible actions, and motivations than traditional models. However, it is possible in any model to ignore the true motivations, information, and possible actions of the subject of a problem and substitute those described in a textbook or an academic exercise.

   Don't fall prey to this error. Useful models reflect practical conditions and constraints faced by actual people. Reality is often somewhat messy. While you can never capture everything in a model, strive to capture what the subject feels is important.

1. This is another illustration of the "first maxim of business economics:" *Remember, it is a social science.*

   See Anderson (2004).

Chapter 3. Approaches to Decision Problems

POSSIBLE APPROACHES

People with training in accounting, finance, business management or economics have been introduced to at least one approach to solving decision problems that involve time, risk, and money. In fact, it may appear from reading the literature in these fields—as well as popular investment advice columns and websites—that there is one universally accepted decision rule in finance.

However, before discussing that common approach, we should note that there are many possible methods to address decision problems where risk and uncertainty are present. These include random selection, appeals to moral authority, listening to experts, “follow the crowd” and “don’t follow the crowd” advice; as well as selecting choices on the basis of avoidance of risk, seeking of risk, diversification of risk, highest possible outcome, best typical outcome, and least bad worst outcome.

We will introduce a novel approach that can be used to solve many sequential decision problems. However, first we consider the most common approach.

FOUR SPECIFIC DECISION RULES

In this *Guide,* we will focus on sequential decision problems that involve important financial and economic considerations, where we assume that the courses of action available are consistent with the laws of society and the philosophy of the subject person. For these, it is worth discussing the following possible approaches to financial decisions, which have been proven to be powerful enough to be applied in a wide variety of situations. Each is very briefly described below:

1. The DCF method, with the net present value rule

   This method requires an analyst to identify on scenario for the future, usually involving forecasting future cash revenue and expenses. From this forecast, the analyst discounts the future earnings and expenses back to the current day, using a discount factor that (at least in principle) incorporates the risk in the investment.

   DCF analysis is probably the most widely-used tool for analyzing financial decisions, as well as many managerial, entrepreneurial, and life decisions. Countless loan applications, business plans, valuation schedules, and cost-benefit analyses include variations on the basic discounted cash flow model, and (implicitly or explicitly) rely upon the net present value decision rule. It is discussed further below.

2. Monte Carlo analysis

   One response to the defects of the standard DCF model when faced with real options, as well as to the possibility of responses to random events that are likely to occur in the future, has been the development of Monte Carlo methods. Rather than project one scenario and evaluate it, Monte Carlo analysis begins with one scenario and then evaluates a large number of randomly-generated paths that are based on the one scenario.

These are not truly independent scenarios, but instead are random "trials" that are explicitly based on a single scenario that allows for a different random outcome in each trial. The decision rules that are followed after Monte Carlo analysis can be based on selecting the best likely outcome, the best outcome on average, the best worst case.

3. Financial option valuation methods

   Since the 1970s, an analytical method for pricing certain pure financial options in complete markets has been widely available. As in the classic Black-Scholes-Merton formula, option pricing methods provide a price that is determined by an "intrinsic value" based on net present value, plus an "option premium" based on the likelihood that a randomly varying price will end up "in the money." These methods can be used only for traded securities where option contracts are widely available, prices are quoted on a regular basis, investors can borrow and lend easily at known rates, and counter-party risk can be largely ignored. However, in these cases, option pricing methods are very powerful, and avoid the failings of DCF models in valuing these assets.

   Option valuation methods imply a decision rule to exploit any arbitrage opportunities, meaning selling or buying whenever market prices deviate substantially from the analytical price.

4. The market method.

   The basis of the market method is the principle that similar traded assets should trade in the same market at a similar price. This method is only available where a specific security or asset is similar to others that are traded on a market sufficiently frequently to establish a buying and selling price. When data on comparable asset sales are available, and the asset is actively traded, this is a powerful tool.

   As with option pricing methods, the market valuation method implies the exploiting of arbitrage possibilities.

5. The "value functional" or "recursive" method

   In the 1950s, a theoretical approach to evaluating sequential decision problems was first outlined. This approach involves the specification of a *functional equation* that incorporates both state and action variables, as well as a time index and transition probabilities. The functional equation often has a *recursive* structure, meaning that a multi-period problem is formulated as a series of two-period problems that form a mathematical recursion.

   This method requires an analyst to identify and forecast future cash earnings or other benefits for a set of possible actions and market conditions, as well as a discount factor and a transition function or set of transition probabilities that specify how market conditions change over time.

   The value functional approach is also distinguished by its decision rule: the maximization of *value,* rather than current-period earnings or the highest possible discounted earnings from any one scenario.

As stated in these very brief descriptions, certain types of investments, securities, characteristics of markets, and institutions are required for some approaches. Furthermore, all of these approaches place demands on the analyst to estimate, forecast, identify, and calculate various factors.

THE NET PRESENT VALUE RULE AND DCF

The most widely-cited decision rule in economics, finance, and business management is known as the "net present value rule" in finance and business, or (in economics) the "neoclassical investment rule." That rule is to maximize the expected net present value of any investment opportunity. Versions of this "rule" appear in almost every standard corporate finance textbook, an entire stream of literature in economics, as well as in finance literature dating back at least to the Modigliani-Miller articles of the late 1950s. [1]

Following that rule can appear to be deceptively easy: simply project out the future cash flows for any investment (using your best estimate given the information at the time), and then discount it back to the present date. If the net present value of those expected cash flows exceeds the cost of the investment, then undertake the investment. If not, choose another use for your funds. The practice of discounting cash flows in a manner consistent with this decision rule is so common that it has its own acronym: DCF.

ASSUMPTIONS IN THESE METHODS

While there are many wrinkles and complications involved in each of these tasks involved in all the methods listed above. Furthermore, these and all analytical decision models involve some assumptions. The following critical assumptions are worth noting:

1. The DCF method assumes that the investor has the capital necessary to make the investment, and has a known discount rate.

   The former is often true, although the latter is often a suspect assumption. To some extent, these assumptions are shared across many methods, although vary in importance among them.
2. Furthermore, DCF relies on the assumption that analyzing one scenario (or, perhaps, a handful of scenarios, each with its own cash flow schedule) is adequate as a basis for evaluating the decision. As is apparent by testing the DCF model with any investment involving real or financial options, this is often not true.

---

1. See Anderson (2013) for citations on the ubiquity of the NPV rule in business and in the academic literature, and the intellectual origins of the "rule."

   Anderson (2013) also cites studies showing the failure of business investors to actually follow it. See Schwartz & Trigeorgis (2001) for an assembly of classical papers on the failure of the NPV rule in the face of "real options."

   Complete citations to these works are in "Appendix B. Bibliography and References" on page B-1.

The failure of the DCF method to consider management flexibility—the ability to change course in the future—is probably the most damaging to its use in critical situations. We discuss this further in “Guidance for the Use of Recursive Models” on page 16.

3. The DCF method also relies on the assumption that the investor is just as happy with receiving income 2 and 3 years from now as tomorrow, as long as the amount received is bigger (by an amount represented by his or her discount rate) in the future.

   This is known as “time separable utility.” However, very few investors (or business managers) are actually indifferent to the timing of future expenses or earnings! (See the discussion in “Discounting and Discount Rates” on page 27, and in particular “Two Critical Assumptions: Utility, and Time Preference” on page 28.)

4. Similarly, the market method (and the option pricing method) relies on the key assumption that there are enough trades of similar assets to establish a market price for the subject investment.

   This is often true for commodities such as oil and gold, and close enough to being true to make the method very useful for pricing many used cars and trucks, houses, and commercial real estate, as well as some kinds of collectible goods and certain well-identified businesses. It is not true for most other types of assets, notably including most entrepreneurial and closely-held businesses.

5. The Monte Carlo method assumes that a relatively simple structure of decision nodes and rules can be identified in advance, and that the underlying risk distributions can be well-defined by a standard statistical formula. Under these conditions, Monte Carlo by itself provides a richer analysis of one scenario; Monte Carlo combined with a decision tree (or “event tree”) provides a richer analysis of a handful of scenarios.

   The assumption about well-behaved risks, which can be modeled using standard statistical distributions, might be adequate in certain situations. However, it is clearly not true for most intellectual property, almost all start-up businesses, and many closely-held firms, not to mention investments in new technology, many natural resource leases and rights to segments of the communication spectrum, and companies with significant real options to grow or facing substantial risks that may imperil their future viability. Furthermore, the limitations of an event tree become apparent as soon as you consider more than a few possible states and actions that could be undertaken, or more than a few time periods.

Again, no analytical method avoids relying on at least some assumptions. However, some methods require assumptions that are often heroic and frequently are clearly incorrect. Others rely upon much milder assumptions with smaller penalties involved for minor violations of assumptions. These are worth noting when selecting a model.

MOTIVATION FOR THE RECURSIVE METHOD

One of the methods listed above is at least as powerful as the others, and arguably relies on fewer unrealistic assumptions. It is the value functional or recursive method, and it is the focus of the majority of this *Guide.*

The major reasons why a recursive method is worth pursuing are:

1. As noted above, it does not require certain assumptions that are infrequently true in actual practice. For example, it does not require an assumption of time-separable utility, nor of low transaction costs, perfect information about the market today, or that the subject is widely-traded securities.
2. It does not require the assumption of symmetric distributions of possible outcomes, such as the normal distribution, or even a "bell curve" shape of possible outcomes. Indeed, it can explicitly model asymmetric risks, including the risk of an unusual and potentially catastrophic "black swan" event.
3. It can recognize explicitly the existence of managerial choices, including real options, and natively incorporate them in the decision model.
4. It can incorporate information in a formal, objective manner that is not considered in other models. For example, an analyst may be aware that a real estate parcel could sold for a known price, but cannot include that information in a DCF analysis unless the analyst assumes it is sold.

DIFFICULTIES IN USING A RECURSIVE MODEL

While the recursive model is powerful and has significant advantages over other methods, it also suffers from some disadvantages that have formed significant barriers to entry. In particular:

1. It is novel. Until very recently, the entire concept was obscure except in a fairly narrow set of research and academic settings.
2. Because it is novel, there are far fewer guidelines to usage and examples, not to mention well-established teachings, than exist for all the other methods.
3. The underlying logical paradigm is different from that of the standard DCF model, and involves concepts that are often not highlighted in other approaches. Usage of a recursive model therefore requires an analyst to learn a bit of a new "recursive language" as well as understand a new model.
4. Even though there are typically fewer assumptions about markets, assets, and investors required by for a recursive model, one still must make some assumptions.
5. Until very recently, no commercial software existed that allowed interested analysts to create recursive models and solve them in an accessible manner.

Fortunately, much progress has been made over the past decade on the first three barriers. With the public release of the Rapid Recursive® Toolbox by Supported Intelligence LLC in December 2012, the latter burden was finally overcome. This *Guide* is intended to describe how the Rapid Recursive® software can be used in practical situations, and to briefly explain the underlying logical paradigm of a recursive model and introduce some of the elements of the recursive language. It addresses the remaining barrier, although does not completely overcome it.

In addition to the advice in the main part of this *Guide,* the references included in an Appendix provide extensive theoretical and mathematical background on this and related topics.

Part II. Guidance for the Use of Recursive Models

Chapter 4. Problems with the DCF Model: Ignoring Management Flexibility

In addition to the inherent problems of forecasting future events, the DCF methodology has a serious flaw in cases where the investor or manager has the power to change course in the future.

This is most easily seen with real or financial options. An "option" is an opportunity, but not the obligation, to do something. A financial option (such as a call option) is the right, but not the obligation, to purchase or sell a traded security for a certain price during a certain time period. Such financial options typically have some value, even when they are "out of the money," because they provide insurance against other events, act as a hedge against other risks, or provide the chance to earn money (or avoid losses) under certain circumstances.

REAL OPTIONS

"Real options" are opportunities to buy, sell, wait, reinvest, scrap, or take other actions involving real assets. The term "real options" was coined by Stewart Myers (1977), and the existence and value of real options in business situations has been extensively studied by Kester (1984), Dixit & Pindyck (1994), and in the papers compiled in Schwartz & Trigeorgis (2001). (See complete references in the Appendix.)

Both real options and financial options derive some of their value from the *possibility* that events will produce a valuable outcome for the investor in the future, even if such events do not occur. Because the discounted cash flow methodology relies on projecting the most likely scenario, it inherently ignores other scenarios. If other scenarios are very similar to the one forecasted, there may be little loss in ignoring them. However, if they are markedly different—as is often the case with a real option to buy or sell—the single-scenario DCF model can provide very bad advice.

FAILURE OF DCF IN CAPITAL BUDGETING

DCF methods are the standard approach to evaluating investment opportunities in large businesses. (See the survey in Anderson [2013]). However, the failure of the net present value criteria in the "capital budgeting" decision of companies facing real options or asymmetric risks has been apparent for some time. (In portfolio analysis, the failure of a pure NPV rule was notes at least as far back as Harry Markowitz [1952], who proposed adjusting the expected NPV for risk, creating the basis for the mean-variance analysis.)

The failure of DCF was noted in the economics literature at least as far back as the 1970s (including by Myers [1977]). By the 1990s, Dixit & Pindyck (1994) concluded that the DCF model and NPV rule were inherently flawed:

> Hence the simple NPV rule is not just wrong, it is often very wrong.

By the 2000s, the causes of the failure of the NPV rule had been well-identified:

- DCF requires the creation of a single-decision scenario, which plays itself out over multiple time periods, and for which the net present value of the benefits is the key decision criteria. A single discount rate is used to capture all elements of uncertainty and time preference.
- However, both real and financial options imply a multiple-decision scenario. Management has flexibility that is not captured in a DCF model.
- Financial option pricing formulas, even in simple cases, imply value that is not the same as implied by DCF or the NPV rule.[1]
- Risks are asymmetrical for business operations when real or financial options are present. Typical DCF analyses are based on an assumption of symmetric risks. (This assumption is often based on a presumption that their is a distribution of returns that is similar to that of a bell-shaped normal distribution.)
- For entrepreneurial firms, the standard DCF models are often completely incapable of providing a good valuation assessment.[2]

Thus, standard DCF analysis will often produce, as Dixit & Pindyck put it in 1994, “very wrong” answers.

PREVALENCE OF “ADJUSTMENTS” AND SUBJECTIVE METHODS

Of course, the workhorse DCF model is used by sophisticated investors and analysts. They, presumably, are aware of asymmetric risks and at least some of the real options that are present in business settings. How do they deal with the problems with DCF? Among the common strategies are the following:

- Use DCF as a primary method only when the embedded assumptions are approximately fulfilled, at least for major risks.
- Create an “expanded net present value” (or XNPV) measure, which is the sum of the expected NPV of current operations, plus the value of the embedded real options.
- Create a DCF schedule on the basis of objective assumptions. Then, “adjust” the assumptions until the result matches a subjective assessment of value, or the value implied by another method. By altering a growth rate and discount rate assumptions, analysts can frequently produce vastly different net present value calculations on multi-year DCF schedules.
- Ignore the flaws and plow ahead anyway.
- Create several sets of assumptions, and do a DCF analysis on each. Either compare them using a different decision tool, or select the one that most closely matches a subjective judgement about the value.

1. See, for example, an option pricing reference such as Hull (2008). The companion critique involving real options is amply covered in Schwartz & Trigeorgis (2001).
2. See Smith, Smith, & Bliss (2011), as well as Anderson (2013).

- Make decisions on "gut instincts" or other reasons, and then instruct an analyst to create a DCF schedule that supports the decision.

In some cases, the analyst is forced to use one of these methods due to lack of resources or data. Hopefully, the analyst discloses the limitations and major assumptions underlying his or her estimate in such cases. Unfortunately, however, it is all too common to see naive (and even reckless) use of DCF analyses, as well as undisclosed use of significant subjective adjustments.

ALTERNATIVES TO DCF

In the past, there were often no alternatives to DCF methods in practical work. Managers, regulators, and investors concerned about risk were often forced to use blunt instruments to evaluation asymmetric risks, use Monte Carlo simulation to at least describe the possible outcomes, or ignore such risks entirely.

However, with the discussion of recursive methods in the literature in economics, and the introduction of commercial software designed to make use of these methods practical, they should now be considered by analysts, managers and investors that face problems involving real options or asymmetric risks.

Chapter 5. The Recursive Model

RECURSIVE THINKING

The essence of a recursive model is to separate a multi-period decision into a series of two-period decisions. The theory behind this originated with Richard Bellman's 1957 book *Dynamic Programming.*

Bellman argued that many multi-period decision problems could be separated into sets of the following problems: maximize, across all available actions, the sum of these two quantities:

1. The current-period earnings or other reward; and
2. The discounted value expected in the next period, given the decision made in the current period.

This gives rise to a *value functional equation,* where the value in the current state of affairs is the maximization of the sum described above. This equation is, in mathematical terms, a *functional,* rather than a simple function. A functional is a "function of functions," which was described by Dreyfus (1963) as a function that takes curves as its arguments.

The Bellman Equation

The value functional equation is sometimes called a Bellman equation. Because such an equation leads to a recursive set of calculations, the method of creating a value functional equation and solving it is often called a "recursive method."

In various forms, such methods may also be called dynamic programming, stochastic control, impulse control, the value functional method, or a Markov Decision Process.

The concepts and variables in a value functional equation, and then the value functional ("Bellman") equation itself, are shown below.[3] These equations are general enough to apply to both continuous and discrete cases, although the common terminology differs slightly. The following table provides a quick comparison between the two forms.

3. There is a slight difference in notation between these two involving the time index. In one, time is shown as an argument; in the other, it is an index on the state variable. Both are possible in value functional equations. In cases with explicit, time-sensitive real options, time is an argument to the value function.

TABLE 1. **Functions in Discrete-Time, Discrete Variable Case and Continuous Variable Cases**

	Continuous Variables	**Discrete-Time, Discrete Variables Case**	**Notes on Usage in Rapid Recursive toolbox**
Reward	reward function *f(s,x)*	reward matrix *R*	R matrix is size *SxA,* with each row a state and each column an action
Transition	transition function *g(s,x,e)*	probability transition matrix *P*	P matrix is size *SxSxA,* with each frame of the matrix corresponding to a certain action. Because these are probabilities, each row of the matrix must sum to 1

Functional Equation Variables (EQ 1)

$$
\begin{aligned}
V(s,t) &= \text{ value at time } t \text{ given state } s \\
f(s,x) &\equiv \text{ reward function given state and action} \\
g(s,x) &\equiv \text{ transition function} \\
t &= 0,\ldots,T \text{ time index} \\
s_t &= s_0,\ldots,s_T \text{ state variables} \\
x_i &= x_0,\ldots,x_M \text{ action or control variables} \\
\beta &= \text{ discount factor} \\
\varepsilon &= \text{ random error term}
\end{aligned}
$$

Value Functional Equation (EQ 2)

$$
\begin{aligned}
&V(s_t) = \max_{x\in\Gamma}\left\{ f(s,x) + \beta E\left[V(s_{t+1})\right]\right\}; \\
&where: \\
&s = \text{state}; \\
&x = \text{action}; \Gamma = \text{feasible set of actions.}
\end{aligned}
$$

The Principle of Optimality

The solution to a value functional equation is based on the "principal of optimality" first outlined by Richard Bellman.[4] This can be rephrased as "choose the best course going forward, regardless of what decisions you made in the past, in a manner that maximizes the value to you."

ESSENTIAL ELEMENTS OF A RECURSIVE MODEL OF BUSINESS

Anderson (2013) describes a value functional model of the firm that contains the following elements:

1. Identification of a relevant subject.

   The subject of a recursive model must be an entity (person, company, country, organization) that has certain characteristics. Anderson defines a firm using a three-part definition requiring a separate identity; a motivation to earn some type of reward for its stakeholders; and an operating structure.[5]

2. Identification of the state variables.

   The state variables establish the basis for the subject company's operation.

3. Identification of control variables.

   Some element must be under the control of the subject person (such as a company's manager or owner). That "control" variable must affect the current reward, or the likely change in the state variables.

4. Identification of a transition equation for the state variables.

   The transition equation relates the current state and control variables (and any random elements) to the state variable in the next period.

   Note that the combination of control and state variables, and a transition equation, makes the model recursive.

5. A value functional equation, which is usually cast in a "recursive" form.

   The relevant equation (the "value functional" equation) equates the value of the firm under the current state as the maximization (across the possible actions of the firm's manager and owner) of the following sum: the current period's rewards, such as distributed profits; and the discounted next-period's expected value.

Non Business Models

The tasks involved in creating a recursive model in general are well motivated by those identified above for the value of an operating business. However, the subject of the decision problem need not be an operating business, a business asset, or even an entity for which monetary rewards are the standard metric for success.

---

4. "Principle of Optimality: An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision." [Bellman, 1957, chapter III]

5. Anderson (2013) includes a 3-part definition of a firm, which is an organization with:

   1. a separate legal identity;

   2. a motivation to earn profits for its owners; and

   3. a set of replicable business processes.

   He notes that some legally-organized companies (such as passive investment vehicles, and sole proprietorships that are economically equivalent to wage earners) do not meet this test.

In many cases, personal and organizational decisions can be modeled in a recursive manner.

Questions to Answer

The following clarifying questions are helpful in setting up a decision problem, especially if the decision problem is cast in a recursive fashion:

- *What person (or what organization) is the subject of the decision problem?*
- *What is the reward (or penalty) that the subject person receives, dependent on the state of affairs and decisions that are made?*
- How does the subject person discount future rewards (and costs) compared to current rewards and costs?
- In what time periods does the subject typically recognize the state of affairs, receive rewards, and take actions?
- What actions can the subject person take that affects the current reward, or the future state of affairs?

Worksheets that include these questions are included in “Appendix C. Sample Worksheets” on page C-1.

Chapter 6. Using the Rapid Recursive® Toolbox

The RR tool allows a person to:

a. Compose a decision problem (or “model”), using information gained from the user and other sources of data.
b. Error check and correct the model.
c. Formulate the underlying problem into a mathematical value functional equation.
d. Solve the equation, yielding a value in each state and a value-maximizing action (“policy”) in each state.
e. Report the results.

Practical guidance for using the toolbox is included in this chapter.

FIRST STEPS IN COMPOSING THE PROBLEM

The first steps in using the toolbox to analyze a decision problem often involve thinking about the problem, rather than using the software. The essential elements of a recursive decision problem must be identified. This is discussed in “The Recursive Model” on page 20. Also, we have created a set of worksheets for users that helps outlining a decision problem; and identify states. These are included in “Appendix C. Sample Worksheets” on page C-1

STATE VARIABLES FOR BUSINESSES

Anderson (2012) provides a prototype state vector for an operating firm, with elements from the following four categories:

1. The human capital of the entrepreneur, which he describes as “a critical variable and frequently watched by other employees, customers, investors, and (of course) the entrepreneur herself.”
2. Economic and competitive conditions in the relevant marketplace.
3. Balance sheet information, including invested capital.
4. Employees, technology or other intellectual property.

Some variation of these items are likely to be good state variables for an operating firm.

In addition, there is often a specific state variable that reflects past management decisions, and would be altered if the company changed policy. This state variable would be directly affected by the exercise of a “real option” available to the company, which may or may not be irreversible.

STATE VARIABLES FOR OTHER SUBJECT PERSONS

Possible state variables for individuals, families, nonprofit organizations, schools, countries, and other organizations include:

- Enrollment or attendance
- The number of customers, dues-paying members, or affiliates
- The weather (represented by temperature, wind speed, number of storms, rainfall, or other variables)
- Crop size, type of crops planted, growth of trees or other crops
- Number of predators or prey
- Number of people, families, churches, schools, streets, or cars
- Income, jobs, after-tax earnings, savings or investment balances
- Years of schooling, degree attainment, or other human capital measures
- Years to retirement or graduation

In effect, any variable that is observed and considered to be important to the subject of a decision problem is a possible state variable.

"Importance of Importance" Principle

The last sentence should not be read over lightly. Unless the variable is *observed* and *important to the subject,* it is not a good state variable.

One often sees poor analyses in a variety of fields that suffer from the (often fatal) logical flaw of analyzing variables that the subjects do not consider important, and as a result "finding" unusual results. One common variant of this error is to compare the response of individual residents of a country or area to a change in government policy, when that policy is unobserved or measured in a way that is foreign, unknown, or unimportant to the subjects.[6] The policy may be important to the government agency or other organization that conducted or commissioned the research, but if it isn't observed and considered important by the subjects, it is a poor choice for a state variable in a decision model.

Remember the "importance of importance" principle: if a variable is not observed and important to the subject, it is not important to the decision model.

Difference Between Econometric and Recursive Models

6. For example, consider the use (in economics) of Gini coefficients, top marginal tax rates, "effective" tax rates, and total government program expenditures as explanatory variables for a variety of models. Although these variables, in some cases, change over time in a way that provides some information about conditions affecting taxpayers and employers, very few of them actually face the top marginal tax rate, pay the average "effective" rate, or know what the Gini coefficient is.

Continuing on this line of logic, it is worth noting that econometric models that seek to explain aggregate behavior over time are not the same as recursive models, in the following way:

- Data on individual behavior that is measured by aggregates suffers from the well-known and obvious problem of aggregation and inference. Recursive models can more easily focus on the individual decision maker.
- Econometric models of time-series data usually involve a model that does not change over time, preferences that remain the same over time, and an underlying model that assumes that responses to changes in state variables are additive.[7] Recursive models, on the other hand, can involve transition probabilities that are different in different states or when actions are different. Although the recursive models considered in the examples presented in this *Guide* rely on transition probabilities that do not change over time, they do not assume that the conditions or consequences of actions do not change in a way that affects the subject's decisions.

Of course, econometric methods are sometimes the best available given the data and problem at hand. Furthermore, one may find results from econometric analyses useful in estimating parameters in a recursive model. However, always consider the fundamental difference in the logic between a recursive (two-period, decision-maker-centered) model and a multi-period, aggregate-data econometric model.

USE OF INDEX AND TRANSFORMED VARIABLES

Often the available data is messy, or involves multiple possible metrics. In some such cases, the data can be transformed into an index or representative indicator.

Although this procedure is sometimes necessary, it should be approached with care. Just as noted above, it is quite possible to transform data that the subjects see and find important to an index that does not reflect conditions important to the subject.

7. Experienced econometricians will note that the underlying model can be multiplicative and then transformed (such as by logarithms) into an additive model, or can involve a more complicated nonlinear regression. However, this does not change the main point of this observation.

Chapter 7. Discounting and Discount Rates

RATIONALE AND DATA ON DISCOUNTING

Theories of Discounting

Economists as long ago as Adam Smith conjured up reasons why consumers preferred to receive money now rather than later, and why lenders and investors insisted on what we would now call a rate of return. Smith, in *Wealth of Nations* (1776, book I, chapter 8) described how the potential for loss, the difficulty in ensuring repayment, and the possible alternative earnings all affected the market rate of interest or "ordinary rate of profit."

The early 20th-century economist Irving Fisher catalogued the reasons for "impatience" among consumers in his monumental *Theory of Interest* (1930, chapter IV).[8] Fisher even described why consumers with relatively little income had much stronger reasons to prefer income immediately than those whose income afforded them the ability to patiently invest.

In the latter half of the 20th century, a large number of theories and variations of those theories arose regarding expected rates of return for investors. The most well-known of these (such as the Mean-Variance Framework, the Capital Asset Pricing Model, the Fama-French model, the APT, and many variants of these) have been employed to create models of expected rates or return for certain investors.

For the purposes of this *Guide,* we will merely introduce the theories and authors listed above, and suggest that one or more of these could be applicable to specific decision problems. At least as important, however, is likely to be the availability of data for that problem. We discuss this next.

Data on Rates of Return

The availability of data on rates of return vary tremendously by the type of investment or security.

---

8. Fisher stated that an individual's impatience depended on the following "characteristics of his income stream:"

    1. The size (measured in dollars) of his expected real income stream.

    2. Its expected distribution in time, or its time shape—that is, whether it is constant, or increasing, or decreasing, or sometimes one and sometimes the other.

    3. Its composition—to what extent it consists of nourishment, of shelter, of amusement, of education, and so on.

    4. Its probability, or degree of risk or uncertainty. [*Theory of Interest*, IV, sec. 5]

1. Over the past few decades, price and dividend data from U.S. stock and bond markets have been compiled into data on actual rates of return (and, sometimes, defaults) on traded securities representing many classes of equity and debt in these markets. These data focus only a small fraction of companies (those that are publicly-traded), are often subject to "survivor bias," and typically ignore transaction costs. However, these data provide a very rich source of information to those that can observe both the limitations and the benefits.
2. Unfortunately, nothing remotely comparable exists for investments in privately-held firms. This is a huge area of missing information, as investments in private firms are clearly not the same (in terms of control, liquidity, marketability, and tax implications) as investments in publicly-traded firms through the stock market.
3. The burgeoning field of entrepreneurial finance—and the increasing visibility of venture capital—has provided some public awareness of investment terms that are common for start-up and early-stage firms. However, comparable data here for venture capital investments across industries is very difficult to obtain.
4. Some data exist on the value over time of investments in houses (for many households, their most valuable asset), and on what households report as their retirement savings.
5. Data on consumer loan rates are available. However, these are subject to the considerable limitations, in that many consumers avoid taking out loans (including those represented by credit card balances), and others use the strong differences in terms among these loans to shift funds from one purpose to another.
6. Finally, there are enormous data available on interest rates and yields on Treasury securities for the United States and many other nations. These, of course, do not reflect the borrowing terms of most private sector lenders or borrowers, but do provide an indication of the base rate of interest and inflation expectations.

It is important to note that the wide availability of data for two of the above categories does *not* mean that such data can be readily applied to investments or loans made in the other categories.

TWO CRITICAL ASSUMPTIONS: UTILITY, AND TIME PREFERENCE

The literature of finance, economics, and professional valuation can differ substantially in terms of their embedded assumptions, conventions, and nomenclature.[9] Unfortunately, readers familiar with one literature can easily overlook differences in assumptions in texts from another field.

9. For a pointed discussion, see Anderson (2013, chapter 1).

To avoid this problem, and to further explain the intellectual basis for the recursive approach, it is worth discussing the following two critical assumptions and how they often differ depending on the field of use:

1. "Utility" versus cash flows

   Since the emergence of the neoclassical school over a century ago, economists have typically rested their analyses on the bedrock of a set of preferences of a consumer. These preferences are embodied by a "utility function" that maps the receipt of income and other benefits to subjective benefits as seen by the consumer. Much effort within the profession has focused on the form of this utility function, the implied view of risk, and (recently) the apparent deviations from ideal models of utility that are implied by actual consumer behavior.[10] Depending on the utility function used in an economic model, the first dollar earned in a time period is often "worth" much more to a consumer than the thousandth or ten thousandth.[11]

   On the other hand, most finance and valuation texts largely ignore this entire concept, using the simplifying assumption that money earned is money earned, and that twice as much money is worth exactly twice as much. Of course, for many tasks (such as managing a large investment fund), this is quite accurate.

   The common analytical tasks in each field underline this difference: economists often "maximize utility" in their models, while finance professionals "discount cash flows" in theirs.

2. "Time separable" versus "recursive" utility

   Recursive models incorporate the series of two-period decisions that investors, managers, and consumers make. Each of these involve a trade-off between benefits (or costs) today, and expected value tomorrow.

10.The latter has spawned an entire "behaviorist" school of finance and economics, including the notion of "bounded rationality" and "loss aversion." Kahneman & Tversky (1979) is one of the seminal contemporary articles establishing this school of thought.

It interesting to note the recent re-discovery of the "behaviorist" tendencies of classical economists. For example, Adam Smith clearly outlined "loss aversion" 200 years before it was studied by contemporary economists. In *Theory of Moral Sentiments* (1759), he wrote:

> Pain...is, in almost all cases, a more pungent sensation than the opposite and correspondent pleasure. The one almost always depresses us much more below the ordinary, or what may be called the natural state of our happiness, than the other ever raises us above it. [Theory of Moral Sentiments, III, i]

On the classical behaviorists, see Ashraf, et al (2005).

11.This is the famous "diminishing marginal returns" argument that often surfaces in introductory economics classes. It is most easily expressed in terms of the desire of a person to eat when very hungry: the first several bites can be the difference between life and death; those that follow are beneficial and enjoyable; those still afterwards can result in sickness.

In contrast, the standard finance model involves a stream of future cash flows discounted to the present. The implied trade-off here is between an investment today, and a *sum* of future expected earnings. [12] This implied attitude is known as “time separable” or “additive” utility.

In fact, very few investors (or consumers, or managers) are truly indifferent to the timing of their payments. Instead, they behave as if they have *recursive utility,* meaning that they attach meaning to the timing of income and expenses as well as the net present value of their sum.[13]

Here, only a sliver of the current economics literature observes the difference between time-separable and recursive utility, with the large majority typically assuming time-separable (additive) utility.

ASSUMPTIONS IN RECURSIVE MODELS: UTILITY AND TIME PREFERENCES

With the above discussion in mind, it is important to observe that recursive models embed the following assumptions:

1. Utility in a recursive model can be measured by cash income, cash expenses, or any utility measure that can be defined.

   The analyst can define the measure used for current-period benefits. This can be—and often will be in practical models—a simple expression of income or expenses. It can also be something that incorporates a utility function that incorporates risk aversion, diminishing marginal benefits, loss avoidance, and other commonly-observed human preferences.

   In the Rapid Recursive® toolbox, this can be accomplished by using some type of mathematical construct that implements the utility function, and putting the result into the reward matrix. For example, very strong avoidance behavior can be implemented by using an extremely low number to represent the subjective reaction of the subject to a particularly bad situation that he or she wishes to avoid.

2. Time preferences in a recursive model are inherently recursive, and can mimic time-separable preferences if appropriate.

   The discounting factor in a recursive model is always recursive. This means the appropriate discount factor is the one the subject uses to compare current-period benefits and the expected value in the next time period. Combined with the wide flexibility available for describing benefits, a very large

---

12.The logic of net present value is inescapable: When the discount rate is 10%, $1 today is worth the same as $1.10 next year, $1.21 the year after that, and $2 some years after that. The net present value of each of these possibilities (and of many combinations of them) is the same.

This use of a NPV measure or decision criteria implies that an investor is indifferent to the timing of payments, as long as the net present value of the eventual sum of payments is the same.

13.Systematic treatment of recursive utility is found in Stokey and Lucas (1989, chapter 5), and Duffie (2001, chapters 1 and 2 including notes); Anderson (2013, chapter 8) contrasts time separable and recursive utility in decision models such as those discussed here.

range of preferences (including strong loss aversion, asymmetrical attitudes about earnings over different time periods, and desire to rationally evaluate time-sensitive real options) can be modeled.

This being said, it is quite possible—and sometimes desirable—to use a simple discount factor, straightforward cash earnings for a reward matrix, and structure a model so that the subject is indeed largely indifferent about the timing of cash flows as long as they have identical net present value.

PRACTICAL ADVICE ON DISCOUNT RATES

Given the discussion of theory, and then data, and then critical embedded assumptions above, we can offer the following practical advice on using discount rates in a recursive model:

1. Think hard about the underlying problem before selecting a discount rate.

   The correct discount rate will always depend on the subject of the decision problem, and only sometimes on stock market or other widely-available data. Always think hard about the problem itself before selecting a discount rate.

2. The correct discount rate will be based on factors that are apparent and important to the subject of the decision problem.

   Remember that the discount rate reflects the competing incentives for the subject of the problem, given the other factors surrounding him or her. Take those factors into account

   Avoid relying on a theory of discounting, or a data source on discount rates, that does not match the conditions of the problem.

3. Consider transaction costs and liquidity constraints in any problem.

   One of the strengths of the recursive approach is the ability to explicitly include such factors in the model. When selecting a discount rate, be sure that you don't ignore these factors.

4. The recursive structure of the problem implies the use of a two-period discount rate. This may or may not be represented by a "short term" interest rate.

   In the recursive model, the subject makes a series of two-period decisions. Therefore, they are always choosing between this period and the next. In cases where the subject can repeatedly use short-term borrowing, and faces no liquidity constraints, this suggests the use of a short-term borrowing rate as a discount rate. However, the more likely case for both businesses and households will be that repeated short-term borrowing could only involve costly unsecured credit. The alternative sources of funds are equivalent to equity financing, such as home-equity loans, refinancing of existing debt, implicit use of proceeds from car loans or student loans, and reducing household expenditures. Some of these sources have high transaction costs, and all are limited in some way.

5. Take into account inflation and growth trends consistently throughout the model.

> In times of low inflation, low interest rates, and relatively stable economic growth, one might get away with subsuming inflation expectations into a discount and growth projections without much additional thought. However, such times are the exception, not the rule. Therefore, always consider the growth and inflation trends in the problem, and recognize them consistently throughout the model.

Practical Advice by Subject of Decision Problem

Table 2, "Suggestions for Discount and Growth Rates," on page 33, contains a set of recommendations for growth and discount rates for use in recursive models. The recommendations are separated by type of subject.

Note that the advice in the table primarily reflects important considerations for the subjects of the problem, rather than extensive data or econometric analysis. This is consistent with the recommendations under "Practical Advice on Discount Rates" on page 31.

TABLE 2. **Suggestions for Discount and Growth Rates**

Subject and Problem	**Possible *d***	**Possible *g***	**Notes**
Publicly-traded corporation with access to short-term and long-term borrowing Subject considers whether to invest or re-invest in current business	commercial prime rate; commercial paper rate; "cost of capital" estimate	zero; long-term real growth	This is a common subject in corporate finance, but is the exception rather than the rule in the economy. "Cost of capital" could be estimated in a variety of ways, taking into account debt and equity mix for corporation.
Classic micro-economic household unit Subject considers whether to save or consume	borrowing rate for household assets, such as a mortgage	long-term productivity or wage growth trends	This is a common teaching and analytical example. Transaction costs, short-term uncertainty, and even mortality are often ignored in these examples.
Realistic household unit or individual worker Subject considers whether to save or consume	available borrowing rates, which could include high-cost credit card debt, mortgage rates, home equity loans, and family borrowing	human capital-influenced wage growth rate, which could reflect unemployment and risks related to health, industry, and skills	The actual constraints and market conditions imposed on typical families and workers are often severe. For realism, ensure that you have taken into account transaction costs and liquidity constraints, which can dramatically increase the costs of borrowing, as well as access (or lack thereof) to financial markets and institutions.
Closely-held business in operation for at least several years and with good financial and managerial resources Subject considers whether to re-invest in current industry, perhaps to expand or acquire a competitor, or distribute income to owners		growth trend for that company in that industry, adjusted for "key man" or similar risks inflation trends in industry	Note that even a successful privately-held firm cannot borrow using commercial paper; cannot issue stock on major exchanges; and may or may not be able to obtain conventional bank loans. Note also that the owners may have the benefits of controlling the entity in a manner that stockholders or large corporations do not.
Start-up or entrepreneurial firm; entrepreneur starting firm or creating invention Entrepreneur may be considering whether to continue re-investing in unprofitable company (or invention without current market), or stop and return to regular employment. Firm may be struggling to develop a product and get it to market; or to become profitable after a product is first introduced	Cost of entrepreneur's time and money; cost of "angel" capital; total cost of venture funding	zero; high-growth potential (if properly offset by high discount rate); modest growth rate for productivity of entrepreneur	Liquidity constraints and transaction costs are very significant in this sector. Often, the only financing available to entrepreneurs is "3-F" financing. Terms for VC and angel financing are often harsh. "Bootstrapping" for some period of time is often the only possible path. Risks, and possible rewards, are extremely high.

Source: Author's advice

TABLE 2. **Suggestions for Discount and Growth Rates**

Subject and Problem	**Possible *d***	**Possible *g***	**Notes**
General treatment of inflation in all these problems			Price inflation trends should be considered consistently across all elements in problem, including reward function, available actions, and both discount and growth assumptions embedded in discount factor.

Source: Author's advice

Chapter 8. Additional Topics in Discount Rates and State Variable Diffusions

There are a variety of conventions and practical considerations that develop over time in the use of any method. Because recursive methods are new, these conventions and practical considerations are not well established.

By *convention,* we mean a certain manner of usage that is assumed to be known and followed. Conventions are often the result of tradition, precedence, and reasoning. However, the use of a convention does not mean that there is no other way to do something, or that only one way is correct. For example, people in many parts of the world drive on the right side of the road. Others drive on the left. It is very, very important to have a convention about what side of the road should be driven. However, it is difficult to make a strong case that one is logically superior to another.

One purpose of this *Guide* is to help users of the Rapid Recursive® toolbox quickly become proficient in the use of recursive techniques. In this chapter, we discuss a set of conventions and practical considerations that can be very helpful in actual use.

TIMING OF REWARDS AND ACTIONS

The Rapid Recursive toolbox uses the following conventions regarding the timing of rewards and actions.

1. Rewards are assumed to be received at *the beginning* of each period.
2. Actions take place simultaneously with the reward.
3. The results of the action (and any random factors that are incorporated into the transition probabilities) occur with the arrival of the state at the beginning of the next period. At that point, the process begins again.

As noted below, practitioners in other fields may use this convention for timing of rewards that are cash earnings, or they may use other conventions.

Time Index Starts at t=0

The Rapid Recursive toolbox uses the convention that time starts at $t=0,$ and that the first reward is received at that point in time. Note that some summation formulas in finance textbooks and articles either omit the time subscripts, or start with $t=1$.

Observing this convention is particularly important when the discounting is also captured in the summation. For example, the following summation includes an exponent on the discount factor *beta* that is multiplied by a series of cash flows C_t:

NPV Summation (EQ 3)

$$NPV = \sum_{t=0}^{T} \beta^t C_t$$

$$\begin{aligned} &where: \\ &\quad t = \text{time index}, \\ &\quad \beta = \text{discount factor; } 0 < \beta < 1; \\ &\quad C_t = \text{cash flow at time } t. \end{aligned}$$

If the time index starts with *t=0,* the first cash flow is not discounted. However, if it starts with *t=1,* then the first cash flow is fully discounted.

CASH FLOWS AND DISCOUNTING

One implication of the beginning-of-period convention for rewards and actions is that the appropriate discounting of cash flows should also assume these cash flows arrive at the beginning of the period.

The fields of Finance and Valuation do not always follow the beginning-of-period cash flow convention. As a result, there are many formulas for the discounting of cash flows that assume income or expenses occur at other times during the time period, such as the end of the period or the middle of the period. Furthermore, many Economics texts describe theoretical models that rely upon continuous discounting.

A concept known as the "stochastic discount factor" emerged in the economics and finance literature within the past two decades.[14] The SDF is composed of both a pure discount rate *beta* and a ratio of the preferences for consumption in this period and the next. For analyzing how and why consumer-investors decide to save, the SDF concept is superior to the notion of a "market rate of interest" because it properly captures what Irving Fisher called the "impatience" of the consumer. See the equation "Stochastic Discount Factor" on page 37.

The SDF is often used in what is called the "basic pricing equation" in a two-period setting. Popularized by Cochrane (2005), this is a powerful insight into how consumer-investors value future possible earnings. This is inherently a two-period model, which means it can be naturally extended to a recursive model. However, the basic pricing equation itself is not a recursive model, and

14. John Cochrane's *Asset Pricing* (2005) is an especially good introduction to the SDF approach to pricing of traded securities, and starts with a solid foundation of consumer-investor preferences. Ljundqvist & Sargent (2000, 2012) also cover the topic, although more briefly.

can be used in other settings as well. See the equation "Basic Pricing Equation" on page 37.

Stochastic Discount Factor (EQ 4)

$$m_{t+1} \equiv \beta \frac{u'(c_{t+1})}{u'(c_t)} \text{ stochastic discount factor or pricing kernel;}$$

where :

$$\beta = \text{ subjective discount factor, and}$$

$$\frac{u'(c_{t+1})}{u'(c_t)} = \text{ratio of marginal utilities.}$$

Basic Pricing Equation (EQ 5)

$$p_t = E(m_{t+1}x_{t+1})$$

where :

$$p_t = \text{ price of asset with payoff } x_{t+1}$$

$$m \equiv \beta \frac{u'(c_{t+1})}{u'(c_t)} \text{ stochastic discount factor (pricing kernel)}$$

$$x_{t+1} \equiv \text{ payoff at time } t+1$$

DISCOUNT AND GROWTH RATES

In addition to the differences between beginning-of-period and other conventions regarding the timing of cash flows, there are also a range of possible conventions regarding the incorporation of discount and growth rates (as well as, implicitly, inflation rates) in the discount factor used in a value functional equation.

The discount factor in value functional equations is typically denoted by the greek letter *beta.* It is sometimes called the "gross discount rate" because, in the form used in functional equations, it captures all of the motivations for preferring the receipt of benefits this time period rather than the next.[15]

Practitioners in business, finance and accounting are more commonly exposed to a net discount rate or (net) interest rate, which is commonly denoted by *d* or *r.* In general, the gross discount rate and net discount rate are related by the following equation: *beta = 1/(1+d).* For example, if *d=.10,* then *beta=.91.*

15. It is sometimes called in economics the "marginal rate of intertemporal substitution," because it represents the rate at which a person will substitute consumption next time period for consumption this period.

However, because the gross discount factor *beta* captures *all* of the motivation for substituting income in this period over the next, economic theory does not provide one well-defined method for calculating that factor from component motivations, or even for estimating it.

The following equations provides a set of possible conventions for the calculation of the discount factor in the value functional equation that is the heart of a discrete-time recursive model. In general, we use the standard convention listed.

Conventions for Discounting (EQ 6)

Standard Convention:

$$\beta = (1+g)/(1+d)$$

$$\begin{aligned} &where: \\ &\beta = \text{gross discount rate per period} \\ &d = \text{net discount rate per period} \\ &g = \text{growth rate trend per period in state variable} \end{aligned}$$

Alternate conventions:

$$\beta = 1/[1+(d-g)]$$

$$\beta = (1+\pi)/(1+\rho+\pi)$$

$$\begin{aligned} &where: \\ &\pi = \text{inflation rate} \\ &\rho = \text{real interest rate} \end{aligned}$$

However, as also noted in the equations, there are alternate conventions even for discrete-period calculations. If the numbers for inflation, growth, and discounting are small, the different conventions may produce a very similar result for *beta.* However, if any one is large, there could be a large difference.

Note on Continuous Compounding

If one compares continuous compounding versus annual compounding of discount rates, it is usually necessary to make an adjustment in one or the other even if the net discount rate is small. A further discussion of this topic is in Chapter 16 of *Economics of Business Valuation* [Anderson, 2013], including the use of both continuous and discrete discounting, and the application of different conventions for the calculation of *beta.*[16]

DECOMPOSITION OF TOTAL RETURN AND THE HJB FORMULA

The value functional equation that underlies the recursive model allows for a much richer decomposition of the “total return” of an investment. It is common for practitioners in Finance and Valuation to describe the total return on a stock investment to be the sum of dividends and any capital gain or loss. Ignoring (again, as is common) transaction costs and the timing of capital gains and dividends, this is often a useful approximation.

A richer decomposition is possible for both standard assumptions regarding stock price movements and for the value arising from a value functional equation. In both cases, this arising from the assumption that a critical variable (the stock price, stock price plus dividends, or the value of a state variable) follows a stochastic difference equation or *diffusion.* The implications of this for stock prices are found in many Finance and Economics textbooks, as well as in references presenting derivations of option pricing formulas such as the celebrated Black-Scholes-Merton formula.[17]

Decomposing the change in value from a value functional equation, when the underlying state variable is subject to both trend growth and uncertainty, is possible using an implementation of the Hamilton-Jacobi-Bellman equation that arises from the calculus of variations. This novel application was first presented in Chapter 16 of *Economics of Business Valuation* [Anderson, 2013]. The example business in that publication was also evaluated using a prototype of the Rapid Recursive toolbox. Thus, a back-and-forth comparison between a decomposition using an analytical model and a numerical calculation is possible, for at least some subjects.

In such comparisons, there is a usage convention to note regarding the state variable. It is common in stochastic calculus to work with the differential form of the transition equation for a state variable, such as the following:

16.In one example given in Chapter 16 of *Economics of Business Valuation,* an annual net discount factor of d=.15 was converted to a continuously-compounded discount factor of rho=.14, and a discrete growth rate of g=.05 was converted to a continuous growth factor of 0488.

17.A standard reference for option pricing is Hull (2005), *Options, Futures, and Other Derivatives, 5th ed.* A much more rigorous reference is Shreve (2004), *Stochastic Calculus for Finance.* See complete references in “Appendix B. Bibliography and References” on page B-1.

Diffusions, Differential and Integral Forms (EQ 7)

Differential Form:

$$dS = \mu(s,t)dt + \sigma(s)dz$$

or

$$dS = g(S,x)dt + \sigma(s)dz$$

Integral Form:

$$S(t) = S(0) + \int_0^t \mu(s)dt + \int_0^t \sigma(s)dz$$

Functional Form:

$$S_{t+1} = g(S_t, x_t)$$

or

$$S_{t+1} = g(S_t, x_t, \varepsilon_t)$$

where:

ε = random disturbance

z = Brownian motion

Here, the "g" is not the same as the "g" in the functional form of a transition equation.

The only published example (as of this printing) of the use of the HJB decomposition of total return for a subject company that is also evaluated using an iterative recursive solution is in Anderson (2013, chapter 16). In that case, a "g prime" notation is used to indicate that the relevant factor in the HJB equation is the *change* in the state variable, which could be *g'(s)* in the integral form, and simply *g* in the differential form.[18]

18. In addition, it is important to work carefully with the units and the increments of time and state variables. In the cited example, the "g prime" variable involved a trend growth rate that was multiplied by the value of the state variable to yield the increment in the state variable, which is the conceptual basis for the use of the formula.

Chapter 9. Setting Up the Reward Matrix

This chapter provides practical advice in setting up the reward ("R") matrix for the Rapid Recursive® toolbox.

WHAT DOES THE "R" MATRIX REPRESENT?

The elements of the *R* matrix represent the *reward,* which may also be termed the "profit," or "benefit," or "cost" that the subject receives when in a certain state and performing a certain action.

We consider *R(i,j)* to the be element in the *ith* row and *jth* column of an *R* matrix that is sized *SxA*. (Here, *S* is the number of states, and *A* is the number of actions.) The *R* matrix therefore contains a single entry describing the reward of every state-action pair. For example, if there are 3 states and 2 actions, there are *3x2=6* entries in the *R* matrix.

By convention, in the *R* matrix states are represented by rows; actions by columns.

WHAT IS THE REWARD?

The worksheets in "Appendix C. Sample Worksheets" provide an aid to identifying states and actions, and then rewards. These are intended to help organize your thoughts about a particular problem. We suggest that you start with the thought process described in the worksheets in order to most efficiently organize a problem.

In particular:

- Worksheet 1., "Subject of Decision Problem" in Appendix C, asks about the kind of reward, cost, benefit, or earnings that flows to the subject of the problem.
- Worksheet 2., "States and Rewards" in Appendix C helps identify a "benchmark" state for which the reward can be most easily estimated.
- Worksheet 3., "Actions and Effect on Reward" in Appendix C assists is estimating the reward for the subject under different state-action pairs.

ESTIMATING ELEMENTS OF THE REWARD MATRIX

Below is some additional advice on determining the reward and estimating the elements of the reward matrix:

1. Think about the entire problem (the description, the states, and the actions) before attempting to estimate the numeric values in a reward matrix.
2. When you identify the reward, start with a benchmark state (market condition, economic condition, business demand, weather, cost, or other factor). Then select the most common, or average, or other indicative action the subject is likely to take during the benchmark state.

   This benchmark state and action provide the basis for a benchmark reward.

3. The “reward” (or the cost) should be the *net* benefit that the subject *actually realizes*. This means, among other things, that transaction costs, taxes, retained earnings or commissions, and other factors that directly affect the benefit to the subject should be considered.
4. For a business, you may want to start with a traditional income statement, and then work through to the distributed earnings to shareholders.
5. If possible, you can consider rewards net of any income tax liability. However, if you do the reward analysis on after-tax income, be sure you have used after-tax discount rates and after-tax costs throughout.
6. You may wish to do some calculations in a spreadsheet, or have the relevant data already in a spreadsheet. Sometimes, it is a simpler matter to set up a set of income statements that are identical excerpt for differing market conditions and actions. From these, a set of rewards (one for each of a set of state-action pairs) can be identified.
7. If the data are in a spreadsheet, make use of the extensive data import capabilities of Matlab®. A number of methods for importing data are listed in Table 3, “Importing Tabular (Matrix) Data from a Spreadsheets and Other File Formats,” on page 45.

Often, the documents, data, and information necessary to begin creating the *R* matrix include basic accounting, business plan, cost schedule, or similar information that is readily at hand. However, in our experience, one usually has to search to understand and then estimate transaction costs, changes in rewards caused by various actions.

In addition, there is often a significant difference between the numbers reported on an accounting statement, cost schedule, or cash flow schedule and the actual reward to the subject of the problem. In particular:

- The amount listed on an income statement as “profit” and the actual reward distributed to shareholders are often very different. The dividend policy of the company (or the similar policy of making distributions to members, partners, or shareholders), as well as tax treatment and liquidity constraints should be considered here.
- Individuals may expect future earnings from business investments, increases in human capital (such as acquired through additional time in school or training), or committing to work additional hours or invest. In such cases, the recursive method allows for explicit recognition of risks, delays, and costs. Make sure that you have at least considered these when you set up a model involving these expectations.

Chapter 10. General Advice: Getting Data into the Workspace

Part of the power of recursive models is the ability to bring more information into the problem. With the Rapid Recursive® toolbox, that “more data” includes the *R* and *P* matrices, which can constitute dramatically more information about what could happen in the future than is incorporated in a typical one-scenario discounted cash flow schedule.

The advantages of more information are obvious. Once you begin working with practical models, you are also confronted with the cost: you have to enter more data. Much of this *Guide* involves advice on how to structure the problem, what information is necessary, and how to use the information once it is entered into the workspace. This chapter deals with a more mundane, but very important topic: getting you data into the workspace.

ADVICE ON GETTING TABULAR DATA INTO THE MATLAB WORKSPACE

The Rapid Recursive® toolbox requires MATLAB®, a powerful mathematics, programming, data analysis, and visualization computer software.[19] Matlab includes a wide range of capabilities for importing and exporting data from other file formats. All of these capabilities are available to users of Rapid Recursive® toolbox for Matlab.

In the following table, we summarize a set of possibilities for getting tabular (matrix) data into the Matlab workspace, where it can be accessed by Rapid Recursive commands as well as all the functionality of Matlab. There are often multiple ways to get data from, say, a spreadsheet into the workspace. For data in other formats (including web pages and “flat” files such as simple ASCII files), there are typically at least two methods. One method is typically direct and precise, using a specific command for that type of data. The other, less precise but sometimes more reliable, typically involves pasting data into another file (such as a simple text file, or a spreadsheet) and then importing it.

Keep in mind when you review this advice that computer programs (including both Matlab and the Rapid Recursive toolbox) change regularly, and that file formats also change. Therefore, check the current documentation for the relevant software that produced the data file, as well as for Matlab, to see if there are changes in these methods or new methods.

19.MATLAB® is a product of the Mathworks, whose website is: http://www.mathworks.com.

IMPORTANCE OF CHECKING AND SOURCING

Sophisticated data entry and analysis is far less important that actually knowing your data!

Although this *Guide* cannot address the broad and important topic of data for specific problems, we feel two admonitions are essential:

- Always double-check to ensure that the data you have instructed the computer software and hardware to import is the data you have in your workspace.
- Always properly source your data. This typically involves writing down (in a manner that other users, and you, can see in the future) the source of the data and its precise subject, as well as the date the data were retrieved and other information (such as units, periodicity and revisions) that may also be relevant.

In addition, we also advise those investigating specific issues to budget enough time for the important tasks of identifying, finding, checking, and (if necessary) adjusting and estimating data. For these tasks, we offer the following additional bits of advice:

- A website address alone is almost always *not* an acceptable source note.
- Government statistical agencies, even for the United States and other developed countries, do not have data on most private business activities as well as earnings and expenditures of private individuals. They do have aggregate data, which is of course extremely useful. However, you should expect that private business and household data will require some work to acquire.
- There are multiple measures for many important business and economic variables. Don't assume that "inflation" or "earnings" will be unambiguous concepts for an economy or a business.
- Consider the source of the data you acquire. Remember that organizations—including governments and think tanks—that collect and provide data have incentives and organizational interests, too, and these (unfortunately) sometimes color the data that are provided to the public.

TABLE 3. **Importing Tabular (Matrix) Data from a Spreadsheets and Other File Formats**

Method	Matlab Commands	Notes
Directly read in data from specific cell ranges in a Microsoft Excel spreadsheet and some other spreadsheets such as OpenDocument spreadsheets	xlsread	Ensure that you have properly set up the cell ranges., and have the worksheet in the Matlab path
Get data from “flat” and delimited files, including ASCII files and.csv files	dlmread textscan load	These may work on data that are cut-and-pasted, or downloaded into files with little or no formatting. Careful choices of delimiters (spaces, paragraphs, tabs) and double-checking of data is usually important
Use data import commands to import data from other files, including: scientific data in CDF, HDF, netCDF,	importdata uiimport textscan dlmread fitsread	A wide variety of file s in various formats can be accessed using these commands.
Get data from XML files (including some web pages)	xmlread	
Cut-and-paste		For small sets of numbers, this method is sometimes the fastest. Check for typing errors and for proper matrix syntax, as well as for mixed data types.
Import data facility	importdata	Depending on Matlab version, it may be accessed in different ways
Create script file to read in data	after using “import data” facility in Matlab, save to a script file	This is an excellent method if yo repeatedly import data. Be sure to comment the script file so that you or your colleagues can revise and re-use it in the future.
Save data into (.mat) file once entered; reload when starting session in future	save load	

Source: Author’s advice; Mathworks documentation.
(Note: Excel, Matlab, and OpenOffice are trademarks).

Chapter 11. Troubleshooting Recursive Models

This chapter contains practical advice on troubleshooting recursive models. The suggestions in this chapter arise first from the author's experience in developing the prototype for the Rapid Recursive toolbox, using the toolbox to create solution templates, and from actual use in consulting engagements.

Order of Troubleshooting Advice

We suggest that users consult the information in this chapter in the following order:

1. First, review and confirm all of the "First Checks: License, Program, and Platform" on page 46 if you have any difficulty starting the program, or if you encounter error messages that suggest the toolbox is not running, or not visible to Matlab.
2. Then, if you are certain that the toolbox is properly licensed, appears in the Matlab path, and encounters errors that prevent it from completing the analysis, try "Second Checks: Data Conformance" on page 47.
3. If you are running the toolbox without incident, and successfully composing, error-checking, and solving a sequential decision problem, but are unsure of the results, try "Third Checks: Model Composition" on page 48.

FIRST CHECKS: LICENSE, PROGRAM, AND PLATFORM

If you have any difficulty starting the program, or if you encounter error messages that suggest the toolbox is not running, or not visible to Matlab, or that commands are completely failing to execute, try confirming all of the following:

1. You have a valid license to the Rapid Recursive toolbox, which has not expired. (If it has expired, you can purchase a Maintenance from Supported Intelligence that provides program updates and other benefits.)
2. The toolbox is entirely within the Matlab path. It should show up with the ver command as "Rapid Recursive toolbox" in the Matlab command window. For example, you should see something like the following output:

CODE FRAGMENT 1. Checking Version and Installation

```
MATLAB                                   Version 8.0      (R2012b)
...
Rapid Recursive Toolbox for MATLAB       Version 1.0.0
...
```

3. Use the Matlab utility for checking the path. It should indicate that all the folders in which the Rapid Recursive toolbox files are in the path, including

the solution templates. Also confirm that the path includes any files you have referenced, or are using directly in your routine or solution template.

4. See if you can load and run one of the Solution Templates supplied with the toolbox. If you can, and can look at the output, you probably have a correctly-installed program with a path that covers all the toolbox files. If you can't, start checking the items above again.
5. Then, see if you can "publish" it (create a report from running it) using the "publish" functionality in Matlab. Running it with the (default) format of HTML is usually enough to confirm that the program is installed correctly.
6. If you are having difficulties with one format for a published report (e.g., HTML, MS Word or Adobe Acrobat), try "publishing" it into another format, and then switching formats.

   Note that "publishing" requires interaction among a set of document programs and the operating system; such as Adobe Acrobat, Microsoft Word, and various web browsers, along with OS X or MS Windows; and problems in this interaction are often caused by these other programs or the operating system.) Sometimes, Matlab will create a report but not be able to call the document program to open it. Check using the file explorer (or similar utility) for your operating system, and see the "help" information for the Matlab publish command.

For all of these, you can consult the *User's Guide* to the Rapid Recursive toolbox, and also the Matlab documentation.

SECOND CHECKS: DATA CONFORMANCE

These steps are useful once you confirm that you have the properly licensed product, and all of the toolbox and any files you reference are also in the Matlab path, but you have difficulties composing a problem, solving it, or reporting the results. In this situation, try the following steps:

7. Confirm you are using the built-in function to create an "input" structure, thus ensuring all the required fields exist. If you have chosen to create one in a different way, check that the fields required by any commands that use the input structure exist, and are properly filled with data elements that match the expectations of the problem.
8. Check that none of the required fields of the input structure are "NaN" or complex numbers, and that "inf" and "-inf" either does not appear, or does so only in the specific and narrow instances where such an extreme input is proper.
9. Confirm that you are not inadvertently deleting elements that are required by the toolbox, or deleting the "input" structure entirely, or substituting another. Try the debug command in Matlab to step-by-step through a Solution Template, checking the input structure along the way.
10. Confirm that you are running the RRcheckinputs command, and that there are not error messages.
11. If errors occur when after you have composed the problem and error-checked it, try using both the RRvalueiteration and RRpolicyiteration com-

mands. Read the sections in the *User's Guide* to ensure that you have the necessary elements for both.

12. After you have created the entire "Input" structure, and run RRcheckinputs, try pulling out individual fields in the Input structure and examine them. For example, the following command sequence would allow you to check the state vector, size of state vector, action vector, size of action vector, and both the R and P matrices.

CODE FRAGMENT 2. Checking Elements of Input Structure

```
Input
Input.S
Input.A
size(Input.R)
size(Input.P)
Input.states
Input.statelabels
Input.actions
Input.actionlabels
Input.R
Input.P
```

In your examination, confirm that the parameters {*S*,*A*} properly match the size of the P matrix (*SxSxA*) and R matrix (*SxA*), and that the "states" and "statelabels" fields (for example, Input.states and Input.statelabels) are both the same size (*length=S*), and that the "actions" and "actionlabels" fields similarly have *length=A*.

13. Make sure that the "input" structure you created (and just examined) is actually the one being given to the solution command (such as RRvalueiteration), and that is hasn't been changed further or had fields deleted inadvertently.

14. If you have a real option that can be exercised in the model, ensure that you have correctly identified it in terms of an action, a consequence to that action (including possible elements in the state vector, reward matrix, and probabilities in the transition probabilities matrix).

THIRD CHECKS: MODEL COMPOSITION

These checks are useful if you are having difficulty composing the problem itself. Such difficulties may show up as an inability to get the required inputs together to compose and solve a problem, or as a feeling of being overwhelmed with data and unable to make sense of it.

15. Make a *much simpler* model of the decision problem that is posing difficulty. Remember, the purpose of this model (and all models) is not to capture *all* of the aspects of a problem faced by complex human beings. It is to examine only a few aspects of the problem, preferably the ones that are most important to the subjects!

    By "much simpler," we mean only 2 or 3 actions, and 3 or 4 states (including any absorbing or null states); fairly simple R and P matrices; common-sense parameters for {*beta, g, d*}; and a short description in writing that is understandable to a non-expert in the relevant field of study.

16. Run the much simpler model successfully. Then, add one element of complexity at a time, each time fully describing it and adding the required numerical elements to the model and its parameters.

17. Use one of the supplied Solution Templates (or one of the contributed templates available on the File Exchange on the Supported Intelligence web site. Adapt one of these, step-by-step documenting the changes and observing how the results change.

    If you do adapt one of the Solution Templates contributed by another user (or authored by one of the Supported Intelligence development team), please recognize the original author in any published research or similar document you create based on the adapted work you created. Of course, the results from that adapted work are due to you, but good scholarship requires citing the works of others when you base your work on theirs.

FOURTH CHECKS: MODEL CALIBRATION

The tips in this section are for those that can start and run the toolbox, enter data, compose the problem as intended, solve the problem, and report the results. Congratulations! You are able to do something that eluded generations of researchers familiar with the theory of recursive models, but unable to find a practical tool to use it.

However, you may have completed composing, solving, and examining the results of a model—but do not yet fully believe the results. Consider taking the following steps in such instances.

18. Re-check your input data. Many times, you will start with placeholder assumptions in order to quickly get a model working, and then refine the inputs. At this point, check to see if there are any obvious improvements you can make in the input data.

19. Re-check that you have followed the following conventions:

- In the *R* and *P* matrices, each current-period state is represented as a row.
- In the *R* matrix, each action is represented by a column. Thus, the *SxA* matrix provides an element for each possible combination of current-period state and current-period action.
- In the *P* matrix, each action is represented by a frame in a 3-dimensional matrix of size *SxSxA*.

Often, hand-writing (or adding in comments to the computer file) the states and actions in the margins when you print out the *R* and *P* matrices helps reveal any problems.

20. Re-check that the order of states, and the order of actions, is the same across all elements of the input structure. Sometimes, a null or absorbing state is incorporated as the last state in the state vector. If this is the case, make sure it is always the last row—and not the first—in Input.state, Input.statelabel, Input.R, Input.P and all other data.

21. Re-check the calculations for the reward matrix. Remember that the value function will be solved as a function of a set of possible reward functions, and if you get the reward functions incorrect, you will get the value function incorrect as well.

22. Re-check the probabilities embedded in the transition probability matrix. Make sure you haven't inadvertently forced a result that you did not intend. Also, make sure that, once a real option is executed that should result in the subject remaining in an absorbing state, the *P* matrix puts the subject there with probability=1, and the subject stays there with probability=1.

23. Compare the value calculated in the recursive model with a naive discounted cash flow value, calculated as the simple net present value of the most likely (or expected) outcome. You can use some of the built-in functionality of the RR toolbox for this (for example, by simply capitalizing an element of the reward matrix), or calculate it some other way.

    There is some tension here between methods that is worth noting. To allow a valid comparison, you should make sure that the elements of the reward matrix you are using as the basis of a capitalization formula represent a state-action pair that could be maintained for several time periods. Of course, the implausibility of such an assumption may be one reason you are using a recursive method in the first place! As demonstrated in some of the solution templates that are bundled with the toolbox, naively assuming that current earnings can be maintained in perpetuity can give an investor or manager an extremely misleading—and potentially financially ruinous—picture of the world.

24. Re-document your work in the style of a solution template, explaining the inputs step-by-step. Now, "publish" the solution template, and share it with a colleague. Ask the colleague if he or she can walk step-by-step through your analysis. If that colleague cannot, you may have a deficiency in your explanations, even if the decision problem is properly composed and solved!

    From the author's experience, one of the first consequences of forcing yourself to fully document a model is to acquaint the creator of the model with multiple opportunities to correct and improve it. You may find you have done that well before you share the results with your colleague.

Part III. Selected Solution Templates

Chapter 12. Selected Solution Templates

INTRODUCTION

The following sections contain a summary of a set of Solution Templates that are included with the Rapid Recursive® toolbox, or available on the File Exchange portion of the Supported Intelligence web site.

For each of these, we have presented an overview including the following elements:

- Description of the problem, including a description of the subject of the problem
- The states used in the problem
- The actions available to the subject of the problem
- The transition function that governs how the state evolves over time
- The reward function that determines what reward is provided to the subject of the problem, depending on the state and action selected
- A presentation of some portion of the results
- A discussion of how the template could be adapted for other problems

For some of the templates, additional information is included, such as:

- selected output to the command window
- graphical output

PURPOSE AND LIMITATIONS

These templates have the following purposes:

1. They provide examples of how a recursive model works in practice, to assist users in learning about such models.
2. They provide the user of the Rapid Recursive toolbox a set of ready-to-run models that incorporate all of the elements that are essential to a properly-running recursive model, to assist the user in understanding how such models can be applied to a variety of problems.
3. They are also intended to be adapted to the particular situations that the user might wish to analyze in the future. Thus, they intentionally cover different topics and involve different features.

The primary purposes listed above are for teaching and providing practical bases for applying these models. Of course, this primarily educational purpose carries with it certain limitations. These include:

- Models designed to illustrate concepts will, necessarily, not be representative of most practical situations. Therefore, *all* these models should be adapted before using in any practical situation.
- Simplifying assumptions about representative markets, businesses, and people were used to create these examples. When the focus of the problem is on a topic

that is different, or where the persons involved are different that the abstractions used here, you should change the parameters and assumptions of the model.

- All models—from the simplest to the most complex—are simplifications of reality. Therefore, *every* model will miss some aspect of the problem it is intended to represent.
- No one knows the future. Therefore, any forward-looking model will be "incorrect" in terms accurately predicting the future. Remember that the purpose of a decision model is to help make a decision, not predict the future.
- Putting bad information into a good model is a sure way to get unreliable results. It is much better to have a tractable, simple model with good information than to shove bad information into a sophisticated model.

The judgment of the user, not the sophistication of the model, is most important in any decision.

AUTOMAKER MARKET SHARE

A. Problem Description

A large manufacturer of consumer goods such as automobiles and trucks considers increasing its investment in product development. The new investment could increase its market share in the future, but will also reduce its profits immediately. This could occur due to a difficulty in successfully penetrating a certain large market, due to a product that is considered outdated by the consumers in that market. Updating the product is likely to increase sales and profits in the future. However, it will cost a substantial amount of money up front

States: The states in this model represent the state of the economy, and consumer demand for the product, which in turn determine the potential sales and likely profit of the automaker.

Actions: The automaker is able to choose the investment level corresponding to a modest investment, big investment or nothing more than the minimum required reinvestment.

Transitions and reward: A higher level of investment is likely to produce more success for the automaker in terms of a higher market share, keeping current profits, implies not investing in product development and a continued decline in the market share. Choosing modest investment slightly increases the chances of improving the product enough to gain market share.

B. Suggested Uses

This example could be adapted for valuation and investment problems involving companies such as: a. Any manufacturing company. b. Supplier in a market requiring heavy investment to maintain state-of-the-art technology, or heavy expenditures to remain current with fashion or tastes. c. Companies in industries undergoing structural change.

C. Interpreting the Results

The value of a manufacturer (such as an auto manufacturer) that must operate with huge capital investment is very sensitive to volume. Economies of scale are vital to sustain the heavy R&D, advertising, distribution, and other costs, as well as the expensive costs of operating a manufacturing plant. Therefore, obtaining sufficient market share is vital if the company can operate efficiently. At the same time, product quality and the consumers' perception of value is affected by product features that are affected by R&D, materials, quality assurance, and other expenses. The tension between making investments (and accepting lower or negative current profits) and avoiding them (and reducing future market share) is often the central management focus of such a manufacturer.

Keys to using this model in practical cases include:

- Properly identifying the discretionary expenditures that could be undertaken that would increase product quality or desirability.
- Estimating the amount of such expenses that are irreducibly or reflect the company's current policy.
- Estimating how likely such actions are to increase earnings in the future.
- Formulating the inputs such that the rewards for investing in additional features, quality, or materials are anticipated in the proper time order. Often, such expenses must be incurred 1 or more years before they can expect to result in significantly improved market share.
- Compiling these assumptions (which may be based on recent historical financial statements, along with budgets for investments, and sales projections under different product quality assumptions). This may be done outside of this toolbox, and often involves management input and professional judgment.
- Comparing the value, and the value-maximizing action, from the recursive model with the results of a naive discounted cash flow analysis performed with the assumption that the current management policies are continued, or that the company maximizes current-period profits without considering the effects on long-term value.

Below is the input structure developed for the decision problem.

FIGURE 1. Input Structure for Automaker Market Share Decision Problem

```
Input =
             desc: 'Automaker Market Share'
           States: '-----'
                S: 4
           states: []
      statelabels: {'poor'  'mediocre'  'gaining'  'hot'}
          Actions: '-----'
                A: 3
          actions: []
     actionlabels: {'invest minimally'  'invest modestly'  'invest big'}
       Transition: '-----'
                P: [4x4x3 double]
            formP: 'Automaker Market Share'
           Reward: '-----'
                R: [4x3 double]
            formR: 'Automaker Market Share, from spreadsheet'
         Discount: '-----'
             beta: 0.88696
                d: 0.15
                g: 0.02
   OptionalInputs: '-----'
          epsilon: []
          policy0: []
          maxiter: []
               V0: [4x1 double]
 policyevalmethod: []
          verbose: 0
         TimeInfo: '-----'
```

```
periodicity: []
          T: Inf
      Notes: '-----'
      note1: 'R matrix calculated outside of toolbox.'
      note2: []
       date: '11-Nov-2012 22:14:00'
  TimeIndex: 'years'
```

Below are excerpts from output to the command window when running the solution template that composes the problem, and then solves it. You will note that some key information is echoed to the screen, and that a "results table" is created that shows, for each state, the value and value-maximizing policy.

FIGURE 2. Automaker Market Share Code Fragments

```
Set the action labels
 States:
   'poor'    'mediocre'    'gaining'    'hot'

 Actions:
   'invest minimally'    'invest modestly'    'invest big'

----------
 Reward matrix created.
 Automaker Market Share, from spreadsheet
----------
 Checking inputs for validity.
----------
 RR Toolbox: The size of "P" and "R" conform.
----------
 RR Toolbox: "P" passed all checks.
----------
 RR Toolbox: "R" all checks.
----------
 RR Toolbox: "beta" passed all checks.
----------
 RR Toolbox: "V0" passed all checks.
----------
 RR Toolbox: No value has been entered for the following optional inputs:
   "epsilon"
   "maxiter"
   "verbose"
 Default values will be used for these inputs.
----------
 RR Toolbox: Checks on all inputs have been completed and no errors were
 found.
----------
 Solve functional equation.
 Iterations and calculation time:
       228.00         0.01

'Value'           'Policy'
   'poor'          [2536037960.00]    [2]
   'mediocre'   [3149652544.00]     [3]
```

```
  'gaining'      [3478267078.00]    [3]
  'hot'          [3828268407.00]    [1]
To publish the results in a formatted document, use the "publish" command.

State, value in each state, and best policy:
Automaker Market Share
  'State'     'Value'       'Policy'
  'poor'         [2.536e+09]    [2]
  'mediocre'    [3.1497e+09]    [3]
  'gaining'     [3.4783e+09]    [3]
  'hot'         [3.8283e+09]    [1]
```

State	Value	Policy
poor	2536037959.69	invest modestly
mediocre	3149652544.38	invest big
gaining	3478267078.23	invest big
hot	3828268406.94	invest minimally

BEVERAGE WHOLESALER VALUE

A. Description of Problem

A person is considering acquiring or investing in an established wholesaler of alcoholic beverages. The wholesaler distributes the products of major beer suppliers (such as, in the United States at this time, Anheuser-Busch, MillerCoors, Heineken, or Modelo). The potential acquirer, a privately held, well-capitalized and profitable business, (or investor with substantial resources) is estimating the value of the wholesaler if it is managed in an effective manner.

The wholesaler operates in a market area that is considered to have a relatively stable economic base and is growing slowly in population. The wholesaler has been operating at a relatively stable output level over the past five years. However, if the wholesaler were acquired, the company can choose to expand or reduce operations, or keep operations at the same level as previously.

States: the states in this model represent the state of the U.S. economy, and consumer demand for these beverages, which in turn determine the potential sales and likely profit of the wholesaler.

Actions: the wholesaler is able to chooses the investment level corresponding to: an expansion or reduction of operations, or the same level of output as prior to the acquisition.

Transitions and reward: a higher level of investment leads to lower current period profits, but a higher likelihood of higher profits in the future.

B. Interpreting the Results

The value of an operating, profitable wholesaler is usually positive and directly related to the gross profits of the wholesaler. In industries with strong branded products, the franchise rights are quite important to preserving the ability to earn profits in the future. This analysis presumes that the wholesaler has the expectation of keeping the brands that it is currently distributing.

Tasks that are important when using this model in practical cases include:

- Properly identifying the discretionary expenditures that could be undertaken (or savings found) by management action
- Considering how likely such actions are to increase earnings in the future -- After properly formulating the inputs (P and R matrices) for transitions and the rewards, see if expanding marketing or other efforts produces enough additional earnings to increase the value of the company. The results table will show the value-maximizing action for each state.
- Comparing the value, and the value-maximizing action, with the results of a naive discounted cash flow analysis performed with the assumption that the current management policies are continued, or that the company maximizes current-period profits without considering the effects on long-term value.

Below are excerpts from output to the command window when running the solution template that composes the problem, and then solves it. You will note that some key information is echoed to the screen, and that a “results table” is created that shows, for each state, the value and value-maximizing policy.

FIGURE 3. Beverage Wholesaler Code Fragments

```
--------------------
Beverage Wholesaler Valuation
--------------------
 States:
    'Low Sales'
    'Mediocre Sales'
    'High Sales'
    'Very High Sales'

 Actions:
    'Invest Less'
    'Maintain Investment'
    'Invest More'

--------------------
 Discount and growth rates; discount factor beta:
        0.15        0.02        0.89

 Time index and horizon entered.
 Inf years
 Transition probability matrix created.
   thick; balanced
--------------------
```

```
State, value in each state, and best policy:
Beverage Wholesaler Valuation
   'State'           'Value'        'Policy'
   'Low Sales'            [20113350.00]    [3.00]
   'Mediocre Sales'     [20527977.00]    [3.00]
   'High Sales'            [20842131.00]    [3.00]
   'Very High Sales'     [21251631.00]    [2.00]
```

State	Value	Policy
Low Sales	20113350.31	Invest More
Mediocre Sales	20527977.49	Invest More
High Sales	20842130.76	Invest More
Very High Sales	21251630.76	Maintain Investment

Below is the input structure developed for the decision problem, and the output structure. Note that some fields have not been filled in this problem. For some parameters, the toolbox software will use default values if one is not supplied by the user.

FIGURE 4. Input and Output Structures for Wholesaler Valuation Problem

```
Input =
            desc: 'Beverage Wholesaler Valuation'
          States: '-----'
               S: 4
          states: []
     statelabels: {1x4 cell}
         Actions: '-----'
               A: 3
         actions: {'Invest Less' 'Maintain Investment'  'Invest More'}
    actionlabels: {'Invest Less'  'Maintain Investment'  'Invest More'}
      Transition: '-----'
               P: [4x4x3 double]
           formP: 'thick; balanced'
          Reward: '-----'
               R: [4x3 double]
           formR: 'beer wholesaler income statements'
        Discount: '-----'
            beta: 0.88696
               d: 0.15
               g: 0.02
  OptionalInputs: '-----'
         epsilon: []
         policy0: []
         maxiter: []
              V0: [4x1 double]
policyevalmethod: []
         verbose: 0
        TimeInfo: '-----'
```

```
    periodicity: 'years'
              T: Inf
          Notes: '-----'
          note1: 'R matrix calculated outside of toolbox.'
          note2: []
           date: '11-Nov-2012 22:14:00'
      TimeIndex: 'years'

Out =
             desc: 'Beverage Wholesaler Valuation'
            Input: [1x1 struct]
          Results: '-----'
                V: [4x1 double]
           policy: [4x1 double]
     resultstable: {5x3 cell}
            Other: '-----'
       iterations: 159
  calculationtime: 0.0069057
        algorithm: 'Value function iteration'
            Notes: '-----'
            note1: []
            note2: []
             date: '11-Nov-2012 22:18:09'
    DCF_valuation: 1.9904e+07
```

BASIC "BLACK SWAN" ANALYSIS

A. Problem Description

Setting: A manager or investor must decide whether to operate; operate but pay extra to insure against a loss (or improve the chances of greater earnings in the future), which is a discretionary expense or investment; or sell an interest in a business, operation, or similar enterprise. Market conditions, which are outside the control of this person, could improve (increasing the value of the investment), or decline. In one state, the entire investment is wiped out.

States: The states in this model represent economic or market conditions. Better conditions are about as likely as worse conditions, but one "black swan" possibility of a catastrophic loss exists.

Actions: The subject person chooses to continue, insure, or sell.

Transitions: The consequence of reinvesting is to reduce earnings in the current period, but increase the possible gain (and loss) in the future. Market conditions are about as likely to increase as decrease, with a small chance of a total loss of the investment. Thus, risks are "asymmetric" in a manner quite common in real estate, venture capital, and privately-held business investments, as well as most major life decisions involving substantial economic uncertainties (such as deciding whether to buy a house or return to school).

Rewards: Rewards are the net distributable earnings from owning and operating the operation, business, or other enterprise for each time period. The reward is dependent on the market conditions, and whether the enterprise is incurring the discretionary (insurance of investment) expense.

The reward matrix, which is sized *SxA* where *S* is the number of states and *A* is the number of actions, is represented in Table 4.

TABLE 4. Reward Matrix for Black Swan Template

	Action "Operate"	Action "Insure"	Action "Sell"
State "sold" *Note: this is an absorbing state, from which the subject cannot emerge.*	0	0	1.00
State "black swan"	-2065.00	-2030.00	-Inf
State "revenue=1190"	-122.50	-87.50	600.00
State "revenue=1330"	-17.50	17.50	600.00
State "revenue=1470"	87.50	122.50	600.00
State "revenue=1610"	192.50	227.50	600.00
State "revenue=1750"	297.50	332.50	600.00

Source: Output from Solution Template included in Rapid Recursive Toolbox version. 1.0

Notes on Forcing Certain Decisions or States. This reward matrix includes two techniques designed to force the solution algorithm to recognize that, in certain situations, the action by a subject person is already decided.

Observe that a small positive number is included for "sell" even when the state is "sold," to keep the subject in that state during the optimization calculations. An entry of 10 cents would probably also work, as it is clearly higher than zero. Also, note the "-inf" number for "sell" when the state is "black swan." (This means "negative infinity.") This is to incorporate into the model a mechanism to prevent a person who owned the operation in good years from selling it just before the cost of a known black swan event were incurred. Because *-inf* is smaller than all countable real numbers, the solution algorithm recognizes that the value-maximizing decision in that state cannot the one with an immediate reward of *-inf*.

Results: Value and Policy: The results from solving the properly-composed value functional equation include the value of the operation to the owner, and the best policy: insure more, continue with current investment, or sell entirely.)

An optional comparison analysis uses these same inputs and calculates the capitalized value using a standard discounted cash flow (DCF) method. Both the recursive model and the standard DCF model incorporate expected future income and rely on the same assumptions about costs and earnings. The DCF model ignores the very small chance of a black swan event.

Interpreting Results: In general, the recursive model recognizes the value that an operation maintains even during a temporary period of poor economic conditions, and does not exaggerate the value during temporary periods of high income. Furthermore, it can take into account asymmetric risks and the risks of a very unlikely, but very damaging, event. It is useful to compare these results with that from a DCF analysis based on the (naive) assumption that current market conditions (and income) will persist with little variation in the future.

B. Results for the Example

The results for the example analysis are shown in Table 5. Note that, in the case of the "black swan" event, the value is negative (as the subject cannot escape losses). However, in other situations the value is positive, and the value-maximizing action can be to sell the operation, or to operate it. In the case of the "sold" state, there is a nominal value that arises from the use of a small number (such as $1) in the reward matrix for the appropriate action in such a case. Although we could manually adjust this, we have left it in the table to provide an example of one of the ways to structure "real option" problems.

TABLE 5. Results for Solution Template "Black Swan"

State	Value	Policy
sold	9.50	Sell
black swan	-1434.12	Operate
1190.0	608.50	Sell
1330.0	613.38	Operate
1470.0	718.38	Operate
1610.0	823.38	Operate
1750.0	928.38	Operate

Source: Output from Solution Template included in Rapid Recursive Toolbox version. 1.0

C. Other Suggested Uses

This example could be adapted for numerous situations such as:

- Problems where one unusual event could have a devastating effect on the business or other operation

- Decisions regarding the purchase of insurance or hedging assets, or making risk-reducing expenditures.
- Problems involving "real options."
- Intellectual property valuations, where there is the possibility that the IP will not be commercialized, or that it may be commercialized but much of the earnings lost to infringement.
- Many common valuation problems where some asymmetries exist in the distribution of possible outcomes, and where market conditions can be directly translated into earnings for the entity being evaluated.

CLASSIC REINVESTMENT PROBLEM

A. Problem Description

Setting: An investor or manager has authority over an operating business that is currently profitable, and can choose to make, or not make, discretionary investments toward the end of each time period. Any such investment reduces the investor's current earnings, but increases the value of the firm in the future.

Purpose and Limitations: This is a classic problem in economics, which is sometimes characterized as an individual choice as the "savings problem." It illustrates the tension between the benefits and costs of a decision in the current time period and the future benefits or costs that result from that decision. The purpose of this model is to illustrate these tensions and to demonstrate how a recursive model resolves them.

As a teaching model, the parameters are not intended to represent an actual operating company. However, this template could be modified to represent an actual firm.

States: The states in this model represent the state of the U.S. economy, which is represented by numbers (which could mean GDP, sales in the industry, or income in the region). Better economy produces better investment earnings.

Actions: The different investment action choices available to the investor. Investor can choose to re-invest some of his or her earnings, or just enjoy the earnings.

Transitions and reward: A higher level of investment will lead to lower current earnings and higher earnings in the future.

Results

The solution results for the example decision problem are shown in Table 6. Note that, as one would expect from both intuition and from economic reasoning, it pays to invest when you have a growing company. (The state variable is a natural indication of size.) However, as is often the case with large, mature companies, at some point simply maintaining a consistent level of investment makes more sense**.**

TABLE 6. Results for Classic Reinvestment Problem

'State'	'Value'	'Policy'
'5'	[1754.00]	[4.00] ("invest big")
'25'	[1971.00]	[4.00]
'50'	[2202.00]	[4.00]
'200'	[2464.00]	[2.00] ("maintain current investment")

Source: Output from Solution Template included in Rapid Recursive Toolbox version. 1.0

MATLAB Note. You might notice that, in some cases, the output from a MATLAB® command will include numbers enclosed by square brackets, as in Table 6. This is a Matlab convention for numerical data in a vector. We have left this in some tables in this *Guide* to give the user some familiarity with how the output could appear in their Matlab session. In some cases, we also left the single quote marks that are used by Matlab to designate a string in the excerpts shown in this *Guide,* for the same reason.

MACHINE REPLACEMENT PROBLEM

A. Problem Description

Subject Company or Person. A company has a large number of assets that are used in the business. These could be computers, engines, machines, structures under heavy use. The company management must decide whether to operate and repair these machines, or to replace them. Once a machine is replaced, it needs less repairs than an older machine. However, replacement is more expensive than repairs.

This is a classic use of dynamic programming methodology, and is often featured as the first or second example in a textbook on optimization. Often, this is a fixed-period problem, in which there is an explicit assumption that at a date certain in the future, all machines are sold and the operation is shut down. In this case, it has been presented as a indefinite-horizon problem.

Inputs. One set of inputs for this problem are based on the actual costs of repairing and replacing a server for a business that served a consulting firm that employed between one dozen and two dozen employees in two offices, and which was used to store files that were necessary for production as well as handle other typical officer server tasks. The business manager provided some of the cost figures (including an estimate of the cost of a breakdown,) while an IT vendor provided estimates of duration of operation before repair and replacement.

During this investigation, it became clear to the author of this particular analysis that much of the information on the reliability and best-time-to-replace of servers came from rules of thumb, rather than hard analysis. In this case, the business manager already had replaced the server before having this analysis completed, and the manager observed that a significant saving could have been achieved by waiting a little longer before replacing it.

FIGURE 5. Possible Inputs for Machine Replacement Problem

```
%   Cost set 2: server for small or medium sized enterprise
    Input.costset = 'Server for SME';
    Input.repairscost = 3200;
    Input.replacecost = 22000;
    Input.breakdowncost = 7500;
    Input.shutdowncost = 7500+2500;
% Baseline revenue and costs
   Input.baserev = 10000;  % this is assumed to be the revenue from operating;
% it is lost when shut down.
% This can be set to an approximate loss of
% output, or to an estimate of income from operations.
    Input.basecost = 500;   % this is assumed to be minimum costs to operate.
```

B. Results

The results of the analysis were presented graphically and with tabular data. The Solution Template makes use of the data that is captured in the Output structure and Input structure to create an interesting graph. This is an illustration of the benefit of having the data in an environment that allows for additional analysis and visualizing.

Table 7 shows the results. Note that, as expected, new machines should not be

TABLE 7. Results for Machine Replacement Problem

'State'	'Value'	'Policy'
'shutdown'	[34309.00]	[1.00] (“Replace”)
'breakdown'	[50722.00]	[2.00] (“Repair”)
'4 yr old'	[59711.00]	[2.00]
'3 yr old'	[61422.00]	[3.00] (“Operate”)
'2 yr old'	[62911.00]	[3.00]
'1 yr old'	[64205.00]	[3.00]
'New'	[65331.00]	[3.00]

replaced or repaired. However, over time it becomes necessary to repair them, or risk a very expensive shutdown. At some point, the machine must be replaced. The Rapid Recursive toolbox makes identifying these decision points easy, by automatically creating the "S-v-p" (State-value-policy) table containing the data shown here.

FIGURE 6. Graphical Output from Machine Replacement Problem

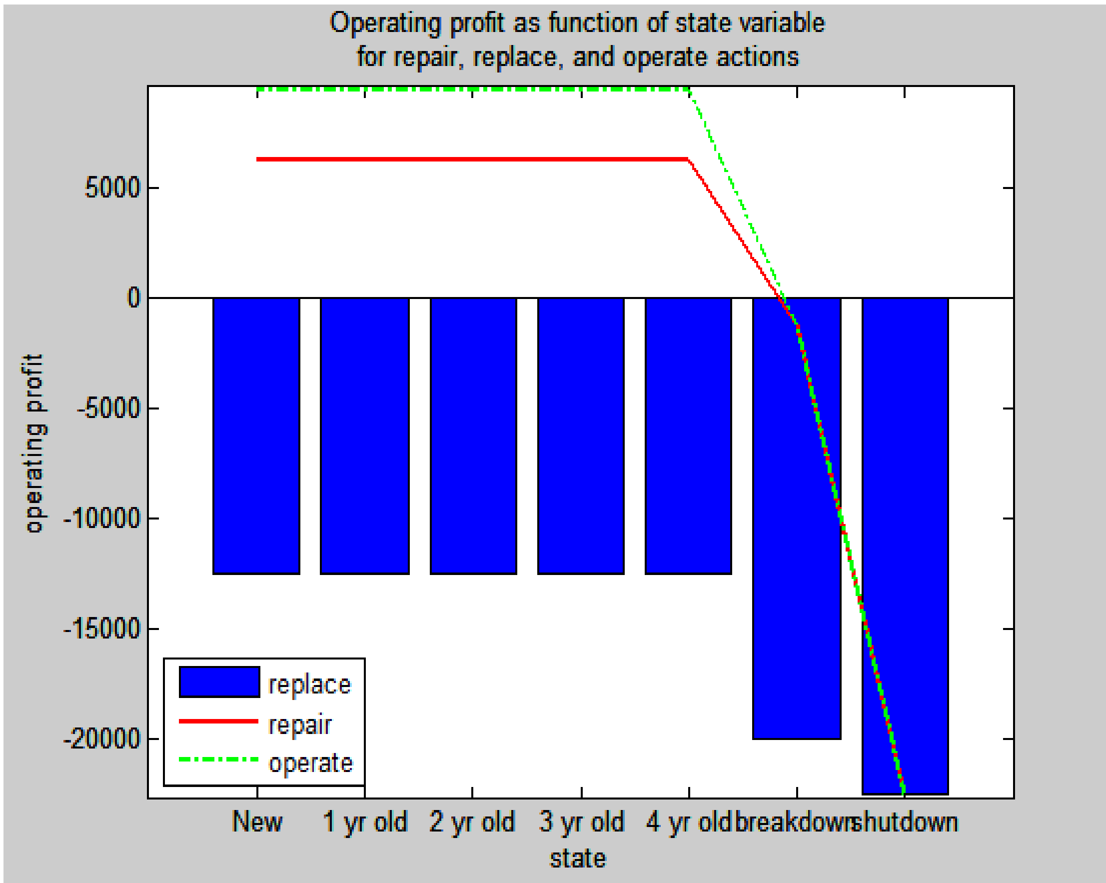

Part IV. Appendices

Appendix A.
Contacting Supported Intelligence LLC

AUTHORS, SCHOLARS, RESEARCHERS

If you are an author with a specific research project with a goal of publishing the results, you can contact us regarding this special assistance. Please include in your inquiry information about the topic and time line of your research, and the expected form of publication.

TRAINING FOR ACCREDITED CONSULTANTS

Supported Intelligence LLC provides training to individuals and organizations that wish to use the Rapid Recursive toolbox in their work. Those completing this training may receive accreditation from SI LLC. Accredited consultants that regularly provide consulting services to their clients using the Rapid Recursive toolbox have access to a licensing agreement with SI LLC that provides enhanced services and a fee schedule that matches the use of the product with the licensing fee.

ACCREDITED CONSULTANTS

For organizations that wish to create, use, and periodically update customized models, accredited consultants may be able to provide assistance on a consulting basis. Please consult the listing of accredited consultants on the SI LLC website.

Please note that accredited consultants are independent businesses. Organizations wishing to retain their services should formalize their arrangement in an agreement that covers such topics as confidentiality of any proprietary data, quality assurance, rights to use models developed during the course of a consulting assignment, quality assurance, work plan, and fees.

GENERAL HELP AND TROUBLESHOOTING

Assistance to licensed users is provided under the terms of the license agreement. Please refer to the *Maintenance and Support Addendum* for your license for full terms and conditions of use. If the Support service time has expired for your original license purchase, you can renew your Support agreement (and receive any updates for the Software) as well.

Many issues are addressed in the *User's Guide* that is bundled with the software, including reference information for all commands. A version of the *User's Guide* is also available on the Supported Intelligence website, which is:

http://www.SupportedIntelligence.com/support

In addition, there is a list of known issues on the Supported Intelligence website.

The Help Desk is available to licensed users during hours that are posted on the website. The help desk email address is: help@SupportedIntelligence.com.

Please note that you may have to provide evidence of a valid license, and should be prepared to carefully describe the problem; as well as the platform, operating system, and version of software you are running.

Planned Improvements; Warranty

Supported Intelligence LLC stands behind its products. We intend to continue to improve them, correct any errors, provide work-arounds to issues that cannot be completely solved, and provide our licensed users with information on best practices. We welcome suggestions on the product, or advice on areas that may or may not be "bugs," but result in difficulties for our users. If you encounter a bug (or something that you just can't figure out) in the Rapid Recursive® Toolbox, please help us improve the software by reporting the issue to: bugs@SupportedIntelligence.com.

However, if you find a defect that we cannot resolve, and are unsatisfied with the product, the License Agreement includes a limited warranty. Please consult the text of the Agreement for this limited warranty.

Help from Accredited Consultants

General assistance does not include the specialized assistance that accredited consultants can provide in evaluating the data necessary to solve a decision problem, or setting up a recursive model with that data. For such special assistance, see the section "Accredited Consultants" on page 2.

Appendix B.
Bibliography and References

Included in this Appendix are an annotated bibliography that illustrates the intellectual history of the recursive models in economics and finance, as well as a set of complete references.

AN ANNOTATED BIBLIOGRAPHY

Control Theory and Mathematics

The pioneering texts on the use of recursive models in economics include Bellman (1957) and Dreyfus (1963). They describe thought experiments involving costs and revenues for theoretical (and extremely simplified) companies. The work of Bellman and Dreyfus brought acclaim to the technique within mathematics and control theory, even as they noted the practical difficulties involved in its use, including the "curse of dimensionality."

Although Bellman is responsible for the modern theory of dynamic programming, and the "Bellman equation" is named in his honor, there were uses of what would now be identified as recursive methods in other fields stretching back centuries. Anderson (in press, 2013) observes that control theory was necessary to allow for the steam engines that powered the industrial revolution in the 1700s to operate without exploding, and that the Apollo moon launch program in the United States in the 1960s relied upon a recursive guidance method (the Kalman filter) as well as the planned use of "mid course corrections."

Economic Applications: Investment Behavior

Perhaps the first notable use of recursive models in Economics arose in Robert C. Merton's 1973 article on the intertemporal capital asset pricing model. Merton describes there a theoretical model of investors choosing between income today, and future income or capital gains.

Earlier articles using the dynamic programming method to, at least theoretically, evaluate the multi-period problem faced by investors include the article by Paul Samuelson on lifetime portfolio selection (1969), and a similar topic in an article (cited by Samuelson) by Jan Mossin (1968). Harry Markowitz, in his brilliant 1959 book on portfolio selection, established what became its dominant paradigm. Markowitz also foreshadowed the later application of dynamic programming in his discussion of multi-period investor problems in chapters 13 and 16.

Specific uses in theoretical models of finance were developed by Avinash Dixit & Robert Pindyck (1994). Dixit & Pindyck showed how the net present value rule failed as a decision criteria in cases where real options were present, and compared multiple approaches (including contingent claims and dynamic programming) with each other.

John Cochrane (2005) surveys a wide variety of valuation methods, and demonstrates how the use of a recursive approach leads to a formulation of a stochastic discount factor.

Mathematics of Recursive Models

The 1989 publication of *Recursive Methods in Economic Dynamics*, by the Nancy Stokey & Robert Lucas (with Edward Prescott) led to the inclusion of recursive models into the many graduate economics programs. Stokey & Lucas proved the existence of solutions for a wide class of problems, and laid the theoretical basis for much wider applications in the future.

Stokey (2009) is a recent and similarly powerful analysis of the underlying dynamics of economic decisions in a wide variety of subfields. Stokey demonstrates how real options, transaction costs, and other factors contribute to creating “zones of inaction” where rational actors wait, rather than exercise management authority, and further demonstrates the power of recursive models to evaluate such situations.

Macroeconomics and Agricultural Economics

The use of recursive techniques in macroeconomics has been developed by Lars Ljungqvist and Thomas J. Sargent (2000, 2004, 2012). They describe the growing use of the method in economics as an example of intellectual “imperialism.”

A number of applications in agricultural economics as well as numerical techniques were described by Mario Miranda & Paul Fackler (2002). Jerome Adda and Russell Cooper (2003) provide an easy-to-understand argument for recursive methods and a number of examples.

Business Valuation and Forensic Economics

The use of recursive models for the valuation of private businesses, including the idea of the firm as a controlled entity, was introduced by Patrick Anderson (2004). Brief discussions of the use of the technique in forensic economics and in the valuation of privately-held businesses in the U.S. also appear in Anderson (2005) and Anderson (2009).

An extensive description of the method, and comparison with other valuation methods, is in Anderson, *Economics of Business Valuation,* (in press, 2013).

Real Options and Failure of Discounted Cash Flow

Following Dixit & Pindyck (1994), a number of authors have developed the theory of real options and outlined the failure of the standard discounted cash flow model. A number of important historical papers on the topic are collected in Schwartz & Trigeorgis (2001). An early examination of various methods of dealing with real options is Vollert (2003). Schulmerich (2010) discusses the under-appreciated topic of interest rate regimes in real options models. Hull (2008) is a standard reference in financial options, and contains a short description of real options.

Smith, Smith, & Bliss (2011) examine the use of standard DCF techniques in the valuation of entrepreneurial companies, and find them wanting.

REFERENCES

Adda, Jerome, and Russell Cooper (2003), *Dynamic Economics: Quantitative Methods and Applications*, MIT Press.

Anderson, Patrick L. (2004). *Business Economics and Finance.* CRC Press.

Anderson, Patrick L. (2009). “Value of Private Businesses in the United States,” *Business Economics,* vol. 44 no. 2.

Anderson, Patrick L. (2005). “New Developments in Business Valuation,” in *Developments in Litigation Economics,* Elsevier.

Anderson, Patrick L. (2013, in press). *Economics of Business Valuation.* Stanford University Press.

Ashraf, Nava, Colin F. Camerer & George Loewenstein (2005). “Adam Smith, Behavioral Economist.” *Journal of Economic Perspectives,* volume 19, Number 3 (Summer 2005); pages 131–145.

Bellman, Richard E. (1957). *Dynamic Programming*. Princeton University Press, Princeton, NJ. Republished 2003: Dover

Brown, Lawrence D. (2000). “An Essay on Statistical Decision Theory,” *Journal of the American Statistical Association*, vol. 95, No. 452 (Dec., 2000), pp. 1277-1281

Cochrane, John H. (2005). *Asset Pricing, revised edition.* Princeton University Press.

Dixit, Avanish, & Robert S. Pindyck (1994). *Investment Under Uncertainty.* Princeton University Press.

Duffie, Darrell (2001). *Dynamic Asset Pricing Theory, 3d ed.* Princeton University Press.

Fisher, Irving (1930). *Theory of Interest.* New York: Macmillan. Reissued by the Library of Economics & Liberty and available online at: http://www.econlib.org/library/YPDBooks/Fisher/fshToI.html.

Hull, John C. (2000 [2002, 2005, 2008]). *Options, Futures, and Other Derivatives, 5th edition.* Prentice Hall.

Kester, W. C. "Today's Options for Tomorrow's Growth." Harvard Business Review (March-April 1984): 153-160.

Ljungqvist, Lars & Thomas Sargent, (2000, 2004, 2012). *Recursive Macroeconomic Theory,* MIT Press.

Merton, Robert C. (1973). "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, 867-887, 1973.

Markowitz, Harry. (1952). "Portfolio Selection," *Journal of Finance* 7(1), 77-99.

Markowitz, Harry (1959). "Portfolio Selection: Efficient Diversification of Investments." Cowles Foundation Monograph no. 16; reissued by Wiley & Sons.

Judd, Howard (1998). *Numerical Methods in Economics.* MIT Press.

Kahneman, Daniel & Amos Tversky. (1979). "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*. 47:2, pp. 263–91.

Miranda, Mario & Paul Fackler, (2002). *Applied Computational Economics and Finance*, MIT Press.

Mossin, J., "Optimal Multiperiod Portfolio Policies," *Journal of Business* 41, 2 (Apr. 1968), 215-229.

Myers, Stewart (1977). "Determinants of corporate borrowing." *Journal of Financial Economics* 5:147–75.

Poole, David & Alan Mackworth (2010). *Artificial Intelligence: Foundations of Computational Agents.* Cambridge University Press.

Puterman, Martin L. (1994). *Markov Decision Processes: Discrete Dynamic Programming.* Wiley.

Samuelson, Paul A. (1969). “Lifetime Portfolio Selection By Dynamic Stochastic Programming,” *The Review of Economics and Statistics,* vol. 51, No. 3 (Aug., 1969), pp. 239-246.

Schwartz, Eduardo, & Trigeorgis, L. (editors.) (2001). *Real Options and Investment Under Uncertainty: Classical Readings and Recent Contributions.* MIT Press.

Schulmerich, Marcus (2010). *Real Options Valuation: The Importance of Interest Rate Modelling in Theory and Practice, 2nd ed.* Springer.

Shreve, Steven (2004), *Stochastic Calculus for Finance, vol. I, vol II.* Springer.

Smith, Adam (1759). *Theory of Moral Sentiments.* Reissued 1976, Oxford University Press. Available online at: http://www.econlib.org/library/Smith/smMS.html.

Smith, Adam (1776). *Wealth of Nations.* Reissued 1904, edited by Edwin Cannan. Available online at: http://www.econlib.org/library/Smith/smWN.html.

Smith, Janet K., Richard L. Smith, and Richard T. Bliss (2011). *Entrepreneurial Finance: Strategy, Valuation, and Deal Structure.* Stanford University Press.

Stokey, Nancy L. (2009). *The Economics of Inaction: Stochastic Control Models with Fixed Costs.* Princeton University Press.

Vollert, Alexander (2003). *A Stochastic Control Framework for Real Options in Strategic Valuation.* Birkhauser.

Wald, Abraham. (1947). *Sequential Analysis*, New York: Wiley

Wald, Abraham. (1950). “Bayes Solution of Sequential Decision Problems,” *Annals of Mathematical Statistics*, 21, 82-89.

Appendix C.
Sample Worksheets

LIST OF WORKSHEETS

The following worksheets are designed to assist the users of the Rapid Recursive® toolbox in setting up a model, and are included in this Appendix:

1. Subject of Model (Worksheet 1: Subject of Decision Problem)
   This worksheet helps identify the proper subject person (company, individual, or organization) and the essential tension between rewards received today, and value expected in the future.
   We recommend that you start with this worksheet, as items identified here become part of the structure of the model, and transfer directly to the other worksheets.
2. States and Rewards in Each State (Worksheet 2: States and Rewards)
   This worksheet helps organize the possible state of affairs, economic or market conditions, or holdings and rights of the subject into a discrete set of "states." These states become part of the structure of the model. This worksheet starts with the identification of a typical or benchmark state. For this benchmark state, the typical reward to the subject (given a typical action) should also be identified here.
   This worksheet also helps clarify how the exercise of a "real option" (such as reinvesting, scrapping, or selling an operation) affects the state, if such a real option is part of the model.
3. Actions, and Effect on State and Reward (Worksheet 3: Actions and Effect on Reward)
   This worksheet helps identify, and then clarify, the possible actions that the subject can take. It also helps organize the effects of these actions on the state and the reward received by the subject.
4. Transition Probabilities
   This worksheet helps organize information about the transition probabilities that become part of the structure of the problem. The sources for information to develop these probabilities includes subjective assessments, objective ("frequenting") data, and other empirical data.
   One of the strengths of Rapid Recursive toolbox models is the ability to natively incorporate asymmetric risks, such as the chances of a "bad" event being different than the chances of a "good" event.

For many decision problems, outlining the problem in a thinking, pencil-and-paper session is often more productive than immediately attempting to use the software. We recommend that users liberally use these worksheets in such thinking sessions when they first approach a decision problem.

USAGE RIGHTS GRANTED TO LICENSED USERS

Licensed users can use, print, and reprint as many of these as they desire for their own private purposes. However, they are copyrighted products of Supported Intelligence LLC, and should not be posted or redistributed outside of the set of licensed users.

WORKSHEET **1. Subject of Decision Problem**

Question	Possible Answer	Notes
How would you describe this problem?	**Title of Decision Problem:**	*A sequential decision problem involves a sequence of events that can be affected by a person's decisions.*
What person (or what organization) is the subject of the decision problem?	**Subject:**	*Clarifying questions:* *What person can execute decisions that affect the future state of affairs?* *What person receives benefits (rewards) that change based on the state of affairs and the relevant decision?*
What is the reward (or penalty) that the subject person receives, dependent on the state of affairs and decisions that are made? *How is the "reward" measured?*	**Type of Reward:**	*Make sure the "reward" is net of all costs, and that it is measured consistently across all states and actions. Consider whether the true measure of the reward is money.*
How does the subject person discount future rewards (and costs) compared to current rewards and costs? *In what time periods does the subject typically recognize the state of affairs, receive rewards, and take actions?*	**Discount Factor:** net discount rate *d:* growth trend *g:* **Time Index:**	*A trend growth rate in the state; a discount, borrowing or equity earnings expectation; or other factors unique to the subject can be the basis for this concept. Focus on the subject's "Euler tension" between rewards today and value in the future.*
What state of affairs or market conditions affect the subject's welfare or reward?	**States:**	*Worksheet 2 will expand on states.*
What actions can the subject person take that affects the current reward, or the future state of affairs?	**Actions:**	*Worksheet 3 will expand on actions.*

WORKSHEET 2. States and Rewards

States	Reward in this State	Notes

Question: What market or economic conditions, or business factors, or holdings or rights, largely determine the ability of the subject person to gain rewards now an in the future? These are the states in a recursive model.

Choose a benchmark state first, and then identify those that are higher and lower in terms of the subject's ability to earn rewards if they follow a typical course of action.

What is the current state, or benchmark against which other states could be compared?

Benchmark State:

What is the typical, current, or other benchmark action in this state?

Benchmark Action (policy):

What is the reward in this benchmark state, with the benchmark action?

Benchmark Reward:

States in which the reward is generally lower:

State __:

State __:

How is the reward lower, compared to the benchmark, in these states?

States in which the reward is generally higher:

State __:

State __:

How is the reward higher, compared to the benchmark, in these states?

Is there an absorbing state that occurs after exercising a real option?

Null State:

If so, what is the action that exercises the option and results in moving to the absorbing state?

Action (Real Option):

WORKSHEET 3. Actions and Effect on Reward

Actions	**Reward with this State-Action Pair**	**Notes**

Question: What actions can the subject take that affect the immediate reward, and the likely future state of affairs, or both? These are the actions in a recursive model, and the choice of such actions given the state of affairs is a policy.

Review the identification of actions, and effect on rewards, in the other worksheets. Explain (or quantify) further here. It is often useful to start with a benchmark state and action pair first, and identify differences from the benchmark when actions or states change.

What is the typical, current, or other benchmark action, in the benchmark state? **Benchmark Action:**	**Reward in Benchmark State, with Benchmark Action:**	*See the other worksheets to ensure they are consistent.*
Actions that would tend to increase the immediate reward: **Action 1:** **Action 2:**	*How is the reward different, in the same states, with these actions?*	
Actions that would tend to decrease the immediate reward, but that are likely to improve the state of affairs and allow higher rewards in the future? **Action 4:** **Action 5:**	*How is the reward different, in the same states, with these actions?*	
Is there a specific real option requiring a specific action to exercise it? **Action (Real Option):**	*What is the consequence, in terms of immediate payoff, of exercising it?* **Reward for Exercising Real Option:**	

WORKSHEET 4. Transitions

State-Action Pair (this time period)	Likely State (next time period)	Notes
Question: What is the likely consequence of being in the current state, and taking a certain action? An intrinsic element of recursive models is the possible movement from one state of affairs to another. It is this movement that generates a large number of possible scenarios that are evaluated when solving the model. *Review the identification of actions, and their effect on likely future states, in the other worksheets. It is often useful to start with a benchmark state and action pair first.*		
What is the benchmark action, in the benchmark state? **Benchmark State and Action:**	**Chances of:** **Remaining in Benchmark State:** **Moving to a "lower" state?** **Moving to a "higher" state?**	*These chances must add up to 100%. You can use subjective and objective information here.*
What is the action that would tend to increase the immediate reward, in the benchmark state? **Benchmark State, Actions 1 & 2:**	**Chances of:** **Remaining in Benchmark State:** **Moving to a "lower" state?** **Moving to a "higher" state?**	*See Worksheet 2.*
What is the action that would tend to decrease the immediate reward, in the benchmark state? **Benchmark State, Actions 4 & 5:**	**Chances of:** **Remaining in Benchmark State:** **Moving to a "lower" state?** **Moving to a "higher" state?**	*See Worksheet 2.*
Is there a specific real option requiring a specific exercise action? **Action (Exercise Real Option):**	*What is the consequence, in terms of moving to another state, of exercising it?* **Likely State if Exercise Real Option:**	*This may result in moving to the null state identified in Worksheet 2.*