Introduction

Suppose you want to teach someone about baseball. It's pretty hard to describe the point of baseball to someone who's never seen a game, although there are better and worse ways of going about it. Those of us of a certain age might remember Bob Newhart's classic skit of an executive at game manufacturer Olympic Games.[1] Newhart picks up the phone and we hear his comments as he tries to make sense of Abner Doubleday's description of his new game. Newhart has puzzled through the fact that the game is played outside and that there are two teams of nine players. The conversation continues:

Newhart speaking to Mr. Doubleday: You got a pitcher and a catcher. They throw this ball back and forth. That's all there is to it?

..... [Mr. Doubleday responds out of our hearing on the phone]

Newhart: All right. A guy from the other side stands between them. With a bat. I see. And he just watches them?

.....

Newhart: Oh, I see. He swings at it?

.....

Newhart: He may or he may not swing at it. Depending on what?

.....

Newhart: If it looked like it were a ball.... Uh, what's a ball, Mr. Doubleday?

.....

Newhart: You've got this plate. Uh-huh.....

1 The skit "Nobody Will Ever Play Baseball", is from the record album *The Button Down Mind of Bob Newhart*, Warner Brothers, May 6, 1960. I transcribed these excerpts.

Things spiral out of control as Newhart tries to understand what happens if the batter swings and hits a ball, and what happens if the ball stays fair: " *... and what's fair, Mr. Doubleday?*" In the end he rejects the game: how could anyone expect people to learn and play such a complicated game? Surely no one will ever play baseball!

Of course when most of us learned about baseball, we watched from the stands or saw games on TV, and with a little explanation here and there we got the basic idea of the game. Probably no one learns about the infield fly rule on his first visit to the ballpark, but most will understand how runs are scored, and what it is to win or lose. Over a period of time you can learn some of the intricacies — how fielders shift depending on the hitter and what pitch is being thrown; or what type of pitch is thrown depending on the hitter and whether or not there are runners on base. What's interesting is that you can learn a huge amount about the game, you can hold your own in discussions about strategy and second-guess the manager, and yet not have a clue about how to actually *play* the game. That is, you can't hit, catch, throw or run.

And when it comes to actually playing, there are again layers within layers. You might be able to hit a fastball, but not a curve. You might be great at a Sunday afternoon pickup game, but wouldn't know how to actually turn a double play, or throw to the right base in the heat of action of a more competitive game.

I suspect that most people who have money to invest have heard an incoherent sales pitch like poor old Abner Doubleday's. Perhaps they've done a bit of watching from the stands. They read newspapers, and maybe follow some of the investment programming on TV. But I'm fairly sure that most people who have started to learn about the investment game in this way do not have a clear idea of what the game is really about, and don't understand some of its very basic rules and parameters. A poor understanding of baseball doesn't really matter, and certainly doesn't have a significant impact on your life. But a poor understanding of investing can have a serious impact on your future, and it shouldn't be surprising to find that investing is quite a bit more complicated than baseball.

So my goal with this book is to elucidate some key aspects of the most important part of investing, what you might call *the long-term investing game*. I will cover some of the background information that you need to know, and some of the rules or parameters you need to understand, in order to monitor and evaluate a long-term investment plan. While things get pretty detailed at times, I am not really trying to teach you all of the "on field" skills you need to play the game seriously. However, there are a lot of hints if you are motivated to build and monitor your own investment strategy.

To begin this process, I want to go back to basics and start by asking what investing really is. Defining it in a very general way might help identify the kinds of things you have to know or believe in order to understand the process. Here is a definition from one of my favourite finance textbooks. In its broadest sense,

> "Investment is the sacrifice of certain present value for (possibly uncertain) future value."[2]

That goes by really quickly, doesn't it, so let's slow it down a little. The *certain present value* is an asset you have today, such as money you have earned, but it could be other things as well. You sacrifice it — give it up, trade or exchange it — for some future benefit. That future benefit or value is the return on your sacrifice or investment. And here's one of the really crucial points: in most cases that future benefit is uncertain. You may or may not actually receive that return or future value.

All sorts of quite ordinary activities can be described as "investing". For example:

- A farmer plants part of his stored wheat, the sacrifice being that he can't sell it now or eat it. Assuming that the crop grows, he will receive many times the amount he planted at harvest. But the crop could fail entirely, produce a small amount due to pests, disease or drought, or at the other end of the scale, produce a bumper crop. What he actually gets is uncertain.

2 *Investments*, 3rd Edition, William F. Sharpe, Prentice-Hall, Inc. Englewood Cliffs, New Jersey 1978, p. 3.

- We often say that someone is *investing in his education*. You might take a night course such as a part-time MBA, investing both money (for tuition and books), as well as time (that you might otherwise have spent at leisure or at a part-time job). The future benefits might be a higher-paying job and a more lucrative and satisfying career over the long run. Of course there is uncertainty, not only as to whether you graduate, but as to the extent that the more lucrative career materializes.
- You can view putting "sweat equity" into a new business venture as an investment, perhaps giving up the income from your current job, in the hope of building a viable business with a larger income. The range of possible outcomes covers the spectrum of abject failure to wild success.

Notice that in all three cases the investment outcome is something attractive and desired, a gain of some kind, but that isn't actually built into the original definition. For example, many forms of institutionalized gambling can be viewed as investments: you place your bet in exchange for an uncertain future return. The fact that the average return is negative due to the known house take might make gambling irrational in some obvious sense, but that doesn't prevent casinos from being full, and lottery tickets being sold like hotcakes.

While those are all valid examples of investing in the broadest sense, this book will focus on a subset of investing that can be called *financial investment*. In a financial investment the investor purchases a security — a GIC, stock, bond or other asset — or deposits money in a bank, in exchange for a more or less uncertain future financial outcome. It is worth noting that I will only be discussing financial investments that are marketable or traded in active regulated markets, although most of the principles discussed here apply to financial investments that are difficult to trade.

All of these examples of investing, including financial investment, share three elements: (1) the sacrifice or outlay of some asset that you now have, that is, the cost or investment; (2) a future benefit or return; and (3) a greater or lesser uncertainty surrounding that future benefit. The basic idea is pretty clear, and in a very general way it describes what happens with an investment. But up to this point I have avoided using a very important term when outlining

parts (2) and (3): the term is *expectation*. This is a term that you need to understand if you want to treat investing as a sensible, rational endeavour.

Let me return to the wheat farmer. Were I, a non-farmer, to plant wheat, I might believe that it would grow and that I would end up with more wheat than I had originally planted. I might *say* that I expect 20 bushels an acre, but with my lack of knowledge and experience that's just a vague hope, perhaps based on what a neighbour has told me.

The experienced farmer is in a very different situation. Perhaps he believes that *on average*, his best field yields 35 bushels an acre. In an exceptionally fine year it might yield as much as 42 bushels per acre, and in a drought year it might yield less than 20 bushels. Of course hail could wipe out the crop completely. He understands the range of possible outcomes of planting (investing), and has an *expectation* or *expected outcome* that is essentially the average of those possible outcomes, with each outcome (intuitively) weighted by the likelihood of its occurrence.

In fact, he might be very poor at developing his expectations or estimating the average outcome. For example, his expectation might be biased to the positive side. By that I mean that other experienced farmers might take the same information and formulate an expectation of 30 rather than 35 bushels per acre from that same field. Or we might find that if we looked back at the 35 years of records kept by him and his father, the historical average yield was just 25 bushels per acre, and that would make us wonder why he expected 35 bushels.

On the other hand, the farmer might have excellent reasons for thinking that future returns will be more like 35 bushels per acre rather than the 25 bushels from his past. He might have instituted new farming practices, such as irrigation, new fertilizers, better equipment, and so on, that justify such expectations.

So you can see that the sense of "expectation" in *expected return on investment* is very different from my *vague hope*. You can also see what I believe is a key point, one that I think many sophisticated

investors do not understand, that a true *expectation* of return is implicitly the weighted average of a range of possible outcomes. This formulation might be very formal and (perhaps misleadingly) precise, or it may be informal and based on intuition and experience.

With this in mind I will refine the concept of investing. There is not much point in discussing investments that are merely based on vague hopes, and there is certainly no point in writing a book about them. The context here will be to think about investment as a considered or rational process:

> Reasoned or rational investing involves (1) the sacrifice or outlay of some asset that you now have, that is, the cost or investment; (2) an expectation of a future benefit or positive return; and (3) an expectation of the level of uncertainty with respect to receiving that future benefit.

Those expectations in (2) and (3) might be biased, naive, based on misinformation, and so on, but they are the result of some kind of considered process — rational in an everyday sense of the word. I have added that the expected return must be positive, for surely no rational person would invest, sacrificing their current resources, without a positive expectation or forecast of return.

Yet many investors seem to be reluctant to make forecasts, or at least to acknowledge that they are forecasters. You might come across one of the amusing shtick lines of certain "value" investors: "We don't forecast — we buy what is cheap today". I've often heard value investors making fun of forecasters and forecasting, and this may be partly because forecasting is hard to do. But if you think about it for two seconds, you will see that the comment that "we don't forecast" is downright silly.

Assuming that such an investor actually has a meaningful definition of what it is to be "cheap today", why should his clients buy or continue to hold those currently cheap assets? Is it because he believes that those assets will be even cheaper (that is, sell at a lower price) in a few years? No, presumably he thinks they will rise in price and provide a good, healthy return, ideally a better return than average. And if you probe and pry a little deeper, you might find that

those investors act as though the cheaper an asset is today by their standard, the higher its future returns will be. They "prefer" cheaper stocks, the cheaper the better. That's a forecast of future returns, although it may not be a very clear and explicit one.

Moreover, the shtick obscures the fact that a good value investor takes into account a view of the distribution of possible future outcomes. Not only does he believe that greater value (as he defines it) implies higher return and conversely, he often suggests that the downside outcomes of cheap (high value) stocks are limited. He is doing exactly what the revised concept of rational financial investing suggests, forming expectations of return and risk. If a value investor is a good and successful investor, it is because he is a good forecaster.

Not only do we have to make forecasts, we have to make quite explicit forecasts. First, we need to make forecasts of return in order to estimate how much our investment will grow, which in turn is necessary for determining how large an investment we must make in order to achieve our goals. This is important, because at any point we have limited resources that we can either consume now, or save and invest for future consumption. Explicit forecasts of expected gains are a crucial part of making that allocation decision.

A second related point is that forecasts of both return and risk are required in order to choose some subset of investments from a range of possible alternatives. On the one hand almost no investors literally try to hold every possible asset, and their beliefs about expected return and risk are the basis for this. On the other hand most investors intuitively understand that holding just one asset can be very risky. I won't discuss the idea of investment choice until Part 4, but I'm confident that all investors who choose assets understand at some level that their return expectations (no matter what they are based on) are a key part of this decision. The part that needs explanation, though I think it is quite intuitive once it is pointed out, is that expected risk is, or should be, an important part of making more effective choices as well. In sum, making explicit forecasts is part of a disciplined approach to investing.

In the following chapters I will be working towards the goal of formulating expectations of long-term return and risk, first for equities, and then bonds. This will be followed by a discussion of how forecasts can be turned into an investment portfolio that best meets the needs of a particular investor. But I will begin by examining what has happened to markets historically.

You should not think that my study of history implies that I believe that the future will mimic the past in any simple or straightforward way. For example, as I write these thoughts in the early months of 2011, which part of the past would you think will repeat? The last ten years ending in December 2010, in which the S&P 500 equity index returned 1.41% annually including dividends? What about the ten years ending in December 1999 in which the S&P 500 returned 18.21% annualized including dividends? Probably neither. Or perhaps the future will average the whole of whatever arbitrary portion of history that I happen to have in my databases. Surely not, but that is often how people come up with estimates for the future.

Yet there is a lot to learn from history. You have likely heard the well-worn aphorism that "History doesn't repeat itself, but it sometimes rhymes".[3] I take this as a succinct expression of a very complicated, but in the end very obvious, idea. Conditions are always changing and evolving, but the underlying principles that govern change or generate outcomes, remain more or less the same. We couldn't function in any part of our lives if we didn't implicitly believe that future conditions are the result of current conditions, transformed through the laws or principles of nature and human behaviour.

From this very general point of view, I see the description and analysis of past market behaviour as having many educational benefits. Bearing in mind that investors are interested in forecasting the returns of their potential investments, I will mention two of them. First, if enough financial history has been gathered, it is very likely that future performance will fall somewhere within the

3 Often incorrectly attributed, it seems, to Mark Twain, whom I have correctly quoted in the Preface.

range of that past performance — the past begins to show *the range of what is possible*. For example, the recent downturn in equity markets was by no means unprecedented: there are several similar examples over the history of the US market, and if you look beyond the US to international markets there are many more. Knowing such facts is one step towards understanding the potential risks you face by holding equities. Moreover, if you had reason to forecast returns that stepped outside the bounds of the past that we know of, isn't that an important thing to acknowledge?

Second, studying history can sometimes help investors *understand the mechanisms that can lead to desirable or undesirable consequences*. I may be able to discover relationships that have held across markets and over time, which in turn may help me forecast some part of what might happen in the future. So keep in mind that in examining market history I am both looking at the range of possible outcomes, as well as searching for possible relationships that might hold over time.

My emphasis on both history and on forecasting is focused on returns to asset classes or broad market indexes, such as returns to Canadian, US or international equities, or Canadian fixed income. One useful way of thinking about this is whether or not it is likely that an asset such as Canadian equities is likely to outperform or underperform its historical average return over some specific period.

I am also going to make the case that most investors have little hope of reliably outperforming those broad market averages through active management or stock selection within those asset classes. So my view is that a strategy that employs a mix of broad asset classes at the lowest possible cost will be the ideal approach for most investors. You will not be surprised to find that the majority of industry practitioners, especially those selling high-cost investment products to retail investors, will reject this approach. My views are more closely aligned to those who oversee large institutional portfolios, although even in that area there will be disagreement. I hope that at least some investors will have the patience to think through the main lines of the argument, so that even if they do not accept my conclusions they are challenged by the reasoning.

In Part 1, I present a comprehensive history and analysis of equity markets. Chapter 1 summarizes 140 years of history of the US equity market, while introducing the concepts of price returns (capital gains) and total returns, along with real or inflation-adjusted returns. In the subsequent three chapters I examine US inflation and introduce the concept of volatility, using the US market as a case study. In Chapters 5 and 6, I extend the perspective to cover 19 international equity markets, applying the concepts that have been introduced, and fine-tuning conclusions based on that broader experience. I add another interesting level of detail to the historical analysis in Chapter 7 by developing a return attribution analysis that will eventually provide the necessary framework for forecasting equity returns. In Chapters 8 and 9, I analyze the effect of currencies on investments. All of this is retrospective.

In Part 2, I develop a framework from which I can begin to make forecasts of future equity returns. By way of a preliminary analysis, I examine the relationship between economic growth and equity returns in Chapters 10 and 11, to see whether forecasting economic growth can help us in any way with our equity forecasts. The results are largely negative, but they provide some insights into how equity markets respond to information. In Chapter 12, I discuss some of the problems with assuming that average past returns provide a good forecast of future returns. I show how the attribution framework developed in Part 1 can be used as a framework for forecasting equity returns in Chapter 13, while details of this framework are examined in Chapters 14 through 17, as each component of return is handled separately. These ideas are summarized in Chapter 18, and a full set of forecasts is developed to illustrate the methodology. I also show how the framework can be used as a sounding board against which forecasts made by others can be evaluated.

In Part 3, fixed income assets are examined in some detail, along with a small set of other investment opportunities. I begin by discussing the concept of diversification in Chapter 19: why investors should want to diversify, what the benefits might be, and how those benefits are gained, if indeed they are. Chapter 20 examines the historical returns and risks to nominal fixed income

assets, along with the impact of inflation. In Chapter 21, I develop an attribution framework for fixed income that is analogous to the framework developed for equities, and use this to clearly show why the returns and risks of bonds are different from those of equities. In Chapter 22, this is turned into a forecasting framework for nominal bonds. Chapter 23 adds the concept of inflation-linked or real return bonds, modifying the fixed income framework both for historical attribution and forecasting. Chapter 24 concludes Part 3 by presenting correlations between equities and nominal fixed income assets. But I also develop a conceptual framework for weighing the benefits of adding additional asset classes to an investor's opportunity set, and give several examples of this.

Finally, in Part 4, I discuss how all of this work can be put together into a comprehensive investment strategy. My belief is that for most investors, a sensible asset mix, implemented by investments in the lowest-cost investment vehicles that are available, is the ideal approach. Up to this point, active management within asset classes has not been examined, and Chapter 25 discusses both the difficulties of adding value through active management within asset classes, as well as the difficulty faced by investors who try to choose specific active funds managed by others. Chapter 26 extends the analysis by showing why, again for most investors, a guarantee of lower costs is more beneficial than holding out the hope for added value from active management. Chapter 27 highlights the importance of asset mix in determining investment returns: *on average* asset mix determines 90% of the variability of investment returns and 100% of the investment returns themselves. In Chapter 28, I develop a framework for portfolio construction based on the historical returns and risks to asset classes, and show how you might think about choosing from the myriad of opportunities available. The natural result is an ongoing *static mix* of assets, that maximizes your expected benefit or utility given your aversion to risk. Chapter 29 extends this analysis by applying my dynamic forecasts, and some of the consequences of using those forecasts are explored. The natural result of this more refined approach is an asset mix that slowly *evolves*

as your forecasts adapt to changing market conditions and other information. I summarize some key issues in Chapter 30.

Historical or retrospective analysis can help investors understand the past and answer the question: *how did we get here?* But investors need to make forecasts, and my systematic approach to forecasting both allows me to make objective forecasts and to justify them. Just as useful is the fact that my framework allows me to analyze the forecasts of others, and show what has to happen for those forecasts to come true. That is, in both these cases it helps answer the question: *how do you get there from here?*