

Mastering Uncertainty in Commodities Trading

Mastering Uncertainty in Commodities Trading

How investors and firms can use trend following
techniques to profit sustainably from trends in financial,
currency and commodity markets

ALEX KRAINER

EQUILIBRIUM
MONACO

First published in Monaco in 2015

Copyright © 2015 by Alex Krainer

ISBN 978-2-9556923-0-1

Material contained in this book may be reproduced with permission from
its author and/or publisher, except for attributed brief quotations

Cover page design, copy editing by Alex Krainer. Content editing by Nico
Alexander, Emily Krainer.

Set in Times New Roman, book title in Imprint MT shadow

*For my parents,
Radivoj and Svetlana
Krainer*

*Look deep into nature, and then
you will understand everything better.*

Albert Einstein

Contents

Introduction.....	1
Acknowledgments	4
Chapter 1: Meeting the markets	6
What's our position?.....	6
Trading on margin.....	8
Futures	9
Chapter 2: Limitations of knowledge	14
A little knowledge can be a dangerous thing	15
Market facts vs. market narratives.....	16
Peak oil and Saudi oil wealth	17
Is it even true?	20
Chapter 3: Economics and the sciences of forecasting	23
Economists and their forecasts.....	25
Almost (but not quite) as good as tossing a coin.....	29
The brick wall of complexity	30
Forecasting hurricanes	32
Chapter 4: The value of expertise	37
Investment experts.....	41
Chapter 5: Prices, time series, and technical analysis	46
Price discovery process.....	46
Statistical/probabilistic analysis of time series	49
Technical analysis	53
Chapter 6: Markets move in trends	62
Trend followers	64
Momentum investing.....	67
Market trends and value investing.....	68
Human psychology: the driver of trends and bubbles	72
Chinese Warrant Bubble	74
Trends in the broader economy	75
Chapter 7: Speculation and human psychology	78
The losing game and its lessons	80
Rogue traders	81
A dangerous mismatch: the human brain at speculation.....	84

The psychology of loss aversion	85
A matter of judgment	87
Deeper mysteries of human psychology	91
Chapter 8: Quantitative modelling	97
Conceptual challenges	98
Model risk	101
Chapter 9: Speculation in the wild	105
Nature's speculation	106
Speculation in the life of predators	108
The anatomy of a hunt	109
Decision-making in predators and speculators	110
Nature's risk management	111
Chapter 10: Building the I-System.....	113
Outlining the solution space	113
Price charts – the external environment	114
Recognizing price trends.....	116
Determining trade entry and exit signals	128
Strategy risk profile	129
Chapter 11: Building the I-System, again.	134
Chapter 12: Learning to fly	142
Formulating robust strategies	145
It all hinges on your confidence	146
Expectancy and the psychology of trading	147
Designing a diversified investment portfolio	148
Chapter 13: So far so good	156
Biosphere Strategic Capital Fund.....	156
Galstar Derivatives Trading	158
Altana Inflation Trends Fund	160
Future applications	167
Chapter 14: Trends and corporations	169
Hedging and profitability among oil and gas producers	170
Market trends and value creation	172
Yes, but we do not wish to speculate... ..	173
Hedging and competitive advantage	175
Organizing to manage uncertainty	176

Chapter 15: Speculation and society178
To be beneficial, speculation must be regulated 180

Chapter 16: Advice for the young at heart182
Formal education..... 183
Informal education 184
Keep a journal 184
Career path..... 185
Mental posture..... 186
Hygiene..... 186
Go systematic 187
Don't trade 188
Cultivate discipline 190
Money isn't everything..... 190
Advice for organizations and nations 191

Preface

In his book, “My Life as a Quant,” Physicist and quantitative analyst Emanuel Derman wrote that, “*in finance there are few beautiful theories and virtually no compelling ones...*”

I wrote this book in the hope of outlining a hypothesis as well as a practical solution to the problem of speculation in financial and commodities markets that are both beautiful and compelling.

My book traces the path of discovery that led me from the position of a humble trading assistant at a family-owned oil trading firm to a hedge fund manager.

In writing this text I have taken pains to make it as readable and as un-technical as I knew how. To that end I’ve dispensed with the politically correct use of “he or she,” when writing in the third person, using instead mostly “he.” This is not out of disrespect to women, quite to the contrary.

While trading and investment speculation are not exclusively male-dominated domains, men *do* tend to be the protagonists of virtually all rogue trader scandals through history. On the other hand, women occupy several spots in the ranks of the world’s best speculators, including at least two fellow trend-followers: the late Liz Cheval and Leda Braga.

Monaco, July 2015.

Introduction

I don't exactly remember the first time I looked at a price chart of something, but I know I was a young teenager – this was in the mid-1980s – and my thinking went something like, “ooh, if I bought *here* and sold *here*, I'd make *that much* per share... times *so many* shares... equals *big bucks*...” To the uninitiated, price fluctuations captured in a snap-shot of time look like a shortcut to wealth, and so they did to me. My fantasy was stirred on by Aaron Russo's brilliant film “Trading Places” with Eddie Murphy and Dan Aykroyd.

What's one of my very favorite comedies was also an excellent portrayal of the opportunity and risks inherent in trading in commodities markets. Although I never really expected that one day I would actually be trading stock shares, gold, wheat, frozen orange juice or treasury bonds, the world of trading seemed irresistible to me. The lure of easy money wasn't the main attraction. For a boy growing up in socialist Croatia¹, the notion of living the lifestyle of “Trading Places” protagonists like Louis Winthorpe III or the Duke brothers was so surreal, it didn't affect my ambitions in any serious way. Rather, I was seduced by the idea of being in the midst of things, having my eyes and ears focused at the pulse of the world markets, being in the know about the economic and political news, about the prices of gold, corn, oil, stocks, interest rates... Not to mention the ego-inflating idea of picking up the phone and shouting *buy* and *sell* orders to some broker half a world away. It was this macho, alpha-male dimension of being a trader that shaped my youthful aspirations.

In the early nineties, I picked up the terrible Croatian translation of George Soros's book “Alchemy of Finance” after reading in the news how Soros made an obscene amount of money in a single day by short-selling the British pound. I slowly and studiously read through the book trying to comprehend what it was about as best as I could, but at the time I lacked the cognitive hooks on which to hang much of what Soros was talking about and a good deal of it went straight over my head. I clearly remember not quite grasping terms like *capital markets* or *returns*. It was all rather abstract and I was too far removed from the world of real people like George Soros or fictional ones like Louis Winthorpe or the Duke brothers.

Fast forward to 1996. I landed a job at a Monaco-based oil trading company called Greenoil as assistant trader, which eventually brought me to real trading – first of physical cargoes of oil and oil derivatives and then to paper derivatives like futures, options and swaps. Among my first

¹ At that time, Croatia was part of Yugoslavia.

impressions was that oil trading was a lousy business to be in. With razor-thin margins and high risk, one bad transaction easily lost you the profits of several profitable ones.

As the years passed, margins only got slimmer, risks more pronounced, and the whole business model of being an independent trader seemed to have a bleak future. But by now I was at the foot of my learning curve and sensing that my career path was taking shape, I tried to absorb everything I could about the oil market, money, finance, trading, economics, asset management, capital markets and risk.

During a trip to the U.S. in May 1997 I spent a day at a Borders bookstore where I loaded up on books about market analysis, trading and risk management. I generally chose geeky, university curriculum-type books except for one New York Times bestseller: Victor Niederhoffer's "The Education of a Speculator." I knew that Victor Niederhoffer traded for George Soros. In February 1997, Business Week had a full page article about him titled, "*Whatever Voodoo He Uses, It Works,*" showing a small graph with Niederhoffer's investment performance with the caption, "*Crazy like a fox*".

In the article, Niederhoffer is quoted stating how, "*By paying attention to the little things, the nitty-gritty, the humdrum things in life, you become a great speculator.*" Naturally, I was intensely interested in his voodoo and took time to study his book carefully. I was so dazzled with the man's charisma, I practically wanted to become Victor Niederhoffer and tried to emulate his style, thinking, and analytics in my job as a market analyst.

After a while, I sensed that I was pushing the limits of my mathematical capabilities, so I persuaded my superiors to provide me a budget to hire a team of more capable mathematicians and computer programmers to work with. On the morning of the 18th November 1997 I went over to the nearby Ramada hotel to meet one of the candidates for my team. As I waited for him in the hotel lobby, I picked up a copy of the Herald Tribune and found the most astonishing article on the front page. The title read, CONTRARIAN GETS CAUGHT FLAT-FOOTED BY MARKET. Below, the sub-title said, FUND MANAGER LOST ALL IN OCTOBER STORM.

I was dumbfounded. Shocked. Flabbergasted. The article was about none other than Victor Niederhoffer: on 27th October 1997, he sustained a total, 100% loss in a single trading day. This story struck me like a ton of bricks. It was a needle prick to the soap bubble of my aspirations. I was eagerly embarking upon a career path that – for all I could predict – might wind through swamps of mediocrity only to lead to a capital disaster at the end. Imagine a lifetime of ambition, effort, and hope crowned at the end with a humiliating defeat? If this could happen to Victor Niederhoffer,

INTRODUCTION

why not to the next guy? Why not to me? Niederhoffer had a degree in economics and statistics from Harvard, a PhD in finance from the University of Chicago, an assistant professorship at Berkeley... He was the world's number one hedge fund manager in 1996...

Who was I to even dream that I could be better than him? How stupid would it be to dedicate your life to a pursuit that could leave you empty and defeated beyond redemption? These doubts never quite left my thoughts, but by this time, I was committed. I was the market analyst-slash-risk manager at an oil trading company, I had asked for a team of quants and programmers, and I needed to get on with it. I interviewed several fellows, but remained uninspired and reluctant to take up the whole challenge of quantitative analysis of markets.

At one point I thought of one of my colleagues from my high school days in Rijeka, Croatia, Gorazd Medić whom I remembered as an exceptionally bright and hard-working guy. I definitely wanted to talk to him about the project and decided to look him up. Since high-school, I hadn't kept in touch with him, and apart from bumping into him randomly about town when visiting in Rijeka, I wasn't aware of his whereabouts, so I was pleasantly surprised to find out through mutual friends that he was in Paris, working on his PhD in applied mathematics. I gave him a call, and for the next three years we worked together trying to tackle the problem of risk and uncertainty, throwing at it every kind of model and theoretical approach we could lay our hands on. Gorazd was flying from Paris to Monaco and back loaded up with books from the university library and we endeavored to read and work through everything that seemed remotely relevant to our subject matter.

As our project unfolded, we enjoyed a deeply stimulating learning process that slowly scrubbed away the discouragement brought by Niederhoffer's failure and everything that it implied. We had no shortage of ideas to explore or concepts to try out, and when we weren't immersed in building models and running simulations, we were discussing the subject matter almost ceaselessly from every imaginable angle.

Slowly, the information, the theories, and near constant brainstorming led us to a handful of insights and an intuition that guided our efforts toward concrete solutions. The path that emerged was unexpected, but not entirely surprising as it seemed to spring out of some old, deep inner wisdom, or perhaps just common sense about the world that only needed to be freed from the blinders drawn by university education and many of the misconceptions it imparted on us about the world, society, its economic systems and also about ourselves as its human participants. The rest of this book is about our progress down this winding path, the insights we gained, misconceptions we shed, and the solutions that have

crystallized along the way. This process has led me to a point of *certainty* that our solution – the model we built – is a valid answer to the problem of speculation. This statement however, requires a rather significant caveat.

As MIT professor Jay W. Forrester expressed it, “*there is no way of proving that a model or law or theory representing the real world is right... There is only an experimental demonstration that such laws are useful for specific, limited purposes.*”² In this sense, my certainty that our model is a useful (value-adding) answer to the specific, limited purpose of speculating in the financial and commodity markets was a subjective experience. This experience took root long before I had any proof that I was on the right track, and for a time I had a hard time reconciling this certainty with my frustrating inability to persuade hardly anyone to my point of view.

I thought that my insights were so crystal clear and so compelling, that they should be obvious. But most people to whom I tried to convey my ideas were unmoved or even dismissive, which left me wondering if I was speaking nonsense that only made sense to me due to some hidden error in my thinking.

Thankfully, this doubt didn’t discourage me and I proceeded on the assumption that I hopefully wasn’t delusional, reconciling other people’s lack of conviction as an experiential gap between them and myself. The insights that seemed so obvious and compelling to me did take a long time and a huge deal of work, study and contemplation to crystallize. Pushing this project forward was going to be a lonely pursuit and I seemed to be the only volunteer.

Over time, I had the great fortune to come across a handful of individuals who understood our solution sufficiently well to risk their own capital in order to put it to a real life trial. At the time of this writing, we have been able to achieve over eight years of continuous track record, which justifies our hopes and gives us the encouragement to confront the developments that lie ahead.

Acknowledgments

First and foremost, my gratitude goes to my family who have shared and endured many of my project’s ups and downs along the way. My erudite brother Boris provided valuable feedback and quality input that improved the quality and readability of this book. I benefited enormously from Dr. Gorazd Medic’s seemingly super-human work ethic and relentless problem solving drive that propelled our efforts through discouraging

² Keough, Mark and Andrew Dorman. “The CEO as organization designer: An interview with Prof. Jay W. Forrester” McKinsey Quarterly, 1992 Number 2, pp. 03–30

INTRODUCTION

points of stuckness time and time again. My deepest gratitude goes to Boris Brec, who put me through a nearly two-year crash course in software engineering and whose precision, meticulousness, and dedication in building our models leave me in awe. He deserves a monument. Angelo Kirigin, my boss during my tenure at Greenoil gave me and my team the support and nearly full freedom of action to explore as we felt inclined, without interfering with the ideas we chose to pursue. Attilio Chiesa, Francesco Baranzelli, and Simon Gould courageously bankrolled our first, unsuccessful foray into the world of hedge funds. I owe a debt of gratitude also to Andrea Pizzorni, Francesco Picasso, Mahmoud El May, Antonello Filosa, Massimo Filosa, Giangiacomo Serena, Emanuele Baronti, Carlo Palumbo, Carole Robino, Francesca Sacconaghi, Enrico Mazzuca and especially Ugo Chimenti, who have provided us generous support and encouragement. I am grateful also to Tiana Pribanić, Dušan Krautsak, Ken and Angela Constable, Kathleen and Eric Nies, Zsolt Lavotha, Carl George, Jan Haraldson, Bryan Brath, Dr. Christopher Culp, Eliot Goodfellow and Thierry Vignal for their help and encouragement along the way. Also, I am profoundly grateful to Lee Robinson for his extraordinary support through some of the most turbulent times even though this has cost me letting him win at squash more times than I care to count. I'm also immensely grateful to (and for) my children Ethan and Jesse without whom this book might have been finished three years sooner but who brought much joy and a deep source of motivation into my life. Finally, my deepest gratitude goes to my parents, who have provided the rare, unjudging support that gave me the luxury to venture forth without fear of failure or harm.

Chapter 1: Meeting the markets

One minute you're up half a million in Soybeans, the next – boom! Your kids don't go to college and they repossessed your Bentley. Are you with me?

Louis Winthorpe III¹

My initiation to global markets came as I sat at our firm's trading rooms – in Geneva or in Monaco – where I overheard countless hours of traders chatting with brokers, bankers, and amongst themselves. At first, I mostly busied myself with the operational side of the business, but aspiring to become a trader myself, I tried my best to keep my eyes and ears open. I reckoned it would take a long time in this job to start *feeling* the pulse of the markets and understanding the game sufficiently well to become a successful speculator.

Fairly early into my tenure it became clear that much of our business did amount to speculation – you bought a cargo of oil, floated it on the sea in a tanker, and tried to find a buyer for it. If you sold for less than your purchase price, you lost money. It was that simple.

On a typical one million barrel cargo, each dollar's difference was a million dollars gained or lost. Apart for squeezing out some value from the financing, shipping, and insurance, we simply had to generate economic value through speculation. For all it mattered, we might have traded just paper barrels. In fact, in order for us to manage risk, we were obliged to resort to paper derivatives including futures, options, and swaps.

What's our position?

One of the first mysteries I encountered in trading was the term "*position*". Again and again, I heard this word tossed around without

¹ Louis Winthorpe III (Dan Aykroyd) impressing upon Billy-Ray Valentine (Eddie Murphy) the enormous potential for profit or loss in trading commodity futures in Aaron Russo's 1983 comedy, "Trading Places"

knowing what it meant. “*What’s our position in the Med*²?” “*Do we want to take a position in heating oil?*” “*We should reduce our positions...*” The word was so commonplace at the office that it took me some courage to ask what it meant. But I promised myself I’d never shrink from asking dumb questions. None of those traders around me were born knowing what they know, so I figured I shouldn’t be embarrassed by my ignorance.

As my friend Lee Robinson rightly says, “I’d much rather *look* dumb than *be* dumb.” Basically, the term *position* means your possession of and risk in some asset. A position can be *long* or *short*. Taking a *long* position or *going long* means buying, or entering a binding commitment to buy an asset at an agreed-upon price. For example, if you bought a house, or entered into a commitment to buy it, you went long that house.

When you are long some asset, you benefit from its price going up. Conversely, taking a *short* position or going short entails entering a binding commitment to sell or deliver an asset that you don’t yet own at an agreed-upon price. This means that you’ll have to arrange to acquire the asset in the future, so you’ll benefit if its price falls. In the house-buying transaction, a builder who took the obligation to build your house went short that house – he’ll have to build it to deliver it to you.

Market speculation boils down to going long if you expect an asset’s price to rise, or going short if you think its price will drop. In our core business, we tended to be naturally long in oil. If we thought its price would fall, we could *hedge* our exposure by going short the paper derivatives like futures, options or swaps. Derivatives allow you to take directional bets on some asset without the need to take possession of any physical merchandise.

For example, if you bought a million barrels of crude oil, you could hedge that position by selling an equivalent quantity of *paper* oil in the Brent Crude Oil or New York Light Crude futures traded on Intercontinental Petroleum Exchange (IPE) or the New York Mercantile Exchange (NYMEX). In both these markets, a standard contract calls for an exchange of 1,000 barrels of crude oil. To fully hedge a million-barrel cargo, you’d sell short 1,000 contracts (1,000 contracts x 1,000 barrels per contract = 1,000,000 barrels). With this hedge in place, whatever money you lost on the physical cargo, you would gain it approximately³ on your short position in oil futures.

The same operation with options would entail buying (going long) *put* options or selling (going short) *call* options. We call an option to sell a *put*

² Trader-ese for **M**editerranean Sea

³ A hedging transaction is seldom perfect. For a variety of reasons, price fluctuations of the oil futures will only approximately correspond with the pricing of a physical cargo.

option, and an option to buy a *call* option. Alternatively, you can use a swap agreement and exchange a quantity of oil at the present price for the same quantity at some future price. Such agreements are normally settled in cash without moving any physical merchandise.

Each category of derivatives – futures, options and over-the-counter (OTC) swaps – have their advantages and drawbacks in any given circumstances, but their essential purpose is similar in that they allow us to gauge our directional exposure to market prices and transfer our risk to a party willing to assume it, be it in the oil market or any other commodity or financial market. In essence, derivatives enable us to speculate on future price fluctuations. For a host of reasons, I am partial to futures; they are the cheapest and simplest instrument of risk transfer. Brokerage commissions on futures are generally very low, getting in and out of positions is quick, and because futures are traded *on margin*, you can take and offset your positions without tying down a lot of cash or needing to resort to bank financing.

Trading on margin

Trading on margin means that to buy or sell a contract for oil, coffee, gold, or S&P500 futures, you need to put up only a small fraction of the contract's value in cash. For example, to buy a 1,000-barrell contract of crude oil which at \$100/barrel is worth \$100,000 you may need as little as \$5,000 – a mere 5% of the contract's full value. That's because by buying a futures contract, you're not really buying the oil; you are only making a commitment to buy it at a specified price. This commitment is easy to cancel by simply selling the contract.

The purpose of the margin requirement is to cover the losses you might sustain on your position. If your losses exceed the margin requirement, the market clearing house will issue a margin call which means that you'll need to add more cash to your margin account. At each futures exchange, margin requirements are set by the clearing house and are normally expressed as a fixed amount of cash per contract. They normally correspond to between 2% and 10% of the value of the underlying contract. Margin requirements can vary depending on the market conditions; when price volatility in some market increases, the clearing house can increase margin requirements quite significantly.

By requiring all participants to keep cash on margin, clearing houses effectively manage the counter-party risk for all participants. Counter-party risk is the main reason I don't like over-the-counter swaps. In principle, there is no reason why you couldn't draft an OTC equivalent of a futures contract between yourself and your neighbor. If you made a lot

of money on the contract, but you discovered that your neighbor was broke, you'd be hit by counter party risk and you couldn't collect your gains. Imagine holding profitable OTC trades with firms like Enron, Lehman Brothers, or Bear Stearns, all considered top quality counterparties before their sudden collapse, and thinking they were as good as money in the bank. They weren't and I believe counter party risk should be avoided whenever possible, even if your counter-party appears bullet-proof (in the financial industry, nobody advertises that they might imminently be going bust).

The above example of taking a position on an asset worth \$100,000 with as little as \$5,000 illustrates the tremendous potential for gain or loss in trading futures. Suppose you put up \$5,000 to buy a single contract of crude oil at \$100/bbl and from that point on, the price of oil increases 10% to \$110/bbl. You would gain \$10,000 (\$10 x 1000 barrels in the contract) earning a 200% return on that trade. All you need to do to make a killing in futures is to make sure you win more than you lose.

But before we get into the fun stuff, let's take a brief detour and take a closer look at the futures markets, how they evolved, and the purpose they serve in the broader economic system.

Futures

Futures markets evolved in the United States from the centralized grains markets organized in the mid-nineteenth century. Before that time, farmers had to transport their crops to populated areas in order to find buyers. After a rich harvest, the large supply of produce invariably drove the prices down, forcing farmers to sell their merchandise cheap, or if the harvest was poor prices would spike sharply up, hurting the consumers. But while consumers always had the alternatives of reducing consumption or using substitute products, the farmers carried the full risk in their crops.

This was partly resolved through the use of forward contracts, which enabled farmers to negotiate months in advance the price they would obtain for their future crop. That way, the farmer could plan his finances and raise credit more easily. In time, the use of forward contracts became widespread and the contracts themselves became negotiable – they could be bought and sold – and traded in centralized exchanges. In 1848, the Chicago Board of Trade (CBOT) was established to facilitate trading between American Midwest farmers and the East Coast merchants of agricultural commodities.

The next stage in the evolution came with the standardization of the contracts with regard to quality, quantity, and time and place of delivery. Such standardized forward agreements came to be known as *futures*

MASTERING UNCERTAINTY IN COMMODITIES TRADING

contracts. The earliest known futures contract in the U.S. was exchanged in 1851 for the forward delivery of 3,000 bushels of corn. Two years later, CBOT created the first standardized futures contract for corn.

With futures, it was no longer necessary for buyers and sellers to carefully read the terms of each contract and weigh how individual clauses might affect its value. The only thing left for buyers and sellers to negotiate was the price. This made the grain markets more attractive to speculators, and more speculators meant greater liquidity and efficiency of the markets, higher transparency of the price discovery process, and more counterparties willing to assume the risk that the producers and industry participants sought to offload. The innovations from the grain markets were gradually adopted in other commodity markets, and with time, more and more futures markets emerged around the world, including coffee, orange juice, butter, live cattle, lumber, gold, silver, palladium, crude oil, and natural gas. The table opposite lists some of the most popular futures contracts and where they trade:

MEETING THE MARKETS

Futures product	Market (futures exchange)	City
Australian Dollar/US Dollar	Chicago Mercantile Exchange	Chicago
Brent crude oil	International Petroleum Exchange	London
British Pound/US Dollar	Chicago Mercantile Exchange	Chicago
German Bund	EUREX	Frankfurt
CAC 40 futures	EURONEXT	Paris
Cocoa	Coffee, Sugar and Cocoa Exchange	New York
Coffee	Coffee, Sugar and Cocoa Exchange	New York
Copper	New York Commodity Exchange	New York
Corn	Chicago Board of Trade	Chicago
Cotton	New York Cotton Exchange	New York
DAX	EUREX	Frankfurt
Euro/US Dollar	Chicago Mercantile Exchange	Chicago
Frozen orange juice	New York Cotton Exchange	New York
Gas oil	International Petroleum Exchange	London
Gasoline	New York Commodity Exchange	New York
Gold	New York Commodity Exchange	New York
Long Gilt	Ldn Int'l Financial Futures Exchange	London
Heating oil	New York Mercantile Exchange	New York
Japanese Yen/US Dollar	Chicago Mercantile Exchange	Chicago
Lean hogs	Chicago Board of Trade	Chicago
Light crude oil	New York Mercantile Exchange	New York
Live cattle	Chicago Board of Trade	Chicago
Nasdaq 100 futures	Chicago Mercantile Exchange	Chicago
Natural gas	New York Mercantile Exchange	New York
Oats	Chicago Board of Trade	Chicago
Palladium	New York Mercantile Exchange	New York
Platinum	New York Mercantile Exchange	New York
Pork bellies	Chicago Board of Trade	Chicago
Rice	Chicago Board of Trade	Chicago
S&P 500 futures	Chicago Mercantile Exchange	Chicago
Silver	New York Commodity Exchange	New York
Soybeans	Chicago Board of Trade	Chicago
Sugar	New York Cotton Exchange	New York
Swiss Franc/US Dollar	Chicago Mercantile Exchange	Chicago
Thirty year Treasury bond	Chicago Board of Trade	Chicago
Ten year Treasury note	Chicago Board of Trade	Chicago
Two year Treasury note	Chicago Board of Trade	Chicago
Wheat	Chicago Board of Trade	Chicago

Natural gas was one of the most recent markets to spawn an efficient futures market, highlighting the cascading changes it brought about.

Opening up of the US Natural Gas Market

The transformation of the \$100 billion Natural Gas industry started when the US Congress passed the Natural Gas Policy Act of 1978, removing the government control of wellhead prices. The next 20 years saw the emergence of a thriving competitive industry. Almost overnight, new gas trading and marketing companies spun-off from pipeline operators. Large energy suppliers such as Duke Energy, Enron, Reliant Energy, El Paso Energy, Dynergy, and others transformed themselves into trading and marketing companies. The rise of these companies was the most significant development in fostering competition and opening the door to the expanded use of natural gas.

However, the development of the natural gas market wasn't smooth at first. During a good part of the transition period (1985 to 1992), gas pricing was obscure and difficult to follow. The gas financial markets were developing, but computer data networks were not yet efficiently operating. Market prices were posted in a few trade periodicals but price discovery was difficult. Traders spent days at the beginning of each month trying to determine fair prices before actually bidding for monthly supplies.

In spite of the difficulties, growth and competition in the physical business put pressure on producers, pipelines, utilities, and trading companies to meet the demands of the industry. Efficiency and productivity have since become the metrics of the business. The development in parallel of the financial markets helped make price discovery more efficient, enabling the rapid growth of trading in futures, options, and over-the-counter derivatives, and improving transparency in natural gas pricing.

Today, the market for natural gas is vastly different to what it was before 1992. The estimated trading volume in natural gas futures is ten to twelve times the amount of physical product consumed, making the market one of the most volatile, most efficient, and most transparent markets in the world. The experience represents a laboratory example of cascading changes following the liberalization of a market. Besides boosting the overall market growth, improving efficiencies, and intensifying competition, deregulation opened new opportunities for risk management.

Initially, risk management was almost nonexistent, and traders' profits were mainly a function of the volatile nature of the market. But the urgent need for risk management was emphasized after 1992. Natural gas storage in the US is relatively small with a maximum of about 13% of winter demand held in storage. In 1992, hurricane Andrew shut down about 5% of U.S. gas production for more than a week, triggering an unforeseen rally in physical as well as future prices of natural gas. As a consequence of this, and possibly other events, many of the natural gas marketing companies reported heavy losses, while two of the early industry leaders, GasMark, and Centran filed for bankruptcy.

Source: NYMEX

MEETING THE MARKETS

Clearly, the advent of transparent, efficient and well-regulated markets should be very beneficial for the overall economy, offering producers and industry a place to offload risk that speculators could assume instead. However, this beneficial role of speculators in any market depends on adequate regulation and control. In commodity futures markets, this primarily requires strict enforcement of maximum position limits in order to prevent large speculators from being able to *corner* a market and unfairly manipulate prices.

Chapter 2: Limitations of knowledge

If we are to understand the workings of the economic systems we must examine the meaning and significance of uncertainty; and to this end some inquiry into the nature and function of knowledge itself is necessary.

Frank H. Knight

Obviously, a man's judgment cannot be better than the information on which he has based it.

Arthur Hays Sulzberger

As our oil trading at Greenoil increasingly depended on speculation, futures trading became more attractive as an alternative to moving the physical cargoes. With futures you can be in and out of trades in the blink of an eye, without the need to negotiate each purchase and sale contract, each charter party¹, bank financing, and insurance policy anew, deal with shipping operations and the extensive and time-consuming paperwork involved in every transaction.

In addition, moving from the physical commodity where trading involved a lot of industry-specific know-how, to trading futures, meant that we could expand our activity to trading other commodities as well. Soon, we experimented with positions in soybeans, coffee, silver, gold, copper and other markets where we perceived opportunity for profit including currencies and stocks.

The main problem with pure price speculation is making sure you gain more on your trades when you get them right than you lose when you're wrong. And at least in our core business – oil trading – I expected that our

¹ *Charter party* is a contract for hire of a ship for transportation of cargoes.

firm's 20+ years' experience and deep industry knowledge would give us the needed edge. I devoted myself in turn to studying the oil market as thoroughly as I knew how. As I saw it, becoming a great speculator was a question of knowledge and experience: learning everything there was to know about your industry, from production, transportation and refining, to distribution and consumption. The richer and more detailed your mental map of this complex system, the better equipped you are to pick out the relevant bits of information and correctly interpret their likely consequences. Or so I thought.

Today, over two decades since my apprenticeship began, I feel no closer to achieving mastery of the oil market – or any other market – than I then was.

A little knowledge can be a dangerous thing

The natural thing to do before making any kind of forward-looking decision is to get the best available information about the subject matter at hand. Market participants' universal need for accurate, up-to-date information spawned a huge industry channelling a constant flow of information through all the suitable media, from newspapers to wireless hand-held devices. The trouble with information is that it only gives you an advantage if other market participants aren't equally well informed. If everyone knows what you know, asset prices may already reflect the new information so you can't profit from it.

This is the idea behind the *Efficient Market Hypothesis*: in an efficient market, asset prices accurately reflect all the known information relevant to that market. Furthermore, future price fluctuations depend on future events, and are therefore presumed to be unpredictable, making it difficult for any one speculator to derive a systematic advantage from the available information. Things may be different if you happen to have regular access to privileged information – if you are a high-level banker or politician.

The western world's ideology holds that our markets are transparent, that the playing field is level, and that risks and rewards are equally available to all participants. In reality the playing field may be quite a bit *more* level for some participants. Researcher Alan Ziobrowski of Georgia State University looked into the stock-trading performance of US Senators in the period from 1993 through 1998. He became intrigued with the subject after reading that three out of four members of the U.S. Senate had investments in companies directly affected by their legislative activity.

In an eight-year collaborative effort with researchers from three other universities, Ziobrowski found that US Senators outperformed the equity markets on average by 12 percentage points per year. That was even better

than corporate insiders who only beat the markets by 5%, not to mention the Senators' typical constituents who on average underperformed by 1.4% or more. As we'll see further on, being able to systematically beat the markets is difficult and unlikely for any group of investors. Doing so by a whopping 12 percentage points is beyond the reach of even the best professional investment managers. Ziobrowski found that the Senators "*had an uncanny ability to pick the right things on the right days.*"

Like elite politicians, elite bankers also enjoy certain privileges not shared by the rest of the world. Thus, large banks like Goldman Sachs and J.P. Morgan frequently report nearly perfect scores on their speculative trading, having positive performance nearly every single trading day. For example in 2010, Goldman Sachs revealed that out of 252 trading days they only lost money on 11 days. Morgan Stanley had similar results. From 2013 through 2016, J.P. Morgan reported a total of two losing days. Its average daily profits from trading were \$72 million in 2013, \$67 million in 2014, \$70 million in 2015 and \$80 million in 2016.² This kind of performance is unattainable for most of the "ordinary" participants who must tackle uncertainty without privileged information or market access.

We all get news on CNBC, Reuters, Bloomberg, and a myriad of other services that all provide the same information to investors who have to make out what the markets are up to and decide how to manage their investments. But there are deeper reasons why information in itself cannot provide us any systematic advantage. First, we don't act on information per se, but on the way we interpret it. Second, much of the information we receive isn't accurate and some of it may not even be true.

Market facts vs. market narratives

During the two decades of my career in commodities trading I have observed time and again how significant price changes shaped the prevailing market narrative. By *narrative* here I mean a shared interpretation of how key causal forces affect market events. Market fundamentals – the available bits of relevant data – constitute the building blocks that mould our understanding of what's going on, but what we do with them depends on how we evaluate their relevance and credibility. This is never a straightforward process, so at any one time we can entertain more than one possible interpretation of market conditions.

As a young oil market analyst in the 1990s, I pretty much expected that available fundamentals data gave us a factual account of the world: that "bullish" information would lead to a rise in oil prices and "bearish"

² Taggart, Adam: "Banks are Evil: It's time to get painfully honest about this." PeakProsperity.com, 17 March 2017.

information would lead to their decline. Thus, an increase in demand for oil should cause oil prices to rise. So would shortfalls or interruptions to its supply. Conversely, falling demand or increasing production should cause prices to fall. Often however, I observed that price action seemed largely in discord with the fundamentals. For example, in the late 1990s, global economic growth was in full swing and the demand for oil was rising. Meanwhile, funding for oil production and refining tightened globally as capital favored investments in information and telecommunications technologies. As a result, demand was expected to progressively outstrip supply, pushing oil prices significantly higher in the future. Contrary to those expectations, oil prices more than halved from around \$24/barrel in the early 1997 to below \$10/barrel in 1999.

Market participants' need to reconcile the supposedly bullish fundamentals with collapsing oil prices gave rise to stories and rumours about massive stocks of unsold oil and vast tank farms around the world, full to the brim. As prices fell toward \$10/barrel, the bearish narrative became entrenched and many traders thought that oil could halve again to \$5/barrel. But stories about huge unsold oil inventories (in effect rumours given credence by the declining price) proved unfounded and after bottoming out in 1999 oil prices tripled to \$35/barrel over the following 20 months even as the world economy slipped into recession and demand for oil contracted.

Again, the market sought to reconcile these contradictions with a new narrative to fit the events. Now we heard about falling production of oil fields around the world, rising production costs, a shortage of refining capacity and growing demand for oil from emerging economies. One of the biggest stories affecting the market was the peak oil hypothesis. Not that this hypothesis was just then formulated catching everyone by surprise: it was originally advanced by Marion King Hubbert in 1956 and subsequently popularized in the 1970s. Its re-emergence in 2005 and 2006 reflected the markets' need to explain the oil prices, which continued to break new all-time record highs.

Peak oil and Saudi oil wealth

Peak oil refers to the point in time when worldwide oil production passes its maximum point, followed by an irreversible decline. According to various interpretations, this may have already happened between 2005 and 2009³. Given the massive relevance of this hypothesis to an oil trading firm, I made a concerted effort to get to the bottom of the issue. I expected

³ In its 2010 International Energy Outlook, the U.S. Energy Information Administration (EIA) proclaimed that oil production from conventional sources probably peaked in 2006.

to unearth the truth of the matter. Instead, I encountered widely diverging views and dissonant information produced by different agencies and research outfits.

In particular, there was a stark contrast between the views espoused by proponents of the peak oil hypothesis and the conventional view of the market held by the industry⁴. Peak oil researchers held that we are entering a period of terminal decline in oil production and that oil prices will get much, much higher in the future. The industry view held that crude oil was very plentiful around the world and that new deposit discoveries and improved drilling technologies would keep the world abundantly supplied at stable prices for decades.

Happy talk about plentiful oil usually invoked Saudi Arabia's vast reserves and production capacity. For years, the kingdom was believed to have some 260 billion barrels of proven oil reserves together with another 200 billion of probable reserves. It had not occurred to me to question these figures until I started to scratch a bit below the surface. The magic of Saudi oil reserves was that they kept constant (or even increased) in spite of the extraction of close to 3 billion barrels each year.

After twenty years of that, you'd think that reserves would decline by 50 or 60 billion barrels. But no: by 2014, Saudi Aramco claimed that they had 790 billion barrels of *oil resources* and expected this figure to hit 900 billion barrels by 2025⁵. This bonanza did not come about from discoveries of giant new deposits⁶ but from the changing definitions of oil and from a subtle shift in terminology.

While most of the press uses the terms *reserves* and *resources* interchangeably, it is very important to discern between the two. Resources comprise oil from *contingent* and *prospective* sources which include quantities that are *potentially recoverable* from accumulations that are as of yet *undiscovered*⁷. Thus, *oil resources* are by definition wide open to exaggeration and wishful thinking. What we have traditionally understood as "reserves," usually represents only a small fraction of resources that can be feasibly developed.

⁴ By industry, I mean the oil corporations, their bankers and a myriad consultancies and analysts.

⁵ Reuters: "Saudi Aramco's Oil Resources to Grow to 900 bn Barrels by 2025." 19 Nov.2014 - <http://gulfbusiness.com/2014/11/saudi-aramcos-oil-resources-grow-900bn-barrels-2025/#.VONRFSztCYM>

⁶ The last great Saudi oil field was discovered in 1967. To date, only smaller deposits have been found and the bulk of them have not yet been developed.

⁷ "Recoverable" doesn't necessarily mean "economically recoverable," which would imply that the value of extracted oil should cover the costs of exploration, drilling, extraction, transportation and a certain return on invested capital.

LIMITATIONS OF KNOWLEDGE

If we revert to the traditional *Proven Reserves Method*⁸, Saudi reserves appear much less abundant. The last audit of Saudi reserves complying with this methodology was done in 1979 and showed that Saudi Arabia had 110 billion barrels of *proven* reserves, another 67 billion barrels of *probable* reserves and 69 billion barrels of *possible* reserves (reserves are classified as proven if there is 90% confidence of them being recoverable with existing technology and under current economic and political conditions; they are probable if there's a 50% confidence of them being recoverable; for possible reserves, there has to be at least 10% confidence of recoverability under existing circumstances).

Given that 100 billion barrels have already been extracted between 1979 and 2015⁹, Saudi Arabia appeared dangerously close to running dry. Work of peak oil researchers like Matthew Simmons, Collin Campbell and Michael Ruppert corroborated this scenario as did the leak of 2007 confidential U.S. Embassy cables from Riyadh published by the Guardian newspaper¹⁰. Such information was never mentioned in the industry publications and only very exceptionally by the mainstream press.

It seemed to me that, between the conflicting figures and narratives, arriving at an objectively correct projection of future trends in the oil industry was quite out of the question. This research dispelled my ill usion that diligent research of supply and demand fundamentals could conceivably lead to more reliable forecasting of the future price of oil or any other asset for that matter.

There was no reason to believe that the information on other industries was any better. Take the example of South Africa's gold reserves. For decades, South Africa had been one of the world's largest producers of gold. According to a revision in 2001, their gold reserves were pegged at 36,000 tons of the precious metal, about 40% of the world's total.

⁸ This methodology was required by the U.S. Securities and Exchange Commission, but was last performed on Saudi Aramco's reserves in 1979. After the control of Saudi Aramco passed from American management to the Saudi Petroleum Ministry no further surveys using this methodology have been conducted.

⁹ According to the U.S. Energy Information Administration figures, Saudi Arabia extracted 99.76 billion barrels from 1980 through 2014. At a reported 10.05 million barrels produced per day in 2015, the total through 2015 rises to 103.4 billion barrels of already extracted oil.

¹⁰ In 2010 Wikileaks released confidential U.S. Embassy cables from Riyadh that were then published by The Guardian newspaper. One of the cables from 2007 recapitulated U.S. Consul General's meeting with Mr. Sadad al-Husseini, Aramco's Executive Vice President for Exploration from 1992 to 2004. According to this cable, Mr. Husseini asserted that at that time, Saudi Arabia had 64 billion barrels of remaining oil reserves and that these reserves would last 14 years (i.e. until 2021), after which Saudi output would enter a period of steady decline that no amount of effort would be able to stop. A different report by Citigroup in 2012 further confirmed the dire situation with Saudi oil reserves concluding that failing to discover major new oil fields, the kingdom was liable to cease exporting oil altogether by 2030.

However, United States Geological Survey subsequently estimated that South Africa only had 6,000 tons worth of feasibly extractable gold reserves left. Later research by Chris Hartnady of the University of Cape Town showed that the country's true reserves were perhaps as low as 3,000 tons.

The discrepancy between 36,000 and 3,000 tons again puts the whole way we obtain such information in doubt. Neatly tabulated figures published in serious looking research reports had a feel of factual truth, yet I couldn't help wondering how those figures came about. In his book, "The Lexus and the Olive Tree", Thomas Friedman explains how he filed temperature reports for Beirut when working there as a correspondent for the New York Times. "I estimated what the temperature was, often by ad hoc polling," writes Friedman. "Gathering the weather report basically involved my shouting down the hall or across the room: 'Hey, Ahmed, how does it feel out there today?' And Ahmed or Sonia or Daoud would shout back, 'Ya'ani, it feels hot.' ... So I would write, 'High 90 degrees.'"¹¹ Friedman's reports were then duly included in UPI worldwide report from Beirut.

Once published in reputable newspapers as the New York Times or the Washington Post, they appeared as facts, black-on-white, but as Friedman confesses, they were merely his own lazy guesstimates. I wondered how much of the information presented in the compelling research reports I occupied myself studying, came from surveys conducted with similar rigour. Once they were cited by respected institutions however, they gained the validity of hard facts, giving us the sense that we could understand what's going on in the world and why. I was beginning to suspect that often we didn't.

Is it even true?

The arithmetic of government statistics (jobs, growth and inflation) is distorted and dishonest almost beyond measure.

Paul Singer

As if shoddy research, questionable surveying and unstated assumptions built into the so called market fundamentals weren't bad enough, it was clear that much of this information was produced by agencies that had a

¹¹ Friedman, Thomas. "The Lexus and the Olive Tree." New York, Anchor Books, 2000.

direct or indirect financial stake in the industries they reported on. This made it hard to dismiss the suspicion that some of the data was subject to deliberate distortion and fabrications.

Paying a small bit of attention to the news and press releases substantiated this suspicion rather abundantly. Here are a few examples: in October 2012, the U.S. Census Bureau reported an unusually sharp fall in the unemployment rate, from 8.1% in August to 7.8% in September of that year. This was a very unexpected bit of good news as it implied that the economy, which was technically in recession at the time, had miraculously powered forward at the fastest rate in nearly thirty years.

As it happened, this information was favourable to President Obama who, at the time, was concluding the re-election campaign for his second term in office. Not only would the information ultimately prove false, but it turned out that the Census Bureau, which published it, was fully aware of this. It transpired that some of the Census Bureau's surveyors fabricated the data by making up household survey results with fictitious people and jobs. The deception apparently escalated at the time of President Obama's re-election campaign¹².

In another example, during the aftermath of the 2008 financial crisis the Federal Reserve Board was seen repeatedly fudging the figures on U.S. household net worth¹³. In the second quarter of 2009, household real estate wealth was reported to be \$18.3 trillion. Later, the figure was revised down by a whopping \$2.1 trillion. Closer scrutiny of the Federal Reserve Board's reports revealed that such revisions happened in every quarter during the crisis period. The repeated pattern of reporting more positive figures first then revising them downward indicated that these weren't innocent errors but intentional distortions.

This enabled the Fed to report encouraging headline figures and thus curb pessimism during a severe recession. Subsequent downward revision would help the next set of quarterly numbers look better. For example, between the second and third quarters of 2009, household net worth staged a jump of \$2.7 trillion, most of which – \$2.3 trillion – was due to the previous downward revision of the second quarter's figures. Without the downward revision, the \$2.7 trillion improvement would look much less rosy at only \$400 billion. Indeed, this pattern appeared less like honest errors and more like the Federal Reserve Board's crisis-management gimmicks. Borrowing from the same playbook, the U.S. Bureau of Labor

¹² Crudele, John. "Census 'faked' 2012 election jobs report." New York Post, 18 Nov. 2013.

¹³ Durden, Tyler. "Charting The Government's Chronic And Flawed Overrepresentation Of Household Net Worth: A \$2.1 Trillion Downward Revision In One Quarter." ZeroHedge, 11 December 2009 - <http://www.zerohedge.com/article/case-governments-chronic-and-flawed-overrepresentation-household-net-worth>

Statistics similarly engaged in the practice of reporting optimistic unemployment numbers first, then revising them later. Between April and October 2010, the BLS underrepresented the unemployment figures on 22 out of 23 consecutive weeks¹⁴ only to revise them upward later, when they no longer had the news headline impact. According to the New York Times, the total revisions of unemployment figures in 2009 showed that 1.36 million more jobs were lost during the year than originally reported.

The list of similar examples is depressingly long and it is hard to escape the impression that a lot of the information presently circulating in the markets is doctored, spun and distorted. All governments, corporations and individuals for that matter, want to appear more credit worthy than they really are. And just as the U.S. government isn't above fabricating the figures, it's safe to assume that most other governments aren't either. Thus, South Africa's gold reserves could be overstated because the country's ability to service its external debt might be severely impaired if it turned out that its gold reserves were in fact 90% lower than the government said they were. This would be even more adverse to banks that have significant exposure to South Africa's debt.

The country's debt is an asset on the bondholders' balance sheets, and unfavorable information could lead to a credit rating downgrade and crippling multi-billion dollar haircuts for the nation's creditors. The same is bound to be true for other countries, banks, and corporations. Contrary to the Western free markets ideology, it does appear that vested interests have their own agendas and when facts get in their way, the vested interest do their utmost to get in the way of facts. Consequently, investors can't assume that the information they obtain fully reflects the objective reality, even where information sources otherwise appear entirely respectable.

¹⁴ Durden, Tyler. "Charting Statistical Fraud at the BLS: 22 Out Of 23 Consecutive Upward Revisions in Initial Jobless Claims". ZeroHedge, 30 September 2010 <http://www.zerohedge.com/article/charting-statistical-fraud-bls-22-out-23-consecutive-upward-revisions-initial-jobless-claims>

Chapter 3: Economics and the sciences of forecasting

The more I studied economic science, the smaller appeared the knowledge which I had of it in proportion to the knowledge I needed; and now, at the end of nearly half a century of almost exclusive study of it, I am conscious of more ignorance of it than I was at the beginning of the study.

Alfred Marshall

Economics is extremely useful as a form of employment for economists.

John Kenneth Galbraith

How information influences our decisions depends on the meaning we attach to it. That meaning depends on our understanding of how things work and on our convictions. When it comes to investing money, our ideas are likely to be shaped by economics, at least in part. Frank Knight explained the purpose of the science of economics as a way “*to work out, on the basis of the general principles of conduct and the fundamental facts of the social situation, the laws which determine the prices of commodities and the direction of the social economic process.*”¹ Economics is a social science, but over the last century or so, economists have increasingly resorted to methods of natural sciences like physics or mathematics.

The shift from fuzzy analyses of human conduct to pursuing more exact scientific methods compelled economists to adopt numerous assumptions about human nature. The effect of these assumptions was that it confined economics research to an unreal world where human conduct resembles the Brownian motion of inanimate particles. Here are some of the explicit or implicit assumptions economists adopted to make human conduct more suitable to exact scientific study:

¹ Knight, Frank. *Risk Uncertainty and Profit*. New York: Hart, Schaffner & Marx, 1921 (p. 71).

MASTERING UNCERTAINTY IN COMMODITIES TRADING

- participants in an economic system are completely rational;
- they are entirely free to act on their inclinations in the process of production, exchange and consumption of goods and services. No constraints are placed by individuals or by the society on members of the community;
- they enjoy perfect clarity as to the long-term and short-term consequences of their actions
- they are entirely motivated by economic factors;
- communities enjoy perfect competition with constant, complete and costless exchange of information between all participants;
- each member of the community acts as an individual and solely on his own behalf with complete disregard of others;
- community members do not collude amongst themselves at the expense of other members or the community as a whole;
- each member continuously produces a complete commodity which is consumed as fast as it is produced;
- each participant endeavors to maximize his or her own utility;
- members in a community do not engage in fraud ...

While such assumptions may be necessary to describe an economic system in mathematical terms, I think that even termites display more individuality and variation in their behavior than do humans as cast by economists. The contrast between the economists' rational individual and the real humans we all know and love is perfectly captured by the so-called *ultimatum game*. In this game, two players are given a small sum of money to divide between themselves. Player A proposes how to divide the sum and player B can either accept or reject the proposal. If the second player accepts, they split the money as agreed and each gains a share of it. If he rejects, both walk away with nothing.

Now, if ultimatum game participants were wholly rational and strictly intent on maximizing their own utility, we should expect player A to propose a split that's grossly in his favor – say, 80% or more. And since player B is also rational, he should accept anything above zero really, because the alternative – getting nothing – hardly maximizes his utility in the situation. But common sense tells us that real people don't behave that way. Even in the experimental setting of the ultimatum game, people tend to observe rules of fairness and the most common proposal is a fifty-fifty split, while proposals where player B gets less than 20% of the money are routinely rejected. Clearly, the players' sense of fairness in dealing with each other trumps their rationality or any utility maximizing impulse.

But since expressing soft concepts like *fairness* mathematically isn't practical, economists prefer to study a termite-like humanity that does not and never did exist. As a fuzzy, social science, economics has offered sufficiently compelling narratives about the affairs of human societies to be accepted as a legitimate science. As such, it has over the centuries mobilized the creative energies of many great minds who made important contributions to our understanding of how the world works. But in its quest for exactness, it has in part become a jumble of superfluous and often misguided intellectual pursuits. To the extent that its objective is to predict future outcomes, it is unlikely to ever succeed. The following cases offer telling examples of this failure.

Economists and their forecasts

Economists can't forecast for a toffee... They have missed every recession in the last four decades. And it isn't just growth that economists can't forecast; it's also inflation, bond yields, unemployment, stock market price targets and pretty much everything else.

James Montier


Twice a year since 1946, the US Federal Reserve Bank of Philadelphia publishes the Livingston Survey which summarizes the forecasts of 31 prominent economists from industry, government and leading academic and financial institutions. These panelists regularly submit predictions about significant economic indicators including Gross Domestic Product, the unemployment rate, interest rates and the S&P 500 stock market index. Only three months before the onset of the year 2000 recession, these forecasters saw no signs of the imminent economic downturn and stock market collapse. Their forecasts of the unemployment rate, GDP growth and the level of S&P 500 index were widely off mark:

Livingston Survey of 31 prominent forecasters, conducted in 2000 by the Federal Reserve Bank of Philadelphia				
Forecast for:	2001		2002	
	Forecast	Actual	Forecast	Actual
Unemployment	4.3%	4.8%	4.5%	5.8%
GDP growth	3.1%	0.8%	3.4%	1.9%
S&P 500 at year end	1490	1145	1639.5	899

Source: Federal Reserve Bank of Philadelphia

The next, 2007-2009 recession and the 2008 market crash caught them equally unawares. The survey released in June 2007, five months before the onset of the recession, stated that *“the panelists think that real GDP will grow 3.0% annually over the next 10 years.”* They also projected that the S&P 500 index, which traded just above 1500 at the time, would rise to 1600 by June 2008 and 1635 by the end of 2008. In fact, by June 2008, the S&P 500 dropped to around 1400. In light of these events, the Livingstoneans duly revised their next batch of forecasts, only this time they got it even wronger: the S&P 500 lost another 700 points, collapsing nearly 50% below the level predicted by these prominent economists.

Exhibit 1: Stock markets always rise in a straight line - Livingston Survey forecasts for the S&P 500 index



The Livingston Survey, published in December and June of each year, makes economic forecasts for the end of the current quarter and future periods, in half-year time lapses. The forecasts are based on surveys of 31 leading economists from industry, academia and financial institutions. The above chart interposes the forecasts of the level of S&P 500 from two surveys (June 2007 and June 2008) against the actual values of the index.

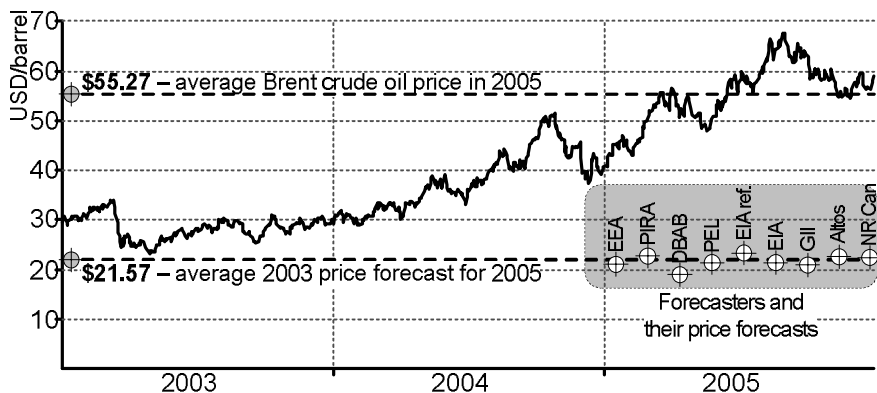
In addition to commissioning surveys, the U.S. Federal Reserve itself retains several hundred economists² who gather economic data and feed it into elaborate economic models that seek to describe how the economy works through complex mathematical algorithms. These impressive troops of learned economists and sophisticated models they built have equally failed at predicting on important occasions.

As hedge fund manager Paul Singer expressed it ever so impolitely in his October 2013 letter to investors, "... the Fed's models and predictions were catastrophically wrong about the financial system, financial institutions and risks in the period leading up to and during the [2008] financial crisis."

The seeming impossibility of successful prediction of economic growth, employment or stock markets is consistent with economists' inability to forecast future commodity price levels as well. The oil market, the world's largest and most closely studied commodity market, offers another example of the failure of forecasting. Every year, the U.S. Energy Information Administration (EIA), the statistical and analytical agency within the U.S. Department of Energy, publishes an exhaustive report titled International Energy Outlook that, amongst other information, provides long-term oil price forecasts. The forecasts are generated by the EIA as well as a group of the industry's leading research institutions.

² According to some reports in 2012, the total number was about 730: 189 worked for the Federal Reserve Board, another 171 at different regional banks; adding in statisticians and support staff – generally also economists, the total arrives at 730. (Source: "How the Federal Reserve Bought the Economics Profession" by Ryan Grim, Huffington Post, 23 October 2009.)

Exhibit 2: A flat outlook – EIA (Energy Information Administration) 2003 oil forecast (for 2005)



Forecasts of the average 2005 oil price, submitted in 2003 to the EIA by the leading researchers: Altos, DBAB (Deutsche Bank Alex Brown), EEA (Energy and Environmental Analysis), EIA (Energy Information Administration), IEA (International Energy Agency), GII (Global Insight), NRCan (Natural Resources Canada), PEL (Petroleum Economics), and PIRA. In 2005, the average oil price rose more than 150% above the average forecast by these institutions.

In 2003, as oil was still trading between \$20 and \$30 per barrel, all the submitted forecasts³ for 2005 were clustered between \$19 and \$24 per barrel. Indifferent to these authoritative predictions, crude oil continued rising with the year’s average vaulting to over \$55 per barrel – 2.5 times higher than the average EIA forecast.


Realizing perhaps the futility of generating specific price forecasts, the EIA subsequently changed the way it projected the evolution of the oil price in the near-term and longer-term future, projecting the likely outcomes in a broadening band between the low and high world oil price. As it extends into the future, the band widens covering as much as \$70 per barrel and more.

With such a broad brush, there’s a better chance of hitting the right answer. Sadly, this may well be the best way of going about predicting future outcomes, short of giving up altogether. Even more sadly for the EIA, even this broad-brush approach put the limitations of forecasting on

³ The forecasts were produced by Altos, DBAB (Deutsche Bank Alex Brown), EEA (Energy and Environmental Analysis), EIA (Energy Information Administration), IEA (International Energy Agency), GII (Global Insight, formed in Oct. 2003 through the merger of Data Resources Inc. and Wharton Econometric Forecasting Associates), NRCan (Natural Resources Canada), PEL (Petroleum Economics), and PIRA. Source: Energy Information Administration “*International Energy Outlook 2003*.”

display: while its 2014 forecast projected the low oil price falling no lower than about \$70 per barrel, within two years – by January 2016 the actual price dipped below \$30!

Exhibit 3: Broad-brush forecasts of North Sea Brent prices, 2014 - 2040



Source: Energy Information Administration International Energy Outlook 2014

Of course, not all forecasts are wrong. Many will turn out correct, but to seriously entertain the notion that economic forecasting adds value in trading, hedging or investment management, the majority of forecasts would need to be right most of the time.

Almost (but not quite) as good as tossing a coin

The only attempt to systematically track market forecasts that I am aware of is a rather admirable study conducted by the Manassas, Virginia based CXO Advisory Group. From 2005 through 2012 CXO tracked over 6,500 forecasts for the U.S. stock market provided by 68 experts including such names as Marc Faber, Jeremy Grantham, Laszlo Birinyi, John Mauldin and Charles Biderman⁴. Their report found that for all graded forecasts, on aggregate only 46.9% were accurate – *almost* as good as tossing a coin! This is quite a concept to ponder: why is it that some of the world’s most learned market analysts with decades of experience, many of them armed

⁴ “Guru Grades” CXO Advisory – www.cxoadvisory.com/gurus/ (last accessed 18 April 2015).

to the teeth with all the information resources and computing horsepower money can buy, can't average better than fifty-fifty?

What is it about economic forecasting that defies the efforts of our smartest economists and institutions? Intuition suggests that the sheer complexity of markets makes it very hard to predict what might happen tomorrow, let alone months or years from today. To be workable, any economic model must be based on a set of assumptions about the future state of the world including population growth, consumer demand, government deficits, inflation rate, geopolitics, wars, revolutions, natural disasters, etc. If any one assumption proves untrue in the future, the model in question will be off the mark. A relatively novel field of mathematics called the Theory of Computation provides an illuminating support for this intuition.

The brick wall of complexity

The science of complexity considers all living systems, from the life of a single cell to human society and its economic systems, as nonequilibrium, or dissipative systems. These are systems that require a constant flow of mass or energy (or both) to sustain the ordered structure. In this sense we can think of economic structures as being maintained in an ordered state by the constant flows of capital, labor, goods and services.

For a whole range of human endeavors, the ability to accurately predict the behavior of nonequilibrium systems like the economy, climate, earthquakes or volcanic eruptions, would be immensely valuable. Nonetheless, we are unlikely to ever achieve any consistent rate of success. This is not due only to inaccuracy of information, limitations of knowledge, or the available computing horsepower, but to the impossibility of modelling complex systems in sufficient detail. This is the hypothesis of the theory of computation, which concerns itself with so called effectively computable algorithms.

The theory of computation studies nonequilibrium systems as if they were computers carrying out algorithms. According to Santa Fe Institute's physicist Stuart Kauffman, the theory shows that "*in most cases by far, there exists no shorter means to predict what an algorithm will do than to simply execute it, observing the succession of actions and states as they unfold.*"⁵ Stated otherwise, an algorithm is its own shortest description. In computer science terminology, it is *incompressible*.

⁵ Kaufmann, Stuart. *At Home in the Universe: the Search for Laws of Self-Organization and Complexity*. Oxford: Oxford University Press - 1995.

Theory of Computation and the Human Brain

Study of the human brain reveals the many difficulties scientists face in analyzing complex systems. In spite of our deep fascination and desire to understand the brain, our knowledge of this complex organ remains relatively modest. To unlock its deeper mysteries, scientists have increasingly turned to computer technology, to try and simulate its various functions and better understand its architecture and functioning. On the 2nd April 2013, U.S. President Barack Obama unveiled "the Brain Initiative," the most ambitious project yet to map the inner workings of the human brain. In Mr. Obama's words, the project's objective was to give scientists "*the tools they need to get a dynamic picture of the brain in action and better understand how we think and how we learn and how we remember.*" That knowledge, said Mr. Obama, "*will be transformative.*"

That may seem like an exciting prospect, but here's a bit of perspective: in August 2013, only a few months after this grand announcement, a team of Japanese and German scientists working at Japan's RIKEN Advanced Institute for Computational Science in Kobe proclaimed that they completed the largest-ever simulation of brain activity using a machine. The simulation was run on Japan's "K" computer built by Fujitsu. K was ranked the world's fastest supercomputer in 2011 and remained among the world's top five in 2013. It consists of 82,944 processors and has a memory capacity equivalent to that of 250,000 personal computers.

The simulation involved 1.73 billion virtual nerve cells connected by 10.4 trillion synapses. That may all sound rather impressive, except, it took K about 40 minutes to complete a simulation of one second of neuronal network activity. Furthermore, while K simulated 1.73 billion neurons, the average human brain is believed to have about 100 billion neurons. In other words, one of the world's fastest supercomputers needed 40 minutes to simulate only a single second of the activity of less than 2% of the average human brain. To be sure, this wasn't an exact replica of a chunk of the actual brain but a rather crude model in which neuronal synapses were connected randomly. By the scientists' own admission, the simulation was only meant to "test the limits of the simulation technology," developed under the project.

The really useful aspect of this simulation was to show just how very far we are from simulating anything resembling the real human brain. Supercomputers will surely keep getting more and more super, but the point in time when they will be able to accurately replicate the functioning of the human brain in real time is very far – possibly infinitely far in the future.

The implication of the theory of computation is that no compact model or theory accurately describing the behavior of a complex system can exist. To be certain about a system's behavior, our only option is to observe it in action. With regards to future outcomes, we are essentially stuck with uncertainty. People frequently counter this idea by raising the subject of weather forecasting and point out that meteorologists have become very good at predicting weather and in particular, hurricanes. While it's true that over the last 100 years or so, climate science has advanced by leaps and bounds in its understanding of climate, hurricanes included, we are quite far from having attained accuracy in prediction. In fact, just like with the markets, there are good reasons to expect that science will never master the complexity of the Earth's atmosphere.

Forecasting hurricanes


The remainder of this chapter is a rather significant detour from our core subject matter, but I've decided to include it here for three key reasons: (1) hurricane forecasting vividly exposes the limitations of science and technology at tackling complex systems; (2) hurricane forecasting offers another real-world corroboration of the theory of computation; and (3) Atlantic hurricanes have great practical importance for oil traders

If a hurricane hits the coastal regions of Texas and Louisiana in the Gulf of Mexico with sufficient strength, it can cause serious disruptions to oil production and logistics. In the process, it may push oil prices higher. For this reason, oil traders tend to pay close attention to the weather patterns in the Atlantic Ocean during the hurricane season and keep current with the most recent forecasts. In fact, weather prediction could be so significant to oil trading that at one point at Greenoil, we seriously considered adding a seasoned climate scientist to our team.

Hurricane forecasts have greatly improved over the last few decades and the 24 and 48 hour forecasts of their *trajectories* tend to be fairly good. *Good* however isn't the same as *accurate*. Furthermore, our ability to predict a hurricane's *intensity* isn't even good. Scientists' ability to predict a hurricane's trajectory has improved thanks to our greatly enhanced understanding of how hurricanes form and why they move the way they do. Here's a very brief description of what we *do* know about hurricanes. They likely originate from the low and mid-level atmospheric winds blowing from the east across the Ethiopian highlands. As they blow over these high mountains, they form vortices that drift westward. When they reach the Atlantic Ocean, moist monsoon winds from the Gulf of Guinea inject humidity into the vortices. If enough humidity meets a sufficiently strong vortex, masses of clouds begin forming rapidly.

Of the total of around a hundred such easterly waves, twenty to thirty will move each season across the Atlantic with the potential to turn into tropical depressions given the right conditions – primarily high humidity and heat. As water temperatures off the west coast of Africa warm to their late-summer peak, moist air begins to rise high into the atmosphere adding more energy to the system and providing humidity for continued cloud formation. For a large storm to form, the following ingredients are required: the ocean's surface temperature must rise above 26 degrees Celsius; the pool of warm water must span at least a few hundred square miles and be at least 60 meters deep; a large layer of warm, very humid air must also be present, extending from the ocean's surface to an altitude of some 5,500 meters. In such conditions, normally in place from June until December and peaking from August through October, the atmospheric disturbances caused by the easterly waves can trigger the formation of large storms. One element however, must be absent for a hurricane to form: winds in the atmosphere over the large pool of warm ocean must be mild or consistent from the ocean surface to an altitude of at least 12,000 meters because strong and varied winds tend to tear cyclones apart. This is the key reason why hurricanes don't form in the southern Atlantic.

Exhibit 4: A (somewhat) predictable path?



Atlantic tropical storm trajectories (1851 – 2005)

Source: Wikimedia commons <https://upload.wikimedia.org/wikipedia/commons/3/31/Atlantic_hurricane_tracks.jpg>

The trajectory of hurricanes is determined in part by Earth's rotation, so they predictably move from East to West rising gradually in latitude along the way. A large area of high atmospheric pressure usually present around the Bermudas (the Bermuda high) causes the hurricanes to swerve northwards and then continue to move toward the North-East following the Gulf currents.

The regularity of these conditions gives hurricanes a somewhat predictable trajectory as the map in Exhibit 4 illustrates. This is only a crude summary of a vast and detailed body of knowledge scientists have accumulated about the formation of hurricanes. In spite of this knowledge and a great deal of technology enabling real time monitoring of storms, we can never be sure about which disturbances will grow into hurricanes and which ones will dissipate or pass with only minor rainstorms.

The structure of hurricanes is largely variable and can change from one day to the next leaving scientists unsure about how exactly all the factors in the atmosphere interact to cause hurricanes to gain or lose strength. Mind numbing advances in computer modelling of the storms have enhanced meteorologists' ability to work out probable outcomes, but not accurate prediction. Weather forecasting models used by meteorologists today fall in three distinct categories: statistical models, dynamical models and hybrid statistical-dynamical models. Statistical models start with the information such as a storm's location and the time of the year to make a prediction based on previous observed storms at the same location and at the same time of the year.

Such models are based on the assumption that over the next 24 or 48 hours, the present storm will behave similarly to the previous ones. Dynamical models analyze all the available information about the storm and its adjacent weather conditions and use the basic laws of fluid dynamics in the Earth's atmosphere to forecast the future development of the storm. Concretely, scientists use six specific mathematical equations to describe the Earth's atmosphere: three hydrodynamic equations rely on Newton's second law of motion to find horizontal and vertical motions of air caused by air pressure differences, gravity, friction, and the Earth's rotation; two thermodynamic equations that calculate changes in temperature caused by water evaporation, vapor condensation and similar occurrences; and one continuity equation that accounts for the volume of air entering or leaving the area.

With the rapid evolution of computer technology, scientists could handle increasing levels of complexity these equations entail. They have advanced by modelling the Earth's atmosphere as a three-dimensional grid consisting of a number of horizontal data points stacked in a number of atmospheric layers. One of the first such models was developed in the

mid-1950s by the US Weather Bureau. Its grid was rather crude, consisting of a single level of the atmosphere at about 18,000 feet (5,500 meters) and data points spaced 248 miles apart. In the 1970s and 1980s, the Hurricane Center developed a much more complex model consisting of ten layers of atmosphere with grid points 37 miles apart. At the time, computers couldn't handle this model's complexity over a large area, so the grid had to move about with the storm, keeping it in the center of an area covering 1,860 miles on each side (it was called the Movable Fine Mesh model).

In the 1990s, the Geophysical Fluid Dynamics Laboratory (GFDL) within the National Oceanic and Atmospheric Administration research center in Princeton, New Jersey developed a model that analyzed data at 18 levels of the atmosphere within three nested grids, the finest of which covered an area of 345 square miles with data points 11.5 miles apart. GFDL's model thus consisted of some 16,200 points receiving data in time steps of 15 seconds. Even at such fine resolution, the model could only represent a hurricane with an idealized vortex structure based upon only a handful of parameters of the real storm (maximum winds and the distance of maximum winds from the storm center).

In parallel, the US Navy developed its own, somewhat less detailed model named Navy Operational Global Atmospheric Prediction System (NOGAPS). NOGAPS consisted of a horizontal grid resolution of about 52 miles with 45 data points in each layer over an area covering 345 square miles. Running on the Cray C90 supercomputer, NOGAPS took about 20 minutes to produce a 24-hour forecast. If the NOGAPS model were run with the finer grid of the GFDL model, it would require about a week's time to produce a 24-hour forecast.

Since the late 1990s, further advances have taken place both in increased model resolution,⁶ speed of computation and processing capacity. Nevertheless, the practical reality of the problem seems to be converging upon the theory (of computation): as computer models of storms have gone from cruder to finer grid resolution, their complexity has increased exponentially requiring exponentially greater computing power to produce timely forecasts. The problem is not only in the resolution of models or in the processing speeds of computers.

An impossible problem also lies in the complex models' sensitivity to input data. Namely, very small differences in the values of initial variables can lead to very large variations in outcomes. The seemingly insurmountable theoretical problem in modelling complex systems was

⁶ In 2007, NOAA's WHRF modelling system was adopted by the National Hurricane Center as one of its main numerical guidance models using grid-points only three kilometers apart.

discovered by MIT's theoretical meteorologist Edward Lorenz. Lorenz developed a relatively straightforward computer model emulating weather. One day in 1961, Lorenz resolved to rerun the results of one particular simulation starting at the half-way point, using the results he had for that particular point in his print-outs. The new simulation quickly started diverging from the original results and soon bore no resemblance to it. The ultimate explanation for this divergence had profound implications for science: while Lorenz's program took its calculations to six decimal places, his print-outs only showed the values to three decimal places.

The minute difference between say, 1.234567 and 1.234 applied in the second simulation led to very large differences in the final results. Lorenz termed this phenomenon, "sensitive dependence on initial conditions." We have every reason to expect that other complex systems will display a similar sensitivity, implying that the problem of accuracy of measurements poses another stumbling block in science's attempt to get to the bottom of such systems. Indeed, accurate prediction will likely remain unattainable in spite of continued advances in all areas of research. As Bob Sheets, the former director of the National Hurricane Center in Miami put it, "*The grid for the computer models does keep getting smaller and smaller, but we're still taking in terms of miles, while the actual weather is taking place at the level of molecules.*"⁷

Both in natural sciences and in economics, our efforts to predict how a complex system will behave are up against a brick wall of complexity. For traders and investment managers, as well as for policymakers, this has sobering implications. It is hard to escape the conclusion that the whole business of economic forecasting amounts to educated guesswork at best. At worst, it could be less than useless.

All the same, well over 60 percent⁸ of investment managers report that they rely most heavily on economic forecasts for their investment decisions. This may help explain another phenomenon that we will explore in the next chapter – the fact that expertise adds little value in investment management and that most investors by far underperform market benchmarks.

⁷ Sheets, Bob and Jack Williams. "Hurricane Watch." New York: Vintage Books - 2001.

⁸ This is according to a 2006 global survey of asset managers and pension funds from 37 countries managing some \$30 trillion in assets co-sponsored by T. Rowe Price Global Investment Services Limited and Citigroup. Questioned about what would drive their investment decisions over the next five years, majority of respondents indicated they would most heavily rely on the "medium term outlook in the bond markets," (67%) and "global/regional economic prospects" (62%).

Chapter 4: The value of expertise

After nearly 50 years in this business I do not know of anybody who has done it successfully and consistently. I don't even know anybody who knows anybody who has done it successfully and consistently.

Jack Bogle on the ability of managers to outperform market indices through market timing.

During a trip to Russia in 1993, William Browder discovered that the whole of the Russian economy – a treasure trove containing some of the world's most abundant reserves of natural gas, oil, coal, iron ore, tin, lead, gold, silver, diamonds, timber, rare earth minerals and arable land – was being privatized at a valuation of \$10 billion, corresponding to one sixth of Wal Mart's market cap at that time. This was a discount of 99% or more on the book value of assets being sold, and the government of President Boris Yeltsin imposed no restrictions on who could purchase the privatization vouchers.

Browder rushed back to Salomon Brothers in London, his employer at the time, to try to convince his bosses and colleagues that they were “giving money away for free in Russia.” But his co-workers showed very little interest. None, writes Browder, “*could divorce themselves from their own narrow mind-set... for weeks I just kept presenting my idea over and over, hoping that by repetition I would eventually get through to someone. ... Instead, I completely ruined my reputation inside Salomon Brothers. No one wanted anything to do with me because I was that 'crazy fuck who wouldn't shut up about Russia.'*”¹ Ultimately Browder set up his own hedge fund, Hermitage Capital Management, which became one of the world's best performing emerging markets fund, gaining 2,679% from 1996 through December 2007.

In 2012, 15 year-old Jack Andraka made an invention and wrote to 200 top doctors and cancer researchers at the National Institute of Health and Johns Hopkins University. He discovered a new test for lung, ovarian and pancreatic cancer which was 168 times faster, 26,000 times cheaper, and

¹ Browder, Bill. “Red Notice.” London: Penguin/Random House – 2015.

over 400 times more sensitive than the standard test used by doctors. He received 199 rejections and only one acceptance.

Browder's and Andraka's stories have two elements in common: a compelling investment opportunity and an astonishingly myopic reaction on the part of supposed experts who should have been interested in such an opportunity. As Browder presented his Russian discovery to investment professionals at Salomon Brothers, virtually all ignored him or peppered him with irrelevant questions about trading spreads on privatization vouchers or advisory fees that could be earned on investment deals.

Jack Andraka's story shows an even more egregious failure of expertise. A fast, sensitive and accurate cancer test costing only \$0.03 vs. nearly \$800 for the standard test is an innovation which, at the very least deserves a second look. The fact that 99.5% of experts failed to recognize this innovation may have had something to do with Andraka's age. However, this in no way absolves the failure of their expertise. An expert should be able to judge a case on its merit and take the correct decision regardless of who presented the case.

Expertise is an important subject in many, if not all domains of human activity, and this includes investing. To negotiate the complexities of our world, we tend to rely on the opinions and judgments of experts for many of the decisions we must take along the way. Expertise gives us a refuge from uncertainty and reassurance when a person knowledgeable in some domain helps us resolve our dilemmas. In many cases, this makes perfect sense. I'd rather not attempt to pilot a jumbo jet or set a broken bone myself – I'm quite happy to rely on the expertise of a trained pilot and a qualified physician in such situations.

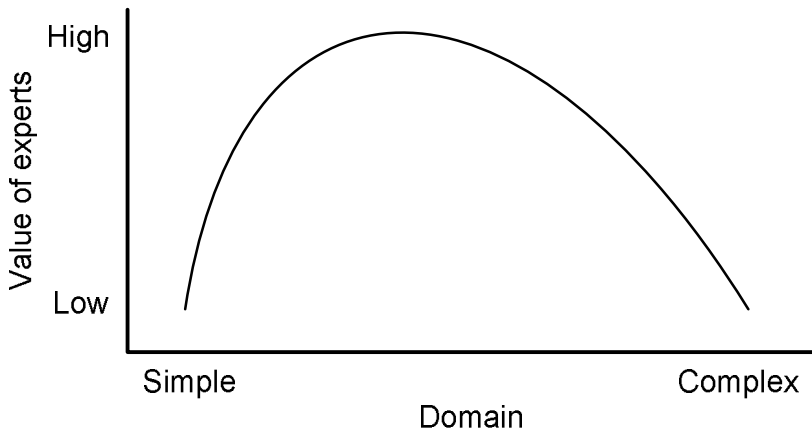
In organized economic life where fragmentation of function and specialization have become pervasive, the reliance on expertise has become indispensable. In "Risk, Uncertainty and Profit," Frank Knight writes that, *"In the field of organization, the knowledge on which what we call responsible control depends is not knowledge of situations and problems and of means for effecting changes, but is knowledge of other men's knowledge of these things. So fundamental to our problem is this fact that ... the problem of judging men's powers of judgment overshadows the problem of judging the fact of the situation to be dealt with."* Indeed, sourcing experts, consulting their know-how, and using their services has become critical to our ability to solve most of our problems in daily life. So accustomed have we become to relying on experts that we don't always discern whether their expertise actually provides the value we seek.

An expert's ability to provide value depends in part on the kind of problem we need to address and the domain in which this problem arises.

THE VALUE OF EXPERTISE

In his 2005 paper titled, “Are you an expert?”² Michael Mauboussin proposes that problem solving domains span a continuum from simple, rule-based systems (like credit scoring and simple medical diagnosis³), to highly complex systems that can’t easily be reduced to a finite set of rules (economic forecasting and stock market investing).

Exhibit 1: Where experts add value



Michael Mauboussin proposes that experts provide the greatest value in moderately complex domains.

Mauboussin suggests that experts tend to add the most value in moderately complex domains, but not in the simplest ones where machines can do the job better, cheaper and faster. Experts are valuable in domains that are too complex for simple algorithms to carry out, but as the complexity of the problem-solving domain increases, the value of expertise begins to diminish and in the most complex domains, expertise is again of little value. I partly disagree with the conception of expertise as Mauboussin presents it. Part of the problem lies in how we recognize expertise. Typically, we recognize trained meteorologists or physicians as experts, but not fishing boat captains or nurses. This dependence on academic credentials and labels as a way to recognize expertise will often prove mistaken. Experts *can* add value even in very complex domains, but this

² Mauboussin, Michael. “Are You an Expert?” Legg Mason Capital Management, 28 Oct. 2005.

³ This strikes me as an unfortunate choice of examples; I’m not sure that very much of medical diagnosing is simple, unless we’re talking about injuries like torn ligaments, broken bones, contusions or burns.

depends on whether they can arrive at their judgment by directly observing the relevant situation or not. Weather prediction falls into the most complex category of problems but expertise can be of considerable value here. An experienced sailor can make fairly good near-term predictions about the weather fronts coming his way simply by observing cloud formations, wind, humidity, and possibly a myriad of other subtle clues like the pain in his joints.

This, in fact, was how the first recorded forecast of an approaching hurricane was made⁴. Namely, while sailing in the West Indies in July 1502, Christopher Columbus watched cirrus clouds moving over the sky from the southeast and an unusually long ocean swell coming from the same direction. He also saw a large number of dolphins leaping from the water at the mouth of the Ozama river just outside the Santo Domingo harbour. In the ten years since his first journey to the West Indies, Columbus learned much about tropical weather. On his second journey to the region in 1495, a similar set of clues preceded a storm that ended up sinking two of the three ships under his command.

Taught by experience, Columbus now expected a large storm and sent a warning to the governor of the Spanish colony asking him to delay the dispatch of thirty ships that were due to sail for Spain and to keep them sheltered until after the storm's passing. The governor was Columbus's rival for the favors of the Spanish crown and to spite Columbus, he disregarded the warning and sent the armada off toward the homeland. Two days later, the storm caught up with the fleet and within hours, 21 of the ships and over 500 sailors were lost. Columbus himself was denied access to the Santo Domingo harbour, but he anchored his four ships in a sheltered bay and all four survived the storm intact.

Another complex domain where expertise can be demonstrably valuable is medical diagnosing. The human body is a complex system and when something is wrong, determining the cause requires a high degree of expertise (unless we're talking about a simple defect like a cut or a sprained ankle). In 1989, Beth Crandall of Klein Associates studied how intensive care nurses make decisions⁵. Crandall interviewed 19 nurses who cared for newborns in distress at the neonatal ward of Miami Valley Hospital in Dayton, Ohio. For example, premature babies are at risk of septic infections that can spread rapidly throughout their bodies and kill them. Recognizing the infection quickly is critical in saving their lives. The nurses' testimonies indicated dozens of cases where this condition was recognized upon a glance and emergency measures were taken,

⁴ Sheets, Bob and Jack Williams. "Hurricane Watch." Vintage Books, New York 2001.

⁵ Breen, Bill. "What's your intuition?" FastCompany issue 38 - September 2000, p 290.

saving the baby's life. When asked how they knew the baby was succumbing to an infection, the nurses invariably replied, "you just know." Upon further investigation, it emerged that the nurses were able to instantly recognize a variety of cues – some of them extremely subtle – that indicated that a baby was in the early stages of an infection. But when Crandall went over the list of cues with specialists in neonatology, she found that half of these cues were not even described in the medical literature at the time. The nurses, which many of us might not recognize as experts, really *just knew*. And by saving the babies' lives, their expertise clearly provided the greatest conceivable value even though they were dealing with a complex problem. The reason they were able to do this – besides their training and experience – was because they were able to observe the babies directly, just as Columbus was able to directly observe the weather cues in his immediate environment.

Had Columbus and the nurses been limited in their reading of the situation to numerical measurements and statistics, they would be looking only at a very rough sketch of the actual conditions, leading perhaps to very different decisions and fewer happy endings. This is the handicap that economists and investment professionals have with regards to their domain of expertise. Market professionals have no way to directly observe the economy or the markets in the same way a seafarer can observe the weather. Instead, they largely depend on the rough – and often distorted – sketches of the economic system through various econometric measures, statistics, prices and the news flow. It should hardly be surprising that their expertise adds little value if any at all. We know in fact, that in most cases it tends to destroy value.

Investment experts

We know that most investment experts tend to destroy value because in the asset management industry, the score can be kept objectively and rather accurately. Our departure point for measuring value in asset management are market indices like the Standard and Poor's 500, Eurostoxx 50, or Nikkei 225. These benchmarks give us a proxy for the aggregate valuation of various global stock markets.

Over the past 100 years or so, stock markets have mostly trended upwards, in spite of periodic crashes and corrections. In this sense, stock markets generated economic value for investors over time, enabling them to benefit even if they invested only passively. Passive investing requires very little expertise and generates results that correspond with the overall performance of the reference stock market. From this base, we might expect that expertise in stock selection and market timing should enable

active investors to outperform⁶ their benchmarks over time. Because this is such an enticing goal, investment management industry deploys staggering resources in trying to attain it. Many of the active fund managers are among the best educated, most experienced and highest paid professionals in the world, with vast information resources and analytical talent at their command. The question is, does all that expertise actually lead to outperformance? The answer is that in the vast majority of cases, it does not.

Over the last few decades, one study after another found that active asset managers have a fairly robust tendency to fall short of their benchmarks. The studies consistently paint a picture that can be summed up as follows: in any given year about two thirds of all active managers underperform their benchmarks. Among the managers who do outperform in any given year, most fail to repeat their success from one year to the next. Measured over longer time periods, roughly 10-15% of the world's investment managers succeed in outperforming their benchmarks year after year.

Percentage of U.S. equity funds outperforming benchmarks

Fund category	Benchmark	1-year	3-year	5-year	10-year
All Dom. Equity Funds	S&P Composite 1500	12.77	23.23	19.18	23.46
All Large-Cap Funds	S&P 500	13.56	23.75	23.75	17.93
All Mid-Cap Funds	S&P MidCap 400	33.77	29.52	14.63	10.29
All Small-Cap Funds	S&P SmallCap 600	27.08	19.60	13.45	12.25
All Multi-Cap Funds	S&P Composite 1500	16.26	23.69	15.98	15.97
Large-Cap Growth Funds	S&P 500 Growth	3.99	29.92	8.50	10.48
Large-Cap Core Funds	S&P 500	20.72	22.68	11.23	15.70
Large-Cap Value Funds	S&P 500 Value	21.41	19.02	13.33	41.24
Mid-Cap Growth Funds	S&P MidCap 400 Growth	43.75	33.52	10.58	8.19
Mid-Cap Core Funds	S&P MidCap 400	41.61	29.57	14.38	11.58
Mid-Cap Value Funds	S&P MidCap 400 Value	26.39	27.27	21.65	14.29
Small-Cap Growth Funds	S&P SmallCap 600 Growth	35.51	28.50	14.05	8.29
Small-Cap Core Funds	S&P SmallCap 600	32.08	17.87	10.75	11.86
Small-Cap Value Funds	S&P SmallCap 600 Value	5.69	10.00	10.32	13.41
Multi-Cap Growth Funds	S&P Com. 1500 Growth	12.15	27.80	15.49	16.67
Multi-Cap Core Funds	S&P Com. 1500	13.40	18.38	14.06	14.80
Multi-Cap Value Funds	S&P Com. 1500 Value	30.58	31.65	17.53	20.36
Real Estate Funds	S&P US RI. Est. Inv't. Trust	19.86	13.77	8.51	21.92
Average – all fund categories		22.81	23.82	14.30	16.04

Source: S&P Dow Jones Indices LLC, CRSP. Data as of 31 Dec. 2014

⁶ In investment management-speak, this added value is referred to as alpha.

The largest of these studies is produced by S&P Dow Jones Indices research which has since 2003 produced a semi-annual report titled “SPIVA U.S. Scorecard.” The table above summarizes the results from the 2014 year-end report⁷. Broadly, these figures are in line with other similar studies done in the past, but while they give us the idea about how few active managers outperform their benchmarks, they omit an important insight: the distribution of managers’ performance, or how many managers outperformed or underperformed and by how much.

This insight is provided in a study by the investment management firm Daniels and Aldredge which quantified the performance of 658 global equity funds over a ten-year period and compared it to an index of securities from all developed and emerging markets⁸. Their findings are summarized in exhibit 2, below. Besides the fact that only 9% of all funds outperformed the QGS Index, two other significant points stand out from Daniels and Allredge data.

First, the performance range spans more than 14 percentage points below the benchmark, but no higher than 6% above it. This suggests that managers’ tendency to underperform is much greater than their ability to outperform. Second, the distribution has a fat tail, but only on the left side. This tells us that while there is some likelihood of extremely poor performance (10 managers falling short by more than 14%), achieving high positive performance is very limited (no managers outperforming by more than 6%).

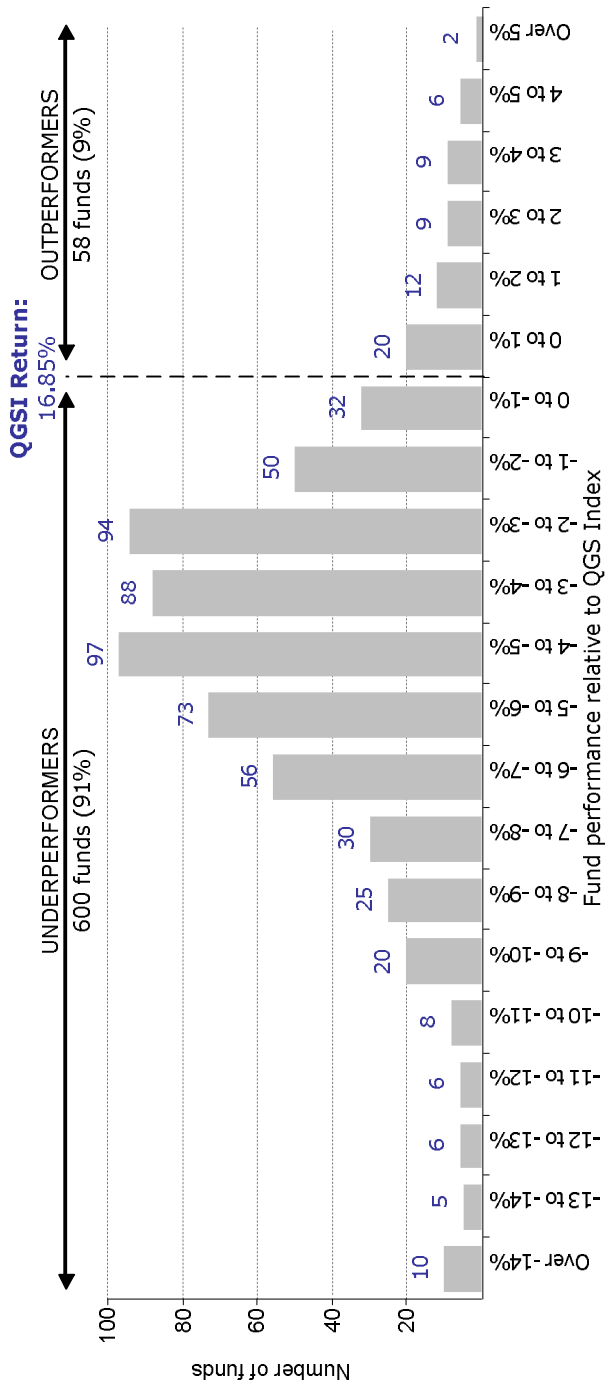
This leads us to the conclusion that consistent outperformance is a difficult and unlikely achievement. But we shouldn’t overlook another significant result of these studies: the fact that a small minority of managers *do* manage to outperform market benchmarks. This would suggest that expertise *can* add value in investment management and the question is, how?

What is it that enables certain managers to do that? I have had the privilege to meet a good number of high performers among hedge fund managers, and there is no question that each one of them has authoritative knowledge about the economy and the financial markets. But this aspect of expertise is not necessarily their edge. Many professionals from among economists, market analysts and even the underperforming asset managers possess impressive expertise in their domains.

⁷ <http://www.spindices.com/documents/spiva/spiva-us-year-end-2014.pdf>

⁸ Malkiel, Burton and J.P. Mei, “Global Bargain Hunting” New York: Touchstone, 1999

Exhibit 2: Expertise destroys value – most active managers underperform



Daniels and Alldredge compared the performance of 658 global stock funds to the Quantindex index, which includes securities from all developed and developing nations. Over a 10-year period (1987-1997), 91% of funds failed to match the index, many falling short by a very wide margin.

THE VALUE OF EXPERTISE

Two key attributes determine the success of high-performing asset managers: their strategy and their risk management discipline. First, all investing is based on clear, strategic thinking in response to some perceived opportunity in the financial markets. More important still is that all investing is subject to rigorous risk management discipline.

A manager worth his salt is always keenly aware that his judgment about the markets could be wrong, and doesn't gamble on being right. For each investment transaction, the potential for gain or loss is determined in advance and losses are not allowed to escalate past a certain predetermined level. On that foundation, investment trading is treated as a long-term process consisting of a patient, disciplined and measured pursuit of investment opportunities sustained over time. In contrast to the constant search for the *trade of the century*, asset management is a marathon, not a fashion show.

Chapter 5: Prices, time series, and technical analysis

*To invent an airplane is nothing. To build one is something.
But to fly is everything.*

Otto Lilienthal

Early into my trading apprenticeship I had to make peace with a few disconcerting insights. First of all, I realized that I would probably never know enough about markets to have much certainty about anything. Second, it was clear to me that I would always have to treat the available information with a good deal of scepticism. And third, I felt that reliable predictions about market prices would remain an unattainable fantasy, and however much effort and diligence I put into this pursuit, I could never overcome the uncertainty about what lies in the future.

This was a rather grim realization for a young man looking to make a living as a speculator. Thankfully, there was one specific kind of market information that was accurate, unambiguous, and almost instantly available: asset prices themselves. Security prices and the data series describing their fluctuations over time offered us important ways to understand markets. By “understand,” I do not mean the kind of understanding that forms opinions or cocktail party discussions about markets, but the understanding that would enable us to make decisions with certain confidence and positive expectancy¹ for speculative gain.

Price discovery process

The concept of price is different in capital markets from what it is in consumer markets. The first kind of price is fluid, the second solid. In everyday life, the price of something is what the seller demands and the buyer pays. If the buyer thinks the price is too high, perhaps he can bargain, or he can shop for an alternative product or a seller with a better price. In organized financial markets, the price of an asset constantly fluctuates as a result of the so-called price discovery process.

¹ In this sense, expectancy is simply an answer to the question of what happens if we continue doing something. Thus, in my mind, a visit to a gambling casino has a negative expectancy – there, the house usually wins, and gamblers usually lose.

This process is driven by an ongoing interaction between numerous buyers and sellers. Buyers come to the market with bidding prices, and sellers with their offering, or asking prices. When any buyer’s bid matches any seller’s asking price, the transaction may take place and the settlement price is recorded along with the number of securities exchanged. The process continues with other bids and offers (or asks) throughout the trading session. At any particular moment in time, a price quotation for a financial product might look like this:

Security	Bid	Ask	Last	Open	High	Low	Chg	Vol.
CME Dec '13EUR futures	1.3074	1.3077	1.3074	1.3004	1.3107	1.2986	+0.007	125000


This is what is called a “level 1” price quote. What we see here are the trading session’s highest, lowest and opening prices as well as the last transacted price and the closest matched bid and ask prices. But there may be many other traders in the market wishing to transact different size trades with bid and ask prices further away from the current price. “Level 2” price quotes provide a deeper picture of the market. Here’s a basic illustration of a level 2 quote:

CME Dec. 2013 EUR futures			
Bid		Ask	
Price	Size	Price	Size
1.3074	3	1.3077	11
1.3073	7	1.3078	7
1.3070	14	1.3080	41
1.3069	9	1.3081	11
1.3065	20	1.3083	15
...

As buyers’ bids match up with sellers’ asking prices, trades are continually transacted with prices fluctuating throughout the trading session. Each trading session is marked by an opening price (the price at which the first transaction took place), the session’s high and low prices, and the last or closing price of the trading session. The volume of trading is also recorded as well as open interest², in the case of futures markets. Each *open-high-low-close* set of prices can be graphically represented by price bars, as illustrated in exhibit 1.

² In futures trading, when a buyer and a seller enter into a transaction, they may open a new contract. This contract remains outstanding or open until it is settled. Open interest in any futures market denotes the total number of such outstanding contracts (or options).


Exhibit 1: Anatomy of a price time-series



A series of bars makes up an asset's historical price chart. Volume and open interest figures are normally plotted in a sub-chart on a separate scale. Price bars can also represent weekly or monthly price ranges, as well as intra-day periods. Thus, with sufficiently granular data, we can construct time series made up of weekly, 60-minute, 5-minute, or any period price bars.

Besides price bars, charts can also be drawn using line-on-close or candlestick charts. Line-on-close removes the “noise” of intra-day price fluctuations and shows a chart plotted only through closing prices. Candlesticks convey essentially the same information as price bars, but make a visual distinction between the “up close” and “down close” days: when the closing price is higher than the opening price the candlestick body is left unfilled, or is colored green; when closing price is lower than opening price the candlestick body is filled solid, or colored red.

Exhibit 2: Three main types of price charts



A less common approach is using point-and-figure charting. Point-and-figure charts are constructed by plotting vertical columns consisting of “X” and “O” symbols where X denotes a price increase and O a price decrease over the period in question – usually daily or weekly. Accordingly, a column of Os implies a possible down-trend, while a column of Xs an uptrend. The peculiarity of point-and-figure charts is that they do not have a linear time-axis and thus focus only on asset price changes. Many of the early trend followers based their trading strategies on point-and-figure charting which apparently gave excellent results for a while. In more recent times however, trading systems based on point-and-figure charts haven’t been as effective.

Statistical/probabilistic analysis of time series

So far as the laws of mathematics refer to reality, they are not certain. And so far as they are certain, they do not refer to reality.

Albert Einstein

Time series consisting of price data, transaction volumes and open interest gave rise to two general approaches in analyzing markets: statistical/probabilistic analysis and technical or chart analysis. Statistical or quantitative analysis of time series usually involves detecting patterns and correlations in the observed price data in order to produce probabilistic predictions about future events. The idea is to allow traders and investment managers to take risks with known odds of favorable outcomes. Here are two basic examples:

- “After two price changes in the same direction, the probability of a third change in the same direction as the second is greater than after two changes in opposite directions... Odds in favor of reversals versus continuations in transaction data normally run at five-to-one.”
- “When the market is down on Friday, chances are three to one that Monday will also decline.”

These regularities in the US stock market were described by hedge fund manager Victor Niederhoffer in his 1997 bestseller, *“The Education of a Speculator”*. Niederhoffer built a notable success as a hedge fund manager by exploiting similar patterns in market behavior. His book was one of the very first texts on trading and speculation I read and I found his approach very intriguing at first. In effect, Niederhoffer treated the markets like they were gambling casinos where each transaction had knowable odds of a favorable outcome. He used statistical analysis of markets to identify trades with high probability of success. By making many such transactions, his trading was expected to steadily generate positive investment returns.

For a time, Niederhoffer’s impressive performance seemed to support this speculation strategy. In February 1997, *Business Week* had a full page article about him titled, *“Whatever Voodoo He Uses, It Works.”* The article included a small bar chart titled *“Crazy like a fox,”* showing Niederhoffer’s investment performance. His returns were about 35% in 1991, over 50% in 1994, nearly 20% in 1995 and over 40% in 1996³, the year in which he received an award as the world’s best hedge fund manager. I took up emulating Niederhoffer’s style with much enthusiasm and started producing my own analyses of energy markets, various currency pairs and interest rates.

After some time, I felt I was reaching the limits of my proficiency in mathematics, so I persuaded my superiors to give me a budget to hire a team of more capable mathematicians and computer programmers to work with. A few months later I had the good fortune to bring to our team my high school friend, Gorazd Medić. At the time, he was working on his PhD in the field of dynamical systems, which was a perfect match for research in quantitative analysis of time series. Gorazd’s intellectual vigour and relentless problem-solving drive proved to be a huge reinforcement to our efforts. For nearly two years between 1997 and 1999 we subjected our price data to just about every known model of data series

³ These are conservative approximations based on the bar-chart exhibit included in the *Business Week* article

analysis, including spectrum analyses, auto-correlations, Box-Jenkins method, ARIMA⁴, and fractal analyses. In spite of our earnest effort to come up with a successful approach to market speculation, the deeper we probed these quantitative/probabilistic models, the less comfortable with them I got. Namely, all these approaches are based on the implicit assumption that the probability of something happening in the future is knowable and quantifiable. That assumption is perhaps the result of a critical, but seldom recognized misconception which holds that correlation implies causation. Consider a simple exercise in logic.

Suppose that our price data shows that an event **B** follows an event **A** in 75 out of 100 observations. This might lead us to the conclusion that there is a 75% probability that **B** will follow the next occurrence of **A**. But that conclusion would probably be wrong. To illustrate, suppose we toss a coin 1000 times and mark the results, **H** for heads and **T** for tails, obtaining a string of 1000 characters like this:

...HTTHTHTHHHTHHHTTTTHHTHHHTHHHHHTTHTTTHHHTT
 HTHTHHHTHHHTHTTTHHTTTHTHTHTHTHHHTHHHTTHT
 HHTHTHTHHHTHTTTTHHTHHHTHTTTHTHTHTHHHTHH
 TTTTHHTTHTHHHTHTHHHHHTTTTHTHTHTTTTHTTHTT
 TTHHTHHHTTHTTTHHTTHTTTHHTTHTTTHHTTHTTTHHTHT
 THTTTTTHHTHTHHHHHTTHTHTHTTTTHTHTTTHHT
 TTTTHHHHTHTTTTT...

Suppose further that we identify 100 occurrences of the pattern TTHH which is followed by a T in 75 of the observed occurrences. We can say that in our sample, a T follows TTHH 75% of the time. But if we take this to mean that there is a 75% probability of tossing a T after each TTHH pattern, we would be in error. The result (T or H) of any coin toss is determined by two equally probable and mutually exclusive outcomes: a coin falls either on the head-side or on the tail-side. Therefore, whatever correlations we can mine out of our data series, each successive toss of a coin will always have a 50:50 probability of landing on either side.

The probability distribution of coin tosses is easy enough to grasp. However, our ability to interpret the cause-and-effect relationships in markets is easily overwhelmed by their complexity. Modern computer technology and the nearly infinite availability of data enable us to identify countless apparent regularities in the behavior of markets. But, unless we fully understand all the determining factors shaping market forces, we can never be sure whether the patterns we observe reflect causal relationships

⁴ ARIMA stands for autoregressive integrated moving average

at work or just coincidences. Probabilistic approaches to speculation spread out over many transactions might conceivably work, provided that the detected correlations in the data remain stable over time. We could be certain about this if we understood what precisely caused these regularities. In a gambling casino for instance, we know what determines the odds at the Roulette table or in the game of blackjack, and we know that those determining factors remain constant: the Roulette table has so many slots each time it is turned; the game of blackjack is played with the same number of cards of the same kind, etc. In securities markets we have no such constants and our grasp of the factors determining market prices is vague at best.

But how then should we explain Victor Niederhoffer's success with probabilistic models? Sadly, his success was relatively short-lived: on 27 October 1997 his hedge fund sustained a 100% loss on one trade – an event that supposedly had a 1 in 2000 odds of happening⁵. It is worth pointing out that Mr. Niederhoffer was not a naive rogue but a highly sophisticated player: in addition to 15 years of trading experience, his education included a degree in economics and statistics from Harvard University and a PhD in finance from Chicago University. His analysis was probably rigorous and precise. The problem was just that it was conceptually flawed.

Another, and at that time rather spectacular example of a great brain trust running conceptually flawed models was the demise of the Long Term Capital Management (LTCM) fund in 1998. LTCM was set up in 1994 by Wall Street's legendary bonds trader John Meriwether. His reputation enabled him to assemble a formidable group of scientists and computer programmers led by two Nobel Prize winning economists, Robert Merton and Myron Scholes. Their core investment approach was also based on probabilistic quantitative modeling. Essentially, LTCM's models scanned the investment universe for pairs of securities whose prices tended to move together. When their prices diverged from one another by so many standard deviations, LTCM would buy the *relatively* cheaper security, sell short the relatively more expensive one, and make a profit as their prices reverted to their historical relationship.

Like Victor Niederhoffer's fund, this worked impressively for a while, until suddenly it didn't. In August 1998 after Russia announced that it would default on her debt, LTCM found itself stuck with a highly leveraged position in Russian bonds which had plummeted in value. As a result, LTCM experienced a spectacular losing streak that led to the fund's

⁵ Segal, D. "Contrarian Gets Caught Flat-Footed by Market." International Herald Tribune. November 18, 1997: p. 13 and 18.

ultimate demise. In a few months' time, LTCM lost \$4.5 billion of their investors' assets and had to be bailed out by a group of 13 large Wall Street banks. Applying probabilistic approaches may make good sense in domains where odds are well understood and unchanging, as in games with dice or cards. Using them for speculation in markets can conceivably work for a time, but will likely fail in the end as correlations observed in the past change.

Technical analysis

The most important tool in investing is a ruler.

Nick Glydon⁶


Technical analysis concerns itself with the study of price charts using a bit of uncomplicated mathematics and simple geometry. Chart analysts use such concepts as trendlines, channels, speedlines, Fibonacci retracement levels or Andrews' pitchforks to divine where the price of some asset might be headed in the future.

They also look for patterns in the price charts such as flags, pennants, double tops, double bottoms or head-and-shoulders reversals. In addition, they normally use a variety of simpler mathematical concepts including moving averages, stochastics, parabolic trailing stops and Bollinger Bands. I started studying these by reading John J. Murphy's textbook "Technical Analysis of the Financial Markets".

At first, I had a hard time keeping an open mind. In fact, I thought the whole concept was a bit ridiculous. To my mind, the subjective nature of chart analysis and a general absence of any scientific rigor placed technical analysis in the same category as astrology and fortune telling. However, after a short time of using an ADP/Aspen Graphics computer program for technical analysis I realized that perhaps it wasn't such a total waste of time. All those strange constructs and patterns I'd read about kept appearing before me again and again, in any market I looked at and on nearly any time scale. Below are just a few examples:


⁶ Quoted by Albert Edwards (Societe Generale Cross Asset Research) in his 25 April 2013 "Alternative View" newsletter.

Exhibit 3a: Trendlines, channels and trend reversal patterns




Swiss Franc futures traded lower for two years from 1999 through 2000, but then started to reverse, first completing a small head-and-shoulders pattern, couched within a similar but larger reversal pattern. From 2002 through 2004, Swiss Franc trended upward within a channel defined by two parallel lines.

Exhibit 3b: Fibonacci retracement levels



A three-year decline in coffee prices sharply corrected in 2014; the correction rapidly covered about 59% of the previous drop, then 61.8% and spent the bulk of 2014 fluctuating between the 38.2% and 61.8% retracement levels acting as support or resistance during the price consolidation period.

Exhibit 3c: Andrew's pitchfork and a double-top reversal



U.S. 30-year Treasury Bond prices advanced during 1997 and 1998 along the lines determined by the so-called Andrew's pitchfork formation, defined by points A, B and C. This advance reversed with a fairly clear double-top pattern.

There are countless examples of asset prices advancing along straight lines or remaining confined within parallel channels. Trend reversals frequently trace head-and-shoulders formations, double bottoms or double tops. Often, significant trend moves correct by 38.2%, 50%, or 61.8% - the so-called Fibonacci retracement levels. And while these occurrences aren't precise, prices do seem to gravitate toward certain technical targets. As chart analyst, I found myself even more mystified by concepts like speedlines or the so-called Andrew's pitchfork, where some chart formation would determine trend lines and support or resistance levels for years into the future.

The recurrence of these patterns in just about any price chart I looked at genuinely puzzled me. Why should prices bounce off of straight lines or remain bound within parallel channels for months or years? Why should reversals so often form double top, double bottom, or head-and-shoulders patterns? Why should trends unfold for years on end within the bounds of speedlines whose slope is defined at the very outset of a trend move? Clearly, all these patterns are merely the result of the price discovery process – the buying and selling activities of traders.

But just how or why this process regularly produced such patterns remained a bit of a mystery to me. Technical analysis makes little attempt to explain the mystery, as it does not pretend to be a science. It is merely a repository of many decades of experience and observations by thousands of market practitioners, within the framework defined by three key beliefs: that prices move in trends, that the market price discounts everything, and that history repeats itself. In turn we examine each of these beliefs, in reverse order.

History repeats itself

History may not really repeat itself, but what this principle entails is that certain chart patterns observed in the past will likely continue appearing in the future with similar implications for price fluctuations. Stated otherwise, various patterns tend to occur repeatedly, offering analysts valid grounds to make a prediction about future price moves. For example, completion of reversal patterns like a double top or head-and-shoulders signifies that the recent trend may have reversed and that prices will continue moving in the opposite direction. Continuation patterns like flags, pennants and various triangles indicate that the prevailing trend will likely continue enabling us to project possible target prices for subsequent moves. The trouble with this belief is that like most of the rest of technical analysis, it is nearly impossible to verify through rigorous science. Price patterns don't always reappear in a precise form, identical to previous

occurrences, so identifying them in charts is a matter of judgment rather than exact science. Still, the experience of many practitioners – and I include myself here – strongly supports the belief that in this sense at least, history does repeat itself frequently enough.

Price discounts everything

Like the Efficient Market Hypothesis, technical analysis also assumes that all the information that's known and relevant to the value of some asset is already reflected in its price. So far as it refers to efficient markets – markets where large numbers of relatively small participants interact on a level playing field – this tenet is not terribly controversial. In efficient markets, the participants' collective knowledge of all the factors relevant to some security will tend to set the price roughly at the correct level. Again, this is a belief, not something we know for sure or even understand with much clarity. Exactly *how* price may discount everything is also a rather mysterious phenomenon. In his fascinating book "The Wisdom of Crowds," James Surowiecki recounts one illuminating instance of the price discovery process at work.

On 28 January 1986, 73 seconds into its flight, the space shuttle Challenger exploded over the Atlantic Ocean off the coast of central Florida. This tragic event triggered a revealing reaction in the stock market. In large part, Challenger's launch was the work of four major NASA contractors: Rockwell International, Lockheed, Marin Marietta and Morton Thiokol. Each of them was a publicly traded company. On the day of the Challenger disaster, stock price of each contractor started dropping some 30 minutes after the explosion, before most people even had the time to digest what had happened.

One firm was hit harder than others: within an hour of the explosion, Morton Thiokol's stock was down 6% and its trading had to be halted. After trading resumed, its stock continued falling and by the end of the day, it was down 12%. By contrast, the stock of other three contractors rebounded and closed with a loss of only about 3% for the day. The reasons why the stock market singled out Morton Thiokol weren't clear; on the day of the disaster, there were no public comments declaring that Morton Thiokol might be responsible for the incident.

On the following day, rumors about what had happened published in the papers did not implicate Thiokol either. In fact, it was fully six months after the explosion that investigators concluded that the Challenger blew up due to the O-ring seals on booster rockets that were built by Morton Thiokol, and that the other three contractors were not liable.

Can it be that within 30 minutes of the incident the stock market determined what took investigators months to ascertain? I must confess that I found this account hard to believe and as I read it, my immediate thoughts were that Morton Thiokol insiders must have dumped their shares in a thin volume session and as the price started dropping, some other participants may have followed suit and that's how the stock ended up battered. However, an analysis of the episode by finance professors Michael T. Maloney and J. Harold Mulherin cited by Surowiecki found that insiders did not sell their stock on that day. In fact, Maloney and Mulherin were entirely unable to come up with a convincing explanation for why Morton Thiokol stock was so quickly singled out by the stock market. It may well be that the market, in its mysterious collective wisdom somehow *knew* the relevant truths and set the price accordingly.

Surowiecki's book compellingly supports this possibility. Of course, as fascinating as it is to ponder the omniscience of collective wisdom, we must also acknowledge that markets periodically manifest manias or panics offering a very different perspective on their wisdom. But as Surowiecki argues, the ability of the collective to reach intelligent solutions to problems depends on certain conditions like decentralization of the flow of information, diversity of the participants, and their independence from one another in making decisions. If either of these conditions is compromised, the wisdom of crowds can – and periodically does – morph into madness.

In modern securities markets, the sources of information are centralized to a large extent, and independence of decision making often gives way to herd-like action. At times when certain momentous events are taking shape, large numbers of individuals follow the action of others rather than think independently. At such times, the wisdom of crowds can get dysfunctional, contaminating the price discovery process with unwarranted fear or excessive enthusiasm that can push prices far beyond levels that could be rationally justified. Wise or not, the psychology of market participants is what ultimately determines asset prices, so the belief that *it's all in the price* remains valid. Whatever the state of a market at any given time – be it rational, depressed, or exuberant – it forms the objective reality and we have no choice but to reckon with it.

Markets move in trends

The third tenet of technical analysis should be uncontroversial to anyone who ever looked at the price chart of almost any market security. Still, numerous learned members of academic institutions have managed to prove that this is not so. Some have gone as far as to claim that those who

think they see price trends in markets are probably hallucinating. I find it perplexing that intelligent people and tenured professors at top universities can find ways to refute something that's obvious even to my golden retriever. Academia's disdain for chart analysts and trend following has a rather long tradition, drawing much of its intellectual inspiration from the Random Walk Theory.

In a nutshell, Random Walk Theory views modern securities exchanges as models of efficient markets where all the information relevant to the traded stocks is already reflected in their prices. Future price fluctuations will be driven by random and unpredictable future developments, which will render those fluctuations random as well. This hypothesis was advanced by a number of theoreticians and academics including, MIT Sloane Business School's Paul Cootner who wrote the book "The Random Character of Stock Market Prices"⁷ in 1964, Eugene Fama who, wrote an influential paper⁸ titled, "Random Walks in Stock Market Prices," (1965) and Princeton University professor Burton Malkiel who popularized the Random Walk Theory – as well as the derision of technical analysis – with his 1973 best-seller, "A Random Walk Down Wall Street"⁹. Malkiel's book has enjoyed remarkable success and has sold in eleven editions through 2012. However, the part of his refutation of chart analysis that seems most compelling at first glance reads like a bit of a fable – not the standard that rigorous science should aspire to.

The fable of the shrewd scientist and a foolish chart analyst

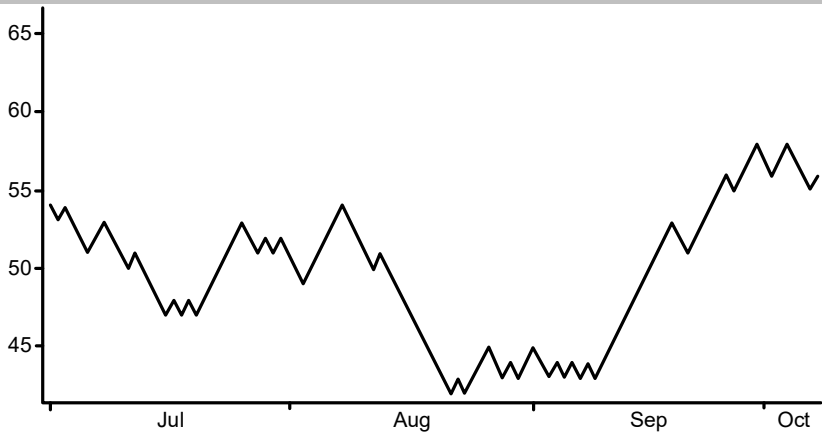
Malkiel's is a fable about the shrewd scientist and a foolish chart analyst. In this story, the shrewd scientist (Malkiel himself) goes to the foolish chart analyst and shows him a chart which he had previously conjured up by flipping a coin. Explaining that the chart represented the price fluctuations of some stock, the shrewd scientist pretended that he was interested in the "wise" chart analyst's divination. Not realizing the scientist's clever trick, the gullible chart analyst looked at the chart and said unto him: "Oh scientist, if you wish to become rich, you must buy this asset at once, for its price is heading higher."

⁷ Cootner, Paul H. (1964). *The random character of stock market prices*. MIT Press.

⁸ Fama, Eugene F. (September/October 1965). "Random Walks In Stock Market Prices". *Financial Analysts Journal* 21 (5): 55–59. doi:10.2469/faj.v21.n5.55. Retrieved 2008-03-21.

⁹ Malkiel, Burton G. (1973). "A Random Walk Down Wall Street" (6th ed.). W.W. Norton & Company, Inc. ISBN 0-393-06245-7.

Exhibit 4: A tosser's "trend"



A "price chart" constructed by coin tosses in Burton Malkiel's experiment

Upon hearing the chart analysts' words, the shrewd scientist laughed and replied, "Do you realize, foolish chart analyst, that this chart is based entirely upon coin-tosses?" Recognizing that he had been outwitted, the chart analyst turned red in rage. The shrewd scientist had unmasked his sorcery and showed it to be futile and worthless for the whole world to see. The foolish chart analyst would now be forever banished from the realm of serious discourse.

Something like that. Namely, Malkiel conducted an experiment where he gave university students a hypothetical stock priced arbitrarily at \$50/share. Each day's closing price was subsequently determined by the flip of a coin: *heads*, the price goes half a point up, *tails*, it goes half a point down. Malkiel took the resulting "price" chart to a chart analyst who promptly advised him to buy that stock. When Malkiel told him that the chart was based on flipping coins, the chartist was allegedly very unhappy. The story of this experiment, the resulting "price chart," and some inept analyst's recommendation was deemed by Malkiel as a solid ground to argue that stock price fluctuations are as random as coin-tosses. A more astute analyst might have caught onto the fact that all price changes occur in equal increments (\$0.50 up or down each day), something you've never seen in real-life daily price charts. Also, a more experienced analyst would have declined to make any recommendations based on only three months' worth of data. Indeed, since the first edition of Malkiel's book, much evidence has emerged suggesting that price fluctuations aren't entirely random, and that market prices do indeed move in trends.

Chapter 6: Markets move in trends

There are three avenues of opportunity: events, trends, and conditions.


Zhuge Liang, "The Way of the General"

We are always subject to a fear, when a market is moving up or down, that others know something we haven't yet figured out. So we feel a strong impulse to do what they are doing.

Robert Schiller, "Herd Behavior"

If you ever lived in a town or a city, attended school, read fashion magazines, or invested in stock markets, then you've inevitably experienced more trends in your life than you could name. The rising or falling popularity of music groups, fashion styles, political leaders or parties, social causes and even spiritual movements all frequently manifest unmistakable trends that surge through a society, reach their peak, and eventually fade. These are indeed such regular occurrences, so firmly rooted in human psychology that most everyone intuitively recognizes and understand them. Because financial and commodities markets also consist of human beings, trends are just as present and pervasive there.

Exhibit 1: Markets move in trends



This should not be a controversial assertion

Still, some humans, particularly those occupying the ivory towers of academia, continue to earnestly argue that there is no such thing as trends. Conceding that trends are real would clash with the theoretical framework that much of the academia explicitly adopted in studying markets. The Efficient Market Hypothesis and the Random Walk Theory maintain that market price fluctuations are random.

If they're random, there can be no trends, and that's that. You would think that a few price charts would suffice to settle this issue. Apparently this is not the case. This attachment of learned men to their theoretical home turf is something of a mystery of human psychology all in itself. An old Hebrew anecdote captures the point perfectly:

While a group of elder rabbis debated a section of Holy Law, a younger rabbi found himself in disagreement. He stated his case compellingly, but the elders disagreed, and pressed him to defer to them on this point. Convinced that he is right, he finally called upon god himself to help him convince the elders, asking god to make the rivers of Israel flow uphill if his position was right. God responded and the land's rivers promptly reversed direction. But the elders were not impressed and refused to change their mind.

Next, the young rabbi asked god to make all trees in Israel bend to the ground, and god obliged him again. Again, the elders were dismissive and unyielding. Exasperated, he finally asked god to speak to the elders directly, at which point the clouds parted, and a booming voice from heavens addressed the elder rabbis: "Hear me wise men, I confirm that the young rabbi is correct. He is right and you are wrong. What he says is what I intended." The young rabbi felt triumphant; surely the elders would now concede... But the elders remained unmoved: "we pay no attention to heavenly voices," they said, "the correct interpretation of this point was written long ago."

It appears that rigidity of convictions and aversion to contrary evidence is as old as history itself. All the same, let's look at some further evidence supporting the notion that trends do exist.

Trend followers

One group of hedge fund managers explicitly uses trends to generate investment returns. Trend followers are often referred to as CTAs (commodity trading advisors) and their investment vehicles as *Managed Futures* funds because as a rule, they tend to trade in commodity futures markets. Rather than cultivating expertise on any specific market,

industry, or geographical area, trend followers seek to identify trends in any liquid securities market and generate returns from advancing or declining prices. If Random Walk Theory adherents were right, then trend followers couldn't achieve positive returns on investment over the long term. But on this count, the random walkers would be wrong.

The table on the following page summarizes the performance of thirteen trend followers with between 16 and 38 years of continuous track record: as we can see, each money manager listed in the above table has generated very high investment returns over the matching time frames, even outperforming the U.S. stock market over the same period. If trends really didn't exist, this achievement would have to qualify as a miracle.

MASTERING UNCERTAINTY IN COMMODITIES TRADING

A partial list of trend followers and their performance vs. the S&P 500 index through April 2015

Trend followers		Track record	Compound annual rate of return *	Value of \$1,000 invested over the matching period
Manager	Fund or program	Years	Fund S&P500	Fund S&P500
Mulvaney Capital Management	Mulvaney Global Markets Program	16	15.77%	10,304
Drury Capital	Diversified Trend Following Program	18	11.01%	6,506
Superfund Group	Superfund Green Q-AG	19	10.63%	6,883
Clarke Capital Management	Worldwide	19	12.29%	9,324
Transtrend B.V.	Diversified Trend Prog. – Enhanced Risk	20	13.37%	12,710
Eckhardt Trading Company	Standard Program	24	13.94%	21,990
Rabar Market Research	Diversified Program	26	11.54%	17,606
Saxon Investment Corp.	Diversified Program	27	12.13%	21,012
Chesapeake Capital	Diversified	27	12.30%	23,408
Abraham Trading Company	Diversified Program	27	16.95%	71,462
Dunn Capital Management	WMA Program	30	14.63%	63,770
Campbell & Company	Financial Metal & Energies	32	11.02%	28,434
Milburn Ridgefield Corporation	Diversified Program	38	15.16%	219,222

* Over the matching time interval from each fund's inception through April 2015


Data source (except last two columns): Institutional Advisory Service Group (www.iasg.com)

Momentum investing

Another category of asset managers that successfully exploit market trends are the so-called momentum investors. In essence, momentum strategy entails buying stocks that have performed best over time and selling short those that have performed the worst. In other words, momentum investors buy stocks whose prices have been trending upwards and short-sell those whose prices have been falling. To test the momentum strategy, researchers from the London Business School, Elroy Dimson, Paul Marsh and Mike Staunton analyzed nearly 110 years' of stock market price history starting with the year 1900.

They constructed investment portfolios by selecting 20 top performing stocks in the previous 12 months from among UK's 100 largest publicly trading firms, and compared their performance to portfolios of 20 worst performers, re-calculating the portfolios every month. They found that previous year's lowest-performing stocks would have turned £1 invested in 1900 into £49 by 2009. By contrast, previous year's top-performing quintile of stocks would have turned £1 into £2.3 million¹, which reflects a 10.3% difference in compound annual rate of return.

Exhibit 2: Investing with trends through momentum strategy



Source: The Economist

¹ These figures correspond to the outcome at the end of 2009, following the 2008 market crash. At the close of 2007, the figures were even more impressive: the portfolio of winners generated a compound annual rate of return of 15.2%, turning £1 invested in 1900 into more than £4.2 million. The portfolio of worst performers would have returned only 4.5% a year, turning £1 into £111.

The gap between investments in best and worst performing stocks was even wider when data from the entire London stock market was taken into account. From 1955 onward, the portfolio of previous year's top performers generated a compound annual rate of return of 18.3% against the return of 6.8% for the portfolio of worst performing stocks.

Dimson, Marsh and Staunton found that these excess returns from the strategy of buying top-performing stocks were "striking and remarkably persistent" as it proved successful in 17 out of 18 global markets studied with data going back to 1926 for America and to 1975 for larger European markets (the sole exception was Japan, where results were based on post-2000 data).²

Professor Marsh's reaction to his own research was symptomatic of the academics' discomfort with objective reality when it fails to conform to theory. In a statement to Financial Times Marsh said that, "It is a very simple strategy, buying winners and selling losers. In a well functioning market it ought not to work. We remain puzzled and we are not the only ones; most academics are vaguely embarrassed about this."³

Market trends and value investing

The success of trend following and momentum strategies seems puzzling from the strictly common-sense aspect. Namely, they both involve buying high and selling low, which is contrary to our natural inclination to buy things at low prices and try selling them at higher prices. After all, this approach is at the core of *value investing* that made Benjamin Graham and his disciple Warren Buffet some of the world's most successful investors of all time.

Benjamin Graham authored "Security Analysis" and "The Intelligent Investor," widely considered as the most important books on investing ever written. He generated an annualized return of about 20% over a 20-year period. During this time the stock-market overall returned about 12% per year. Warren Buffett himself generated a compound annual rate of return of over 18% during 30 years of his career⁴. The S&P 500 index returned 10.8% during the same period.

²Financial Times, "Momentum effect gains new admirers" by Steve Johnson, 23 January 2011 - <http://www.ft.com/cms/s/0/8d7c8a92-26c6-11e0-80d7-00144feab49a.html#ixzz1I0Zl0FZP>

³ Financial Times, "Ignore momentum at your peril" by Steve Johnson, 18. 02. 2008 - http://search.ft.com/ftArticle?sortBy=gadatearticle&queryText=%22hedg%20fund%22&y=9&aje=true&x=16&id=080218000064&ct=0&nlick_check=1

⁴ Sizemore, Charles. "The Worst Investment of Warren Buffett's Career." Forbes, 5/8/2013.

Investment performance of world's best known value investors

	Compound annual rate of return	S&P 500*	Outperformance
Benjamin Graham	20%	12%	8%
Warren Buffett	18%	10.8%	7.2%


*Measured over the same time period

While Graham and Buffett are generally regarded as value investors, a closer look at their performance reveals, that their success had more to do with market trends than with superior value-finding skills. In the “The Intelligent Investor,” Graham observes powerful market trends as they confound his value judgment on stocks.

In 1953, as the US stock market enjoyed one of the longest running bull-markets until that point, he cautioned investors that the stock prices were getting too high. “*As it turned out,*” he later wrote, “*this was not a particularly brilliant counsel. A good prophet would have foreseen that the market level was due to advance an additional 100% in the next five years.*”⁵ By 1959, the Dow Jones Industrial Average reached an all-time high at 58.4, and again Graham warned investors that stock prices were “*far too high.*” Nonetheless, the Dow rose another 26% to 73.5 by late 1961 and after a subsequent 27% correction in 1962, it soared on to 89.2 in 1964.

⁵ Graham, Benjamin. *The Intelligent Investor*. New York: HarperBusiness, 2003. 73

Exhibit 3: Spot the trend – Standard & Poor's 500 index from 1857 to 2015



Since the 1940s, even for value investors, the great bull market trend in the U.S. stocks played a major role in generating investment returns. The effect of the logarithmic scale in the above chart visually diminishes the magnitude both of the massive bull market and of the periodic price corrections that took place along the way.

In sum, Graham thought that stocks were overpriced in 1953 as they were about to treble in value over the next eleven years. Selling your investments ahead of a 200+ percent bull market isn't a good way to earn high investment returns. So how did Graham generate the remarkable results from his investments?

The simple answer: by not following his own investment advice. Instead, Graham inadvertently did what a trend-follower or a momentum investor might have advised him to do: he held onto his best performing investment even though it was overpriced from the get-go. Namely, in 1948, he and his partner Jerome Newman purchased a 50% interest in the Government Employees Insurance Company (GEICO). The \$712,500 purchase was roughly a quarter of their fund's assets at that time.

Here's what Graham says about their GEICO investment in the post-script to "The Intelligent Investor": "... *it did so well that the price of its shares advanced to two hundred times or more than the price paid for the half-interest. The advance far outstripped the actual growth in profits, and almost from the start the quotation appeared much too high in terms of partners*⁶ *own investment standards.*"⁷

Graham explains why he and Newman did not sell GEICO even though they judged its price "*much too high*" from the start. Because of the size of their commitment and involvement in the firm, they regarded it "*as a sort of 'family business,'*" and maintained ownership in it in spite of its spectacular price appreciation. In Graham's words, the profits from this single investment decision, "*far exceeded the sum of all the others realized through 20 years of wide-ranging operations in the partners' specialized fields, involving much investigation, endless pondering and countless individual decisions.*"

In other words, *far* more than one half of Graham and Newman's performance came from the one investment they kept through a two-decades' bull market and *did not* sell it even though it was grossly overpriced "*in terms of partners' own investment standards*". That implies that all their "*investigation*" and "*endless pondering*" contributed less than 10% in annual returns, underperforming the stock market by at least 2 percentage points over 20 years.

That further implies that if Graham and Newman *only* invested in GEICO and spent the rest of their careers fishing and golfing rather than burdening themselves with investigations and endless ponderings, they would have done at least twice as well as they *have* done, generating annual returns of 40% or more from 1948 to 1966!

⁶ When Graham says, "partners," he means himself and Newman.

⁷ *Ibid.*, 532, 533.

For his part, Warren Buffett's style reveals much more of a momentum player than value picker. He made many of his large investments on the back of major run-ups in stock prices. Examples include his investments in Capital Cities (1985), Salomon Inc. (1987 and 1994), Coca Cola (1988), Gillette (1991), Freddie Mac (1991/2), General Dynamics, (1992), and Gannet Company (1994).⁸

When Buffett bought over \$1 billion of Coca-Cola shares, they had appreciated more than five-fold over the prior six years and more than five hundred-fold in the previous sixty years. This decision proved right, as his investment in Coca Cola quadrupled in value over the following three years, far outstripping the S&P 500.⁹

And like Graham before him, Buffett owes much of his success to GEICO. He started buying its stock in 1975 at \$2 per share, and kept adding to this investment even as GEICO's market cap went from \$296 million in 1980 to \$4.6 billion in 1996. This growth in valuation corresponded to a compound annual rate of return of 29.2%, an outperformance of more than 20% per year over the S&P500!¹⁰

Did Warren Buffett sell his stake in this overvalued¹¹ company? To the contrary, in 1996 Buffett bought 50% of it, making Berkshire Hathaway 100% owner of GEICO. This was not exactly a value pick, but the decision again proved a winner: by 2011, GEICO's market cap nearly quadrupled to \$20.5 billion based on Warren Buffett's valuation model.

Even though Graham and Buffett somehow came to epitomize the so-called value-driven investing, both owe their success to market trends. In American stock markets, bullish trends were out in full force through most of Graham's as well as Buffett's careers which were most abundantly blessed by some of their most "overvalued" investments. In essence, Graham and Buffett may both have overtly espoused value investing because it's a rational style that sits well with investors. However, their outperformance was driven far more by trends than by their value picks.

Human psychology: the driver of trends and bubbles

Economic value is central to our decision making and it plays a major role in our intuitive psyche. In daily life, when we buy a loaf of bread or a tank-full of gasoline, we tend to have a good idea about what we think is cheap and what's expensive. We like to find bargains and don't enjoy

⁸ These investments are treated in some detail in R. Hagstrom's "The Warren Buffett Way."

⁹ Hagstrom, Robert G. *The Warren Buffett Way*. New York: Wiley Investments, 1995. (v)

¹⁰ During the same period, S&P500 grew by 8.9% per year

¹¹ At the time, GEICO's book value was \$1.9 billion, which means that the remaining part of its \$4.6 billion market cap was goodwill, rendering GEICO's shares very "expensive."

being ripped off. Just as we are inclined to shop for value as consumers, we find value investing intuitively appealing. However, there's a critical bit of difference between buying goods and investing: shopping for investments is speculative while buying stuff isn't, and speculation activates the part of our mental circuitry that can heat up to a boiling point and overwhelm any rational consideration of value.

When a multitude engages in speculation on some desired asset, their activity frequently produces price trends and in some cases bubbles of great proportions. Here's how this dynamic shapes up: in making investments, our rational goal is to obtain the best possible return with the least risk necessary. If we buy a house or a stock for investment, we want to receive a stream of rents or dividends and preferably the opportunity to sell the asset for a price that's higher than what we paid. Since those outcomes depend on other market participants, we are obliged to reflect on what they might do. Thus, if house prices are going up we infer that people are keen on investing in real estate and that rising demand would push future house prices even higher. If we are convinced that this is the case, we might disregard the fact that houses are already expensive. In effect, led by the actions of others, we might accept inflated house prices and proceed with the investment anyway.

This dynamic was demonstrated empirically in a clever experiment designed by Colin F. Camerer at Caltech's Experimental Economics Laboratory.¹² In this experiment, a group of students were asked to trade shares in a hypothetical company during 15 five-minute periods. The students were not allowed to discuss their actions and only communicated via *buy* and *sell* orders. To start with, each student received two shares and some money with which to buy more shares. At the end of each of the 15 periods, the shares paid a \$0.24 dividend for a total payout of \$3.60 per share throughout the experiment ($\$0.24 \times 15$).

This provision removed any uncertainty about the shares' value: at the start of the experiment, the maximum value of one share was \$3.60 and this amount diminished by \$0.24 after each round, since that amount of dividend was already paid out. The highest price any player should accept to pay for a share should not be one penny more than what that share would yield in remaining dividends. However, Camerer's experiment showed otherwise. When the experiment started the share price immediately jumped to \$3.50, close to the shares' rational value. But rather than steadily declining with each new round, the price remained near that level almost to the very end of the experiment. Even when the value of each share fell below \$1, students were still willing to pay \$3.50

¹² Surowiecki, James. *The Wisdom of Crowds*. New York: Anchor Books, 2004.

to buy them. When the experimenter asked the students why they bought the shares at prices that obviously far exceeded their value, the students would reply that, “*Sure I knew that prices were way too high, but I saw other people buying and selling at high prices. I figured I could buy, collect a dividend or two, and then sell at the same price to some other idiot.*”¹³ A strange confluence of circumstances produced this very same dynamic in a real-life experience that became known as the Chinese Warrant Bubble, described in a remarkable paper by Princeton University’s Wei Xiong and Columbia University’s Jialin Yu.¹⁴

Chinese Warrant Bubble

In an effort to develop China’s financial derivatives market, from August 2005, China Securities Regulatory Commission (CSRC) started introducing a small number of warrants – financial instruments similar to options, issued by publicly traded corporations. Firms were allowed to issue call or put warrants. With call warrants, issuing firms granted investors the right to buy stock from them, and put warrants gave them the right to sell stock back to the issuing company at a specified strike price and time period during which investors could exercise their option to buy or sell stock shares. Between 2005 and 2008, 18 put warrants with maturities from 9 to 24 months were issued to the public.

During this very period, the Chinese stock market experienced a strong bull run and its index vaulted from 1,080 points in June 2005 to 6,124 in October 2007. This rally quickly pushed most put warrants so deep out of the money that they became worthless. In spite of this, feverish speculation on these securities produced an extraordinary financial bubble, unique in the history of bubbles because warrants continued trading at spectacularly high levels of turnover and very inflated prices, even as it became evident that their value had clearly dropped to zero.

Consider the case of a Chinese liquor producer, WuLiangYe Corporation. On April 3, 2006 WuLiang issued 313 million put warrants with a two year maturity and a strike price of 7.96 yuan. The initial price of the warrants was 0.99 yuan and company stock traded at 7.11 yuan. Although the warrant was *in the money*¹⁵ when issued, the dramatic rise in WuLiang’s shares pushed it out of the money in only two weeks after which it never came back in the money. WuLiang’s stock price rose ten-fold, reaching 71.56 yuan in October 2007 before retreating to about 26

¹³ Idem.

¹⁴ Wei Xiong and Jialin Yu. “The Chinese Warrants Bubble.” National Bureau of Economic Research, Working Paper 15481 - <http://www.nber.org/papers/w15481>.

¹⁵ Meaning, the warrant’s strike price was higher than WuLiang’s stock price.

yuan on April 2, 2008 when the warrant expired. Rather than falling in value as they got farther out of the money, WuLiang's put warrants rose along with the company's share price, at one point even surpassing their own strike price at 8.15! Paying 8.15 yuan for an instrument that has a maximum possible payout of 7.96 yuan (if the firm's share price went to zero) makes little sense, but someone did pay that much. Meanwhile, according to the widely used Black-Scholes model, the warrant's value fell below 0.0005 yuan after July 23, 2007 and remained below that level for the remaining nine months of the warrant's maturity.

Still, the warrant continued trading at a price of several yuan, dropping below 1 yuan only in the very last few trading days and dropping to zero literally in the final minutes of the warrant's last trading day. This same phenomenon played out with all 16 put warrants analyzed by Wei Xiong and Jialin Yu. For each, the Black-Scholes valuation dropped to nearly zero (below 0.0005 yuan) where it remained on average for 54 days. During this zero-value period, each warrant traded at spectacularly high turnover levels¹⁶ corresponding to billions of US Dollars per day and at an average price of 1.00 yuan – more than 2,000 times their value.

Chinese warrants bubble provides the clearest evidence to date that in speculative decision making, our views about the actions of others can entirely override any rational appraisal of an asset's value. That in turn gives us a convincing perspective on the reality of market trends: asset prices are not always driven by objective valuation, only to be randomly affected by random external events. Instead, prices are driven by human psychology and its self-stoking collective action capable of sustaining major trends that can last many years. Consequently, as investors and traders we have little choice but to recognize trends as real and legitimate sources of investment opportunity.

Trends in the broader economy

Of course, there is more to trends than just asset prices, which are merely a singular expression of broader economic processes at work. An extensive study conducted by the consulting firm McKinsey & Co. will help us appreciate their importance. McKinsey analyzed the performance of some 100 of the largest US corporations from 17 different sectors of the U.S. economy over two business cycles, from 1984 to 1993 and from 1994 to 2003. The study¹⁷ sought to answer the question: *“How does a large*

¹⁶ During their zero value period, each warrant had an average daily turnover rate of 291% and an average daily volume of 1.26 billion yuan, meaning that each warrant changed owners three times during an average day even though it was essentially worthless.

¹⁷ Smit, S., et al. The do-or-die struggle for growth. McKinsey Quarterly, August 2005.

company achieve and maintain strong growth?” The authors set out to understand which factors made some corporations more successful than others in terms of revenue growth and total returns to shareholders (TRS). They expected that answers would emerge from individual firms’ performance in strategy, marketing, operations and organization. What they discovered instead was startlingly different. From among 102 corporations studied over the 1994-2003 cycle, they identified 32 “growth giants” – firms whose revenue growth outpaced the GDP and whose stock outperformed the S&P 500. Among these growth giants, 90% were concentrated in only four sectors of the economy: financial services, health care, high tech, and retailing.

Those four sectors happened to have enjoyed favorable market trends during the business cycle: financial services benefited from deregulation, increased borrowing and an increasing public participation in equity markets; health care expenditure grew with the nation’s aging population and through innovation; the high-tech industry also enjoyed a massive wave of innovation in the 1990s; retailing grew through growing consumer affluence and format innovation by firms like Wal-Mart, Target, Lowe’s and Home Depot. While the overall economy grew at a rate of 5% from 1994 to 2003, financial services grew by 7%. High-tech also grew 7% overall with services in high-tech industry growing even faster at 9%. Health care expenditures grew at 7%, but most of the growth in health care sector was concentrated in pharmaceuticals, which expanded by 12.5%! In retailing which grew slower than the GDP at 4.5%, growth giants expanded much faster.¹⁸

McKinsey’s analysts wrote that, “*What’s striking for a large growth-minded corporation is just how crucial it is to have this kind of favorable wind at their backs when they try to achieve strong growth.*” Indeed, favorable market developments gave rise to trends that were the key driver of value creation for 90% of the most successful corporations.

By contrast, “*when large companies face slow-growing markets,*” wrote the report’s authors, “*opportunities to change the growth trajectory are limited.*” Warren Buffet anticipated this finding in his famous remark that, “*When a management with reputation for brilliance tackles a business with reputation for poor fundamental economics, it is the reputation of the business that stays intact.*”¹⁹

¹⁸ Only one of the 102 corporations – perhaps the exception to prove the rule – built a big new business without the backdrop of a strong trend of growth in the market: Wal-Mart managed to grow rapidly in the slow growing market for perishable groceries through leveraging of its brand, supply-chain muscle and format innovation.

¹⁹ Berkshire Hathaway Annual Report, 1985, p. 9.

To conclude this chapter, much compelling evidence supports the following simple assertions:

- Markets move in trends.
- Trends shape the price discovery process over the long run.
- Trends represent one of the key drivers of value creation (or destruction) for investors and businesses.

Far from being a figment in the imagination of the unlearned, market trends could well be the single most important element to consider in generating and sustaining investment returns over time. The case, at any rate, appears compelling and it is time we parted ways with elegant but erroneous intellectual contrivances that contradict what is so plainly obvious to most participants in the real world.

Chapter 7: Speculation and human psychology

The person is a conglomerate of independently functioning mental systems that in the main reflect nonverbal processing systems in the brain.

Michael Gazzaniga

For indeed, the investor's chief problem – and even his worst enemy – is likely to be himself

Benjamin Graham

Two or three years into my apprenticeship, I started to embrace two key ideas: that markets move in trends, and that asset prices themselves are a unique means to know something accurate and objectively valid about the markets. The price of an asset gives us a very limited picture about its market, but that picture is timely, clear and unambiguous. The possibility of using this form of knowledge to systematically generate decisions with positive expectancy seemed promising. Importantly, if trading decisions could only be based on an asset's price, this would allow me to transcend my "home turf" of oil markets: by mastering the art of price chart analysis and trend following, I should be able to trade in other markets like coffee, cotton, soybeans, copper, gold, currencies, and stocks. This seemed like a fantastic prospect.

The trouble was, at that stage of my personal evolution, these were merely vague notions suffused with desires and ambitions which I was unsure how to realize. If there was a kind of Holy Grail of speculation for me to uncover, I did not have any idea what it was or how I could advance toward it. By contrast, problems and obstacles were numerous and very clear. Markets may move in trends, but trends don't announce themselves in advance and they only become obvious in hindsight. Trading decisions are made in the present, and their outcome depends on future events that are unknowable. Even when a trend is evident, the problem of *when* and at *what price* to put on a trade is seldom an easy choice.

For example, if a market trend is obvious, the price could be near its peak and due for a correction. A correction may be a good buying opportunity, but what initially looks like a correction could turn out to be the trend's reversal. These issues reflect the unavoidable problem of

uncertainty, which defies easy answers. Technical analysis as a way of analyzing a market's price history has its own limitations. Almost as soon as I started using it in my market reports, a sense of futility curbed my initial enthusiasm.

Interpreting any given price chart is seldom clear-cut; you can always find indicators that suggest one thing, and others that suggest the exact opposite. You might, for instance, be looking at an evident uptrend in some market implying that you should go long, but your momentum indicator could be showing that the market is overbought and possibly due for a downward correction. Then, a correction might look like a good point to buy into the trend, particularly if the price reached some trend line or Fibonacci retracement level, except that the same move might also have completed a double-top pattern suggesting that the trend could be reversing, in which case you should go short. Accordingly, my early reports turned out to be a mix of roughly correct calls, dead wrong ones, and many that were neither. Here's a glimpse of my frustration as I put it in my journal at the time:

21 July 1998

I think one big problem we have with technical analysis... is that rather than observing a situation and considering opportunities (not touching the stuff unless there is one), we sit there day after day, trying to predict what the market might do the next day. Sometimes we get the direction right, sometimes we get the target right and sometimes we get the timing right. We hardly ever get all the elements right...

With time however, I was able to improve my reports quite significantly – not by better divining future price moves, but by identifying good trade ideas with a specific entry point, an expected price target and a stop-loss price level at which we should abandon the trade and accept a loss. I also learned to keep a degree of discipline in resisting the pressure to produce trade ideas if I didn't see any. Gradually, I felt better and better about the quality of my work, but the sense of futility was still there.

Feeling good about my work might have been a subjective impression and it didn't give me a sense of any real accomplishment. Whatever I thought about my analyses, there really was only one true way to objectively measure their quality, and that was through the outcome of actual decisions, once taken and executed. Well, that was a problem because I wasn't sitting at a large investment bank but at a family-owned oil trading company and my boss wasn't about to give me any play money

to test my ideas in live trading. The only other option I had was to trade simulated accounts and I did have a way to do that: the computer system I was using for technical analysis allowed me to set up fictitious trading accounts, place trades in real time and keep track of my profits and losses. I used this system to execute simple directional trades: if I thought the price of something was going up I bought it and if I thought it was going down I sold short.

I started with oil and currency futures, then proceeded to dabble in other commodities like copper, coffee, soybeans, and equity index futures. I wasn't trading with real money, but the sheer desire to see profits and get the sense that I might amount to something as a trader got me emotionally fully engaged in the process. I'd love to tell you that this was fun and that I got all passionate about trading, but it wasn't fun and I didn't enjoy it.

The losing game and its lessons

Gradually it became clear to me that I had a significant tendency to lose money and eventually wipe clean one account after another. This was very disturbing and I tried to uncover any errors of my ways by keeping a trading journal. One of the insights I gained – this may seem obvious, but it was not obvious to me at that time – was that each trade actually consisted of two separate decisions: the decision to commit to a trading position and the decision to uncommit.

Generally, getting into a trade was rather easy. It was getting out that often got messy. I noticed that regardless of how clear-headed I wanted to be about formulating trade ideas and executing them, once there were profits or losses involved, I ended up veering off plan and pulling the trigger for reasons I couldn't easily explain to myself. Closing a trade in profit was satisfying, but this satisfaction would quickly fade if it turned out that I would have made more money had I stayed in the trade. Next, I'd find myself scrambling to reopen the position, but doing so at a price less favorable than the price at which I closed the last trade seemed unbearable. Closing a position at a loss was even more unbearable, and I realized that this activity came with a disconcerting dose of stress.

The feeling of satisfaction was relatively rare and usually short lived while stress fouled up most of the time I spent trading, which sadly was very considerable. There were days when I spent most of my time glued to the screens, watching numbers and charts blinking in front of me, setting up trade orders and price alarms, revisiting my analyses, second and third-guessing them, cancelling my trades then putting them on again. I thought that I could become positively obsessed as I found myself turning down lunch invitations and drinks with friends because I didn't want to be away

from the screens. And I wasn't even trading with real money! I had to ask myself if I really wanted to spend my life in this way, obsessing over something that had a huge tendency to make me miserable most of the time. The answer was clearly, no. That in itself was one of the most useful lessons I gained from the experience. Another was the realization that this game was not so much about mastering the markets or statistics or even the charts as much as it was about mastering oneself. In speculation, markets are the external reality, but what decides the game's outcome is the inner process that determines one's actions.

Key lessons from trading	
Each trade entails two decisions	Getting in and getting out are two separate decisions. Getting out tended to get messy.
Trading is very stressful	Satisfaction in trading was rare and usually short lived, but stress was nearly constant. This was a good way to have no life away from the markets and computer screens.
Losing is the most likely outcome of speculation	I found that I had a significant tendency to lose money. Soon enough I realized I had this in common with most other speculators.
This is no way to spend a lifetime...	... and probably end up broke.
It's me, not the markets	The markets are the outside, objective reality. Winning or losing depended on me, the decision maker. Thus, the Holy Grail entailed mastering myself more than mastering the markets.

With the realization that the Holy Grail was in the decision making process rather than in the knowledge of markets, I became keenly interested in human psychology and especially in the mystery of how we make decisions when facing uncertainty and risk.

This was yet another vast domain to explore for inspiration and solutions. The questions were many: I wanted to understand how we learn, how we know, how our brains form judgment, how we act in complex situations, how we handle risk, and how making or losing money affects us. I also wanted to know if there were particular attributes characterizing successful traders and how these traders did what they did.

Rogue traders

One of the things I soon learned was that I wasn't the only one with a tendency to lose money. It turns out that the vast majority of speculators

eventually do that. My boss at the time always seemed eager to share with me the many stories he knew about high-flying traders who ended up losing everything. He often peppered his stories with remarks like, *“What, you think you’re better than those guys? You think you’re the next George Soros?”* I think he meant well and wanted to impress upon me that speculation is a risky game that usually ends in tears. Whatever his reasons, the result of his influence was that I spent more time pondering failure and how to avoid it than thinking about how to attain success.

The story of Victor Niederhoffer¹ also made an indelible impression on me almost from the outset. Over time I read about many more cases of great traders whose mystique and charisma dazed countless investors and corporate directors, who routinely made huge trades, whose opinions made news and moved markets, and who ultimately came crashing to the ground in disgrace, in some cases blowing up their entire firms. One of the most notorious examples was Nick Leeson whose “brilliance” brought down Barings Bank in 1995 when it transpired that he lost over \$1.4 billion of the bank’s cash. Back in 1995, that was a lot of money! Then there was Sumitomo Corporation’s Yasuo Hamanaka who managed to accumulate \$2.5 billion in losses between 1986 and 1996. Daiwa Bank’s Toshihide Iguchi lost \$1.1 billion. In 2002, Allied Irish Bank’s trader John Rusnak ran up trading losses of about 860 million Euros.

China Aviation Oil was bankrupted by its star trader, Chen Jiulin who lost over \$500 million trading oil derivatives. In 2006, whiz-kid mathematician, Brian Hunter single-handedly lost \$6 billion trading Natural Gas derivatives at Amaranth Advisors hedge fund. Two years later, another star trader went down in flames. His name was John Wood and he came from the UBS bank where he built up a stellar track record and ranked as the bank’s top trader. In 2006, John Wood set up shop in the Principality of Monaco where I had lived. The launch of his hedge fund, the SRM Global Master Fund, generated a great deal of publicity. Seduced by Mr. Wood’s star status, investors piled into his fund with such zeal that SRM became the largest ever European hedge fund at launch. This was in spite of its highly unfavorable terms, including a five-year capital lock-up.

Less than two years after the launch of SRM, the great 2008 financial crisis exploded and sunk the fund, its clients losing more than 85% of their investment. Of course, John Wood wasn’t the only casualty of the 2008 crisis: this was a veritable financial tsunami that engulfed nearly all the major financial institutions, causing losses of such magnitude that they dwarfed anything we’d experienced before. The following table provides a

¹ See the Introduction to this book.

list of the best known speculative debacles since the early 1990s, a list that is almost certainly very incomplete:

The landfill of speculator talent – a very partial list

Year	Company	Market	Loss
1992	Bank Negara Malaysia	Foreign exchange	\$5.5 billion
1993	Metallgesellschaft	Crude oil	\$1.3 billion
1994	Kashima oil (Japan)	Foreign exchange	\$1.5 billion
1994	Orange County, California	Financial derivatives	\$1.7 billion
1995	Barings Bank	Japanese securities	\$1.4 billion
1996	Daiwa Bank	Bonds	\$1.1 billion
1996	Sumitomo Corporation	Copper	1.8 billion
1997	Yamaichi securities	Various investments	¥200 billion
1997	Morgan Grenfell	Equity investments	\$650 million
1998	LTCM hedge fund	Russian securities	\$4.2 billion
2000	BAWAG (Austria)	Foreign exchange	€1.4 billion
2001	General Electric	Financial derivatives	\$1.2 billion
2001	Frostman Little & Co.	Various investments	\$ 2 billion
2002	Allied Irish Bank	Foreign exchange	€860 million
2003	GPIF (Japanese pension)	Various investments	¥6 trillion
2004	National Australia Bank	Foreign exchange	A\$ 360 million
2004	China Aviation Oil	Oil derivatives	\$540 million
2005	HSBC	Interest rate derivatives	\$500 million
2006	Amaranth Advisors	Natural gas derivatives	\$6.6 billion
2007	WestLB (Germany)	Equity investments	€600 million
2007	Bank of Montreal	Natural gas derivatives	C\$680 million
2008	Morgan Stanley	Credit default swaps	\$9 billion
2008	Societe Generale	Equity investments	\$7.1 billion
2008	Peloton Partners	Subprime mortgages	\$2 billion
2008	Aracruz (Brazil)	Foreign exchange	\$2.5 billion
2008	Deutsche Bank	Derivatives	\$1.8 billion
2008	Groupe Caisse d'Epargne	Derivatives	€750 million
2008	Sadia (Brazil)	FX derivatives, credit	\$1.1 billion
2008	Saracen Group	Natural gas derivatives	\$700 million
2009	SEM Group LP	Oil derivatives	\$3.2 billion
2009	Harvard University	Interest rate derivatives	\$500 million
2011	Olympus Corporation	Various investments	¥376 billion
2012	UBS	Exchange-traded funds	\$2.3 billion
2012	JPMorgan Chase CIO	Credit default swaps	\$6.2 billion
2015	Standard Chartered Bank	Commodity loans	\$4.4 billion

Smaller cases likely count in the thousands but often manage to avoid the media spotlight. A McKinsey study² published in 2003 gave us an idea

² Buehler, Kevin and Gunnar Pritsch, “Running with risk” – McKinsey Quarterly, November 2003

about just how frequent such incidents might be. McKinsey looked at the performance of 200 leading financial firms over a five-year period from 1997 to 2002, and identified fully 150 incidences of “significant financial distress” during that time. Authors defined significant financial distress as either a bankruptcy filing, a credit ratings downgrade of two or more notches, an earnings decline of over 50% below analysts’ consensus estimates, or a decline in total returns to shareholders of over 20% below the overall market in any one month.

This high incidence of financial distress suggested that the average financial firm had a staggering 75% probability of experiencing such adversity in any five-year period. The same might approximately be true for any firm that has substantial exposure to commodity price, currency, or interest rate risk. As a rule, financial distress is the ultimate outcome of speculation gone out of control. The perpetrators invariably turn out to be respected managers and highly skilled traders. This begs the obvious question: why is failure so pervasive in speculation? Why do so many smart, respected traders end up with such massive losses? The answer, in large part, lies in our psychology.

A dangerous mismatch: the human brain at speculation

Sustained success at speculation depends on one’s ability to consistently make good decisions about getting into and out of trades. While it’s nearly impossible to make money on every transaction, a successful speculator should get it right most of the time. More realistically, he should make more money when he gets it right than he loses if he gets it wrong so that over time his cumulative gains outweigh the total of his losses.

In practice, this is extremely difficult for most people to accomplish, due to a number of systemic biases in our psychology. Some of these biases are hardwired in our brains by design and can’t easily be cured by education or experience. They include phenomena like overconfidence, anchoring, the endowment effect, loss aversion, and several others that can induce a strong emotional pull on our judgment and distort a reasoned analysis of facts even in the most experienced professionals. Take the overconfidence bias: a large majority of us – close to 90% – rate ourselves above average in our ability and intelligence.

In speculation, a measure of success can give us an exaggerated sense of our own competence, making us prone to taking risks even in situations we understand only vaguely. We are also susceptible to the anchoring bias whereby we tend to rely, or anchor our decisions on a single issue or piece of information while ignoring or underestimating the importance of other relevant factors. The endowment effect predicts that we’ll demand a

higher price for an asset we already own than we would pay for that asset if we didn't own it. This may strike close to common sense, but it has important implications for how we deal with risk. Behavioral economist Richard Thaler studied how individuals evaluate risk to their lives. He asked a group of people two questions. First, how much would you pay to eliminate a one-in-a-thousand chance of immediate death? The second question was, how much would you require to accept a one-in-a-thousand chance of immediate death?

Typically, his subjects would pay no more than \$200 to eliminate the one-in-a-thousand chance of death, but they wouldn't accept the extra one-in-a-thousand risk of death for \$50,000. The disparity between the two answers is intriguing, given that the subjects were evaluating essentially one and the same risk. The ways we interpret and act on new information are also rife with complexity. Fluctuating almost around the clock, modern markets generate a constant flow of news and information enabling traders to keep on alert at all times and remain in control of their positions and risks. This may seem like a good thing, but the reality is that most traders would be better off staying away from the news flow altogether. Numerous empirical studies have shown that even among experts, more information doesn't, in fact, improve decisions. One such experiment, conducted by psychologist Paul Andreassen at the Massachusetts Institute of Technology looked at the way access to information influenced investment performance.

Andreassen divided students into two groups whose participants each selected a portfolio of stock investments. In each group, students were free to buy and sell stocks as they saw fit, but while one group had access to the constant flow of stock markets news, the other group was allowed to monitor their portfolios only through changes in stock prices. The experiment showed that students who got no financial news at all earned double the returns of those who frequently checked the news. This outcome is in part related to the one bias that perhaps more than any other, predisposes us to losing: our aversion to losses.


The psychology of loss aversion

Trading and investment management are long-term pursuits where performance reflects the cumulative results of a long series of transactions. However, rather than considering every decision as just one of many, we treat each transaction as a departure from the status quo, where our fear of loss overpowers our desire for gain. In fact, the logic we apply to decisions about gains is quite opposite that which we apply to decisions about losses. This phenomenon was first described by psychologists

Daniel Kahneman and Amos Tversky, who named it the “failure of invariance.” Through a series of empirical studies, they discovered that we tend to be strongly risk averse with regard to gains, and risk seeking when faced with losses. The failure of invariance thus predicts that we are risk averse when preserving a favorable status quo, but prone to taking risks when dealing with losses. In trading, this creates the disposition to exit profitable trades too early, and “work” the losing trades too long, even taking on more risk in order to try and reverse the losses.

The pressure to recover losses can lead traders to escalate risk to massive proportions, which can precipitate disasters like those we saw in the previous section. That’s what happened also to Victor Niederhoffer, whom we mentioned in this book’s introduction. In 1996, he was rated the world’s top fund manager, but after 15 years of outstanding performance his business came to an abrupt end in October 1997. At that time, his entire fund was wiped out in a single day when the market moved against his short positions in S&P500 put options. The fact that an investor with his credentials, experience and track record should take such massive risk on a single trade is quite astonishing.

Exhibit 1: Niederhoffer Intermarket Fund performance in 1997 (Jan. '97 = 100)



Game over. A large loss in August 1997 led Victor Niederhoffer to take excessive risks to try and recover from his draw-downs. His short trade in S&P500 put options was so large that when the market moved against him, it entirely wiped out his fund in a single day in October 1997.

Mr. Niederhoffer’s fatal trade was partly a consequence of loss aversion: in August 1997 his fund sustained heavy losses on investments in Thailand’s currency and stock market. In September, after recovering

some losses, his fund was still down nearly 40% for the year. The pressure to recuperate the losses led him to excessive risk taking, a mistake which he warns against repeatedly in his book, “The Education of a Speculator”. Loss aversion explains why it is so difficult for most people to follow the often quoted formula for successful investing: “*let the profits run and cut losses short.*” We are predisposed to take profits while they are a sure thing, and let losses run, gambling that the markets will turn in our favor. In other words, we seem to be hardwired to follow the exact opposite formula – we are inclined to cut our profits short and let losses run. This creates a strong tendency in most traders and investors to gradually lose ground against the markets.

Loss aversion underscores the fact that our mental faculties simply aren’t suited to the task of speculating in fast moving securities markets. Human brain is the product of our natural evolution, designed to solve problems of survival we confronted through our evolutionary history. During more than 99% of that time, we lived as foragers in small nomadic bands, and in that environment, loss aversion bias did make good sense. With no refrigerators, bank vaults or stock certificates, most improvements to our natural state had sharply diminishing marginal utility.


More food is good, but there’s only so much you can eat or hoard before it starts to become a liability. By contrast, reduced access to food, or an injury could rapidly spell “game over.” As MIT professor Andrew Lo fittingly put it, “*This notion of loss aversion, being more aggressive when you're losing and more conservative when you're winning, is a very, very smart thing to do when you're being hunted on the plains of the African savannah. However, it's not a smart thing to do when you're on the floor of the New York Stock Exchange*³.”

A matter of judgment

Loss aversion can cause a trader to lose money even when his judgment is correct. Judgment is a discrete process that fluctuates continuously with time and new information. Decisions are binary; they take effect at a precise point in time and determine the results of our actions. Unless his decisions are executed with flawless timing, a trader may have to endure unrealized losses on his positions for a period of time, straining his emotions and putting his conviction to trial. Consider the scenario depicted on the following page:

³ Fitzgerald, Michael “Survival of the Richest,” – MIT Technology Review, 19 April 2006.

Exhibit 2: Even with correct judgment, investors may sustain losses



Even where decisions are based on correct judgment, unpredictable market fluctuations can temporarily lead to considerable losses, eroding an investor's confidence in his judgment and triggering the loss-aversion impulses that tend to have adverse effect on performance.


In 2010 when the Nikkei 225 index was trading around the 10,000 level, numerous analysts thought that Japanese equities were a bargain. Suppose that agreeing with these analysts, in October 2010 you bought some CME U.S. dollar-denominated Nikkei futures contracts at 9,600. Potentially, that would have been a good decision, as exhibit 2 shows. However, in March of 2011 Japan was hit by a massive tsunami that exacted a very significant human and economic toll, and Japan's stock market fell accordingly. For most of the following two years, the Nikkei traded sideways, at one point reaching a low of 8,130.

For the investor who bought Nikkei futures at 9,600, this represented a loss of 1,470 points, corresponding to \$7,350 per contract. At the time, the Chicago Mercantile Exchange set the initial margin for Nikkei 225 futures at \$1,760 per contract. In all, to convert their good judgment into profits, Nikkei traders would have to endure two years of heavy losses without losing faith in their initial judgment. This is easy to say but actually very hard to do; at its highest, their loss per contract (\$7,350) would be more than four times the initial margin (\$1,760), and for any aggressive trader the loss aversion bias might cause them to discard their good judgment and try to reverse their losses with some clever improvised manoeuvres. That usually makes things worse.

I had the unhappy privilege of experiencing a similar scenario first hand with my boss as the drama's protagonist. A veteran with more than 20 years of experience managing an independent commodity trading business, he was a very sharp man and undeniably a successful trader. In early 2003, as the United States and her allies seemed poised to invade Iraq⁴, he judged that the market had already factored the crisis into the oil price and that the invasion itself would lead to a major price correction. Confident in his judgment, in mid-January 2003 he started taking short positions in IPE Gas Oil futures. Unfortunately, the Gas Oil price kept rising through January and February, causing uncomfortable losses on his positions. This led him to second guess his judgment. As he endeavored to recover his losses, his trading became more frequent and more erratic.

⁴ At the time, the American invasion of Iraq was by no means a foregone conclusion and most of the media treated this as the last resort, worst-case scenario outcome that might yet be avoided, so the conflict's consequences for the oil market were far from clear.

Exhibit 3: Good judgment, dismal results – transactions of an oil trader with 20-years of experience



Arrows show buy (B) and sell (S) transactions; the figures above and below the arrows indicate the total number of contracts bought or sold over the course of the day.

Price averaging, intra-day trading manoeuvres and guessing about the next few days' or hours' price moves only made things worse. Ultimately, although his timing was off, his judgment proved correct: from its March 2003 highs, the price of Gas Oil dropped by 39%. In spite of that, his activity produced a large loss. Over a 14-week period he made 46 different transactions. Rather than profiting through his good judgment, he ended up with a large loss.

Similarly, in early 1995, Jeffrey Vinik, the manager of Fidelity Magellan, then the world's largest mutual fund, got trampled by the markets as the internet technology boom was about to take off. At the time, Vinik held over 40% of the fund's assets in technology stocks, proclaiming that most of his investors "*have invested in the fund for goals that are years away... I think their objectives are the same as mine, and that they believe, as I do, that a long-term approach is best.*" But only six months after he wrote this, Vinik dumped almost all of his technology shares, selling close to \$19 billion worth in two frantic months.⁵ In retrospect, it's clear that Vinik was right on the money with his large allocation to technology companies, but fearing that the already "overvalued" tech stocks were due for a large correction, he deprived his investors of the windfall from one of the most spectacular bull markets ever as NASDAQ soared another 400% (from around 1,000 level to more than 5,000) over the following five years.

At the other end of that same bull market, another star manager made a similar and equally unfortunate mistake. While working for George Soros in 1999, Stanley Druckenmiller accumulated a significant short position in internet stocks which he believed to be stupidly overvalued. He was right, of course, but the Nasdaq's meteoric rise eventually made him blink, cover his shorts and join the bulls on the long side.

Shortly thereafter, the dot-com bubble burst and 75% of the internet stocks Druckenmiller shorted eventually went to zero. The rest of them fell between 90% and 99%.⁶ Instead of making an absolute killing in 2000, Stanley Druckenmiller ended up with the biggest loss in his career.

The deeper mysteries of human psychology

These stories underscore the fact that speculation is a problem of human psychology. A speculator's performance depends on his decisions, and at times, speculative decisions shed light on some of the deeper mysteries of

⁵ J. Zweig in commentary on Ch. 1 of Benjamin Graham's "The Intelligent Investor" (p. 37).

⁶ Price, Tim. "The Emotional Investor" – PFP Wealth Management Newsletter, December 2013. (citing also research by fund manager David McCreadie).

our psychology. We might ask ourselves why such learned and experienced men like Stanley Druckenmiller, Jeffrey Vinik and my boss, ended up going against their (much) better judgment to join the investing herd as it was stampeding toward a cliff? Loss aversion clearly played a role, but other, more obscure aspects of the human psychology also contributed to their conduct.

Our thoughts – the nearly constant stream of awareness that determines our conduct and shapes our self identity – is our sole means of consciously knowing anything at all. This conscious thinking is expressed in language (try to think a thought – any thought – without it being expressed in words; it is almost inconceivable). It is this internal monologue that gives us the experience of what being ourselves *feels* like; it's an independent, totally individual, sovereign experience.

In our own minds, we think our own thoughts, arrive at our own truths, craft our free will and choose our conduct. To an important extent however, this feeling is an illusion. Our thoughts and actions appear to be open to outside influences in ways we can't fully account for. This is not limited to just good advice or some new information: thoughts and decisions can quite literally infect our minds from the outside without our conscious awareness.

Part of the mystery stems from the way our brain is designed. It consists of two hemispheres, each specialized in running a different set of processes. Our left hemisphere specializes in processing language and concepts that can be expressed in language. It articulates our speech and generates our internal monologue that we experience almost constantly during our waking hours. Our right hemisphere is the epicenter of our emotional experience. It has *some* language capability, but is largely nonverbal, processing visual information and managing spatial and personal relationships. The two hemispheres communicate through *corpus callosum*, a bundle of nervous tissue that connects them.

Working in concert, the two hemispheres maintain what we experience as our unified system of awareness. When neurosurgeons began to separate the two hemispheres by severing the corpus callosum – as a way to treat patients with severe epileptic seizures – they discovered that each hemisphere had its own separate systems of attention and action capable of independently influencing a person's conduct. A study of these *split-brain* patients by neuroscientists Michael Gazzaniga and Roger Sperry showed what exactly this means. In their experiments, Gazzaniga and Sperry channelled visual stimuli from one side of a patient's visual field to the opposite hemisphere of the brain.

For example, they showed a funny slide to a patient's right hemisphere (by making it visible only to his left eye). On cue, the patient started

laughing, but when asked why he was laughing, he contrived an explanation that sounded credible but was false. This was evident to the experimenters but not to the patient whose two brain hemispheres couldn't communicate with one another: his left hemisphere, which was articulating his speech – was unaware of the slide that triggered the laughter through the right hemisphere.

In another experiment, when the command “WALK” was flashed to a patient's right hemisphere, he promptly got up and started to walk out of the room. When the experimenter asked him why he just got up, he replied quite sincerely that he wanted to get a drink. Again, the patient's left hemisphere unhesitatingly contrived a credible explanation although it was in the dark as to the real causes of the man's actions. These experiments suggest that our left brain is responsible for producing a sense of coherence and purposefulness of our actions, manufacturing it from whatever ingredients it finds, regardless of whether they are true or invented. What's disturbing about this discovery is that even with intact brains, we can't be sure that our left hemispheres are any more truthful with us about our own conduct. It is our left hemisphere's process that produces the chatter in our conscious awareness.

But our conduct might equally be directed by our “mute” right hemisphere whose influence may be indiscernible to us. Sigmund Freud seems to have understood this when he wrote that often our conscious minds do not control how we act, but merely tell us a story about our actions. In his book, “Escape from Freedom,” Erich Fromm offers another telling example of this same phenomenon at work. Fromm recounts an experiment where a subject is put under hypnosis. During hypnotic sleep, the experimenter suggests to this man that after awakening he will want to read a manuscript which he will believe he has brought with him, that he will seek it and not find it, that he will then believe that another person, Mr. C who is also present, has stolen it, and that he will get very angry at Mr. C. The truth of the situation is that the subject never brought any manuscript and that Mr. C is a person toward whom the subject never had reason to feel any anger. Fromm describes the situation after the subject awakens from hypnosis:

“...after a short conversation with the therapist, he says, ‘Incidentally, this reminds me of something I have written in my manuscript. I shall read it to you.’ He looks around, does not find it, and then turns to C, suggesting that he may have taken it; getting more and more excited when C repudiates the suggestion, he eventually bursts into open anger and directly accuses C of having stolen the manuscript.

He goes even further. He puts forward reasons which should make it plausible that C is the thief. He has heard from others, he says, that C needs the manuscript very badly, that he had good opportunity to take it, and so on. We hear him not only accusing C, but making up numerous 'rationalizations' which should make his accusation appear plausible." Again, the subject of the experiment seems fully convinced that he is thinking his own thoughts and acting on his own inclinations; only the observers who have witnessed the entire episode are aware that the subject was manipulated during hypnosis into believing what never happened: that he brought some manuscript, and that Mr. C stole it.

While his anger also seems to have been planted by the therapist, the subject has clearly injected a narrative of his own: he has supplied the rationalizations about why he just knew C was the culprit, and why he was right to be angry at him."

These experiments suggest that we all appear to have an inner spin-doctor charged with giving us a convincing account of our actions. But this spin-doctor seems to have no scruples about telling us lies, which we "hear" loud and clear while we remain largely deaf to our brain's nonverbal processes that can materially influence our actions.⁷ What does any of this have to do with speculation? Here's what: sustained success at trading depends on the decision making process rooted in rational thinking, independent judgment and disciplined action adhering to some form of strategy. We can only formulate and process these elements verbally, which means through our brain's left hemisphere.

At the same time, our actual conduct could well be influenced by our right hemisphere which is nonverbal. The right hemisphere processes emotion, and in speculative trading emotion can strongly influence our actions. These obscure aspects of our psyche could hold the key to the mystery of why intelligent, successful and disciplined traders at some point abandon their better judgment and take action they rationally understand to be wrong. We can clearly see this in the way Stanley Druckenmiller described his failure managing George Soros's Quantum Fund in 2000. Answering the question about what he thought the biggest mistake of his career was and what he'd learned from it, he said:

"... in 1999 after Yahoo and America Online had already gone up like tenfold, I got the bright idea at Soros to short internet stocks. And I put 200 million in them in about February and by mid-March the 200 million short I had lost \$600 million on, gotten completely beat up and was down

⁷ Some psychologists suggest that we can recognize these processes as a *gut feeling*.

like 15 percent on the year. And I was very proud of the fact that I never had a down year, and I thought well, I'm finished.

So the next thing that happens is I can't remember whether I went to Silicon Valley or I talked to some 22-year old with Asperger's. But whoever it was, they convinced me about this new tech boom that was going to take place. So I went and hired a couple of gun slingers because we only knew about IBM and Hewlett-Packard. I needed Veritas and Verisign. ... So, we hired this guy and we end up on the year – we had been down 15 and we ended up like 35 percent on the year. And the Nasdaq's gone up 400 percent.

So I'll never forget it. January of 2000 I go into Soros's office and I say I'm selling all the tech stocks, selling everything. This is crazy. [unintelligible] This is nuts. Just kind of as I explained earlier, we're going to step aside, wait for the next fat pitch. I didn't fire the two gun slingers. They didn't have enough money to really hurt the fund, but they started making 3 percent a day and I'm out. It is driving me nuts. I mean their little account is like up 50 percent on the year.

I think Quantum was up seven. It's just sitting there. So like around March I could feel it coming. I just – I had to play. I couldn't help myself. And three times during the same week I pick up a – don't do it. Don't do it. Anyway, I pick up the phone finally. I think I missed the top by an hour. I bought \$6 billion worth of tech stocks and in six weeks I had left Soros and I had lost \$3 billion in that one play.

You asked me what I learned. I didn't learn anything. I already knew that I wasn't supposed to do that. I was just an emotional basket case and couldn't help myself. So, maybe I learned not to do it again, but I already knew that.”⁸

Day after day, Stanley Druckenmiller watched technology stocks skyrocket and his younger and much less experienced colleagues make huge returns while his fund was just treading water. What they were doing seemed to be working, and what he was doing wasn't. Day after day the markets were telling him that his “gunslingers” were right and he was wrong; that they were smart and he stupid. Eventually he abandoned his discipline and joined the herd even though in his rational judgment he knew he was doing the wrong thing. “*I was just an emotional basket case and I couldn't help myself,*” says Druckenmiller. Any and every would-be speculator should ponder those words, because what happened to him can happen to every speculator.

⁸ Armour, Timothy. “Stanley Druckenmiller Lost Tree Club 1-18-2015” Transcript, 12 Feb. 2015.

The objective of this discussion is not to suggest that being a successful trader isn't possible, but to point out those parts of our mental machinery that make it difficult for us to be consistently successful as speculators over long stretches of time. Conceivably, we *can* learn to be diligent and rigorous in conducting our research, discerning in our judgment and disciplined in making decisions and there *are* individuals out there who manage to outperform the markets year after year over long periods of time (but for his 2000 debacle, Druckenmiller was one of them). But these individuals are very rare – perhaps the proverbial exceptions to prove the rule. Myself, I did not feel inclined to bet my future on the notion that I might be one of these wizards of the trade. If you choose to make a living by walking a tightrope, keeping perfect balance 99% of the way across a ravine is still not good enough.

If I was going to pursue a career in trading, I had to find a way to sidestep the human shortcomings that could spell my doom. There was only one alternative, and that was to go quantitative and systematic.

Chapter 8: Quantitative modelling

Trained economists have never seen a really first-class model. ... In finance, you're playing against... agents who value assets based on their ephemeral opinions... When you take on other people, you're pretending you can comprehend other pretenders...

Emanuel Derman

Misunderstanding of probability may be the greatest of all impediments to scientific literacy.

Stephen Jay Gould

Most market professionals understand the weaknesses of human psychology at speculation. To overcome these pitfalls, many of them turn to quantitative or algorithmic trading strategies. The potential benefits of quantitative strategies are numerous. For example, they offer a solution to our imperfect knowledge of markets and the impossibility of forecasting asset prices.

Quantitative strategies can also eliminate rogue trader risk by imposing decision making discipline. They can even entirely bypass human action if trades are executed directly through electronic trading platforms. Further, because they are based on mathematical algorithms, back-tests of quantitative strategies over historical market data can give us an objective measure of their expected performance. In addition, by virtue of running on computers, algorithmic strategies are accurate and fast, capable of executing trades in fractions of a second. Also, they can run around the clock without ever losing focus or needing a break.

Not surprisingly, financial industry generated an enormous demand for quantitative analysts, or quants.

Advantages of quantitative strategies	
Solution to imperfect knowledge	Quantitative trading does not require perfect knowledge about markets or correct forecasts. Data, a valid hypothesis about markets, and a working model can suffice to generate trading decisions
Decision making discipline	A well modelled algorithm won't lose composure and insist that it is right when markets go against it; it will simply continue to operate according to the predetermined set of rules.
Objective measure of expected performance	Quantitative strategies can generally be back-tested providing an objective measure of success against which we can measure future performance.
Focus	Algorithms don't get distracted. They remain 100% focused on executing their strategy (for as long as they're plugged in)
Work ethic	Algorithms don't get tired, call in sick, take vacation or ask for a raise.

But in addition to their many advantages, quantitative approaches to speculation involve a considerable set of challenges and risks. These partly stem from the conceptual nature of the problem and from practical difficulties involved with modelling trading algorithms.

Conceptual challenges

In formulating quantitative trading strategies, firms typically rely on mathematicians or physicists who work with ideas and theories borrowed from natural sciences. But while applied mathematics and physics deal with the mechanical properties of natural phenomena, markets reflect the aggregate psychology of their human participants. The difference is very significant. Interaction of inanimate particles or fluids might be sufficiently well understood to make the prediction of certain behaviors possible. By contrast, human conduct doesn't conform to the crisp laws of physics or mathematics. In his book, "My Life as a Quant," physicist and quantitative analyst Emanuel Derman¹ reflects on this point:

In physics, the beauty and elegance of a theory's laws, and the intuition that led to them, is often compelling, and provides a natural starting point from which to proceed to phenomena. In finance, more of a social than a natural science, there are few beautiful theories and virtually no compelling ones, and so we have no choice but to take the

¹ Emanuel Derman had been the chief quantitative analyst at Goldman Sachs for 17 years.

phenomenological approach. There, much more often, one begins with the market's data and calibrates the model's laws to fit...²

What Derman relates is a formidable challenge for quantitative analysts and their employers. Starting with data and working backwards toward a working hypothesis hinges on inventiveness and conceptual thinking in a domain that is complex as well as abstract. Mired in numbers and lacking any tangible concepts to grasp upon, quantitative analysts can easily churn out erroneous hypotheses whose flaws could be very difficult to recognize. In such an environment, strained intellectual exertion can cloud common sense and lead analysts to lose sight of clear thinking. The more abstract the subject matter, the more ways we have to reach mistaken conclusions.

In his bestseller, "How the Mind Works," Steven Pinker cites empirical research that shows just how easily we go off the rails when conceptualizing certain types of problems. For example, psychologists Michael McCloskey, Alfonso Caramazza, and Bert Green asked college students to describe the trajectory of a ball shot out of a curved tube. A minority, but a "depressingly large minority" of students, including many who studied physics, guessed that the ball would continue in a curving path, and were even quite prepared to provide scientific explanation for this.³

Dennis Proffitt and David Gildea asked people simple questions about the motion of spinning tops, wheels rolling down ramps, colliding balls, or solid objects displacing water. They found that even physics professors often got their answers wrong unless they were allowed to fiddle with equations on paper. Pinker notes that cognitive misconceptions run deep, but points out that errors tend to arise from "conscious theorizing." When respondents were shown animated illustrations of their answers, they instantly recognized their errors, usually with a burst of laughter.⁴ But if conscious theorizing can get us lost in problems as simple as the motion of objects in the physical world, how confident should we be about our comprehension of more complex problems?

In "The Language Instinct" Pinker provides an illuminating example from the field of early artificial intelligence research.⁵ In the 1970s and 1980s scientists at some of the leading American universities spent tens of millions of dollars attempting to solve the mystery of language in order to

² Derman, Emanuel "My Life as a Quant: Reflections on Physics and Finance" John Wiley & Sons, Inc., Hoboken, New Jersey, 2004

³ "The object acquires a 'force' or 'momentum' which propels it along the curve until the momentum gets used up and the trajectory straightens out."

⁴ Pinker, S. "How the Mind Works" W. W. Norton and Company, New York 1997 (319, 320).

⁵ Pinker, Steven. "The Language Instinct," Harper Perennial, New York, 1995

enable computers to speak. They based their solutions on the notion that language is a discrete combinatorial system (a finite number of words and a finite number of rules about how to form sentences), and advanced the concept of *word chain device*. Word chain devices would construct sentences by selecting words from different lists (nouns, verbs, prepositions...) based on a set of rules for going from list to list. At the time, some psychologists believed that all human language arose from a huge word chain stored in the brain. In their efforts to generate language artificially, scientists painstakingly calculated the probabilities that certain words would follow certain other words in English language and they built huge databases of words and *transition probabilities*. The following sentence is an actual example of what they got out of all that hard work:

House to ask for is to earn our living by working towards a goal for his team in old New-York was a wonderful place wasn't it even pleasant to talk about and laugh hard when he tells lies he should not tell me the reason why you are is evident."⁶

The whole magic ingredient of meaning never made it into these clever models. It is easy for us to recognize the gibberish flowing out of word chains because our brains were designed to process language and they effortlessly detect the meaning it conveys. What our mind was *not* designed to do is process mountains of quantitative data. In this domain, we are not equipped to easily discern sense from nonsense and this can lead us to blindly pursue flawed hypotheses. There are many ways we can misconstrue amorphous reams of data. As we already discussed in Chapter 5, we are susceptible to confusing correlation with causation. If some observation B follows the observation A 90% of the time, we tend to assume that there's a 90% probability that B will follow the next occurrence of A.

In complex domains, this is often not the case. Whatever we are capable of reading out of the market data, the figures can only represent a very limited manifestation of the vastly more complex system, and establishing any kind of causal relationship in the data is bound to be a stretch. We also have great difficulties interpreting probabilities, and this also includes the experts.

Consider the following example: at Harvard Medical School, researchers posed a problem to 60 students and members of the faculty. The problem read as follows: a test to detect a disease that afflicts one

⁶ This word-chain model worked by estimating the most likely word to follow after each four-word sequence, growing the sentence word by word.

person in a thousand has a 5% false positive result. What is the probability that a person found to be positive actually has the disease, assuming that you know nothing about their symptoms? The correct answer to this problem is 0.02. The most popular answer was 0.95 and the average answer was 0.56. Among the experts in this group, fewer than one in five got the right answer.⁷ To be fair, we tend to do much better when problems are presented in terms of *relative frequencies* rather than mathematical probabilities.

As many as 92% of respondents gave the correct answer when the problem was formulated as follows: in a given population, one person in a thousand has a disease and 50 of 1000 test positive. How many who test positive actually have the disease? The difference between the two formulations is subtle, but it goes to show that we may often fail to grasp the substance of complex quantitative problems and that even experts aren't immune to misconstruing mathematical probabilities and arriving at wrong solutions. In quantitative analysis of markets, these issues are highly relevant and represent an important source of risk.

Model risk

All things excellent are as difficult as they are rare.

Baruch Spinoza

Even supposing that we have done a good job analyzing the data and that we reached a valid hypothesis, we still face another daunting challenge: making sure that our models correctly fulfil their intended purposes. This problem spills into the domain of software programming.

Models are normally implemented in software programs that may require thousands of lines of code, large databases and a suitable user interface. Creating such programs involves its own peculiar set of risks which only rarely receive adequate attention. Software code is seldom free of errors, which are often extremely difficult to identify before they cause an adverse outcome.⁸

⁷ Pinker, Steven. "How the Mind Works" W.W. Norton and Company, New York 1997 (344)

⁸ On average, professional programmers make as many as 100 to 150 errors per 1000 lines of code. This is according to a multiyear study of 13,000 programs by Watts S. Humphrey of Carnegie Mellon University's Software Engineering Institute.

If you pay attention to the daily news flow, you'll notice countless examples of model/software issues that result in serious setbacks. Here are a few examples:

- On the first day of October 2013, the U.S. administration under President Barack Obama launched the much anticipated government medical insurance market and its website, *Healthcare.gov*. The government spent some \$600 million developing the website which turned out to be such an unmitigated disaster that fully ten days after launch, not a single person could be confirmed to have successfully enrolled.
- In March of 2013, UK intelligence agency MI5 reportedly scrapped a major IT project to centralize the agency's data stores. The work became such a morass that MI5's director at the time, Sir Jonathan Evans decided to abandon the project altogether and restart from scratch with a completely new team of IT professionals. According to *The Independent*, the abandonment of the project cost MI5 about \$140 million.
- In late 1999, the Mars Climate Orbiter crashed into Mars because an engineer at the Jet Propulsion Laboratories failed to convert British measurement units to the metric system.
- Shortly afterwards, a sister space vehicle, the Mars Polar Lander, also smashed into Mars because the line of software code that was supposed to trigger the vehicle's braking process was missing.
- In 1996, the European space probe Ariane 5 disintegrated 40 seconds after launch due to an error in the computer program controlling the rocket's engines.

The list is long and interesting, including issues with motor vehicles, advanced military hardware and software, communication and navigation technology, medical diagnosing and treatment systems and just about every other kind of technology that uses computer software to function. In the financial industry, software errors don't cause things to blow up, so they can remain hidden or even go undetected for a long time. However, every now and again things get bad enough to attract some publicity.

On August 1, 2012, New York brokerage Knight Capital implemented a trading algorithm that in a very short time caused the firm a direct cash loss of \$440 million and a market cap loss of about \$1 billion. The faulty algorithm bought securities at the offering price and sold them at the bid, and continued to do this some 40 times per second. Over about thirty minutes' time, the algorithm wiped out four years worth of profits. In another example, in June 2010, an international bank's algorithmic trading

system acted on bad pricing inputs by placing 7,468 orders to sell Nikkei 225 futures contracts on the Osaka Stock Exchange. The total cost was more than \$182 million. While the pricing error would have been rather obvious to any human participant, the trading algorithm proceeded to execute approximately \$546 million of the orders before the error was caught. These two quantitative trading debacles are not isolated stories. I believe that model risk events are pervasive, but the vast majority of them remain unknown outside of the firms that experience them.

Over the years I have personally come across a good many cases where an important part of a firm's business process got bogged down due to poorly designed software tools. In each of these cases, frustration with the software dragged on for years and I am not aware of even a single case where the issues were resolved in a satisfactory way. The usual course is eventually to abandon the software tools and return to the old manual process. The main reason these things happen is the lack of appreciation on the part of decision makers of just how difficult it is to build, implement and maintain well-functioning software.

What typically happens is that quants and/or software programmers are hired and simply expected to produce quality tools. Outside the software industry itself, people take best practices in software engineering⁹ lightly, if they are even aware of them, and their approach to building models tends to go straight from half-baked ideas to programming.

With few exceptions, the outcome falls well short of the desired results. In the aftermath of Knight Capital's trading model blow-up, the firm's CEO Thomas Joyce rather flippantly declared on Bloomberg TV that, "*if you get involved in the day-to-day minutia, this will give you a headache occasionally.*" I agree with Mr. Joyce on that, but if you venture to bet money on a trading algorithm, enduring some headaches could prove to be your best investment of time and effort. In any endeavor where performance substantially depends on models, it pays to be thorough.

I've come to believe very strongly that software quality is a strategic issue of the first order in quantitative trading and asset management in general. Sooner or later, failure to adequately manage model risk is likely to have a very meaningful adverse impact on performance. Another important aspect of quantitative modelling involves organizational issues. This is particularly the case in larger organizations where quantitative analysis functions are separate from, and subordinate to the key decision

⁹ There's an important difference between software engineering and software programming. A software programmer is to a software engineer as a construction worker is to an architect. In the financial industry, quants are usually reasonably competent software programmers, but many of them have little awareness of what software engineering entails.

making functions. Particularly in organizations run by clubby management cliques, decisions are frequently based on influence, authority, or group allegiance rather than on a clear-minded analysis of ideas and facts. In such organizations, quality ideas are less likely to be recognized and given support. This is a weakness of many large organizations, even if it isn't directly apparent to outside observers. At times however, we can get a glimpse of them indirectly.

One example came to my attention in 2007 with the growth in popularity of the so-called 130/30 funds, or short extension funds, which were predominantly managed by quantitative managers. A 130/30 fund balances 130% long exposure with 30% short exposure in capital markets. The intent of these vehicles was to outperform traditional equity benchmarks, especially in falling markets.

By 2007 it became clear that most of these funds by far fell short of expectations. When the market fell in the summer of 2007, short extension funds managed by large organizations like State Street Global Advisors, Barclays Global Investors, Goldman Sachs Asset Management, Deutsche Asset Management, JPMorgan Chase, Charles Schwab and ING all left investors with bigger losses than the S&P 500 index. According to Morningstar, only three of the 38 short extension vehicles did better than the S&P 500. They equally disappointed through the aftermath of the 2008 financial crisis. Morningstar reported in April 2009 that 130/30 strategies on average lost 43.1%, compared to a 40.9% drop for long-only funds.

In April 2010, AXA Rosenberg Group, which at one time managed over \$70 billion, told clients that a coding error had affected its computer-driven investment process. Though the error was discovered and corrected in 2009, the company claimed that “high-level investment personnel had kept the problem under wraps”. The direct effect of the coding error was major underperformance of the fund compared to its peers. As a result, in 2011 the SEC handed AXA Rosenberg Group a record fine of \$242 million.

In spite of all these challenges, I much prefer quantitative over discretionary trading so long as the users of quantitative strategies remain mindful of the main issues and risks inherent in this approach to speculation.

As with so many things in life, achieving success with quantitative strategies is a struggle against the odds. To prevail, practitioners must start with clear thinking and proceed with meticulous and disciplined adherence to best practices in systems engineering. Learning about systems engineering, or hiring individuals with this skill set should prove a very worthwhile investment.

Chapter 9: Speculation in the wild

Look deep into nature, and then you will understand everything better.

Albert Einstein

My own experience with discretionary trading reinforced my preference for the quantitative approach. Although I believed I was a fairly good market analyst, I had next to no confidence in my ability to outsmart the markets and little desire to try. The many potential advantages of systematic trading based on trend following seemed much more promising as a way forward. Flanked with a small team of capable programmers and mathematicians, I was in a good position to advance in this direction. Initially I expected that trend following algorithms shouldn't be difficult to formulate, so we set out thinking up trading systems, writing algorithms and back-testing them.

This was an interesting and moderately stimulating experience but I soon found myself sceptical about our initial successes. We'd formulate a set of studies and rules that generated buy and sell decisions, optimized our parameters, and in a day or two we had a systematic strategy that looked like something that might make money. But in fact, all we got through this exercise were strategies that *would have* made money in the past. We had no way of knowing how they might do in the future. This is the inescapable aspect of uncertainty: whatever approach to speculation you adopt, you are always basing your decisions on what you learned in the past, while the results of your actions will depend on what happens in the future.

I was unsure what to make of this problem at first. Even supposing that our strategies turned out profitable in live trading, I expected that they would still eventually fail. How would we recognize the moment in time when a trading strategy was beginning to fail? And if we could recognize that moment, what would we do then? Return to the drawing board and formulate a new-and-improved strategy? That didn't seem like a solution. Quantitative modelling is a fairly error-prone business and I felt that having to repeatedly reinvent the wheel could compound the risks and

uncertainties inherent in this work. I worried that we could spend years stuck in a fruitless loop like a dog chasing his tail.

Truth be told, part of the problem was also my own ignorance. There is a way to go about building trading systems and experienced managers routinely use trading systems based on backtested results. At the time, best practices in systematic trading were fairly well established in the hedge fund world where trend following had already been used for at least three decades, and a good deal of literature was available if you knew where to look. But sitting in a small commodity trading firm in Monaco, we were in the wrong “silo” so to speak, and unaware of much of this.

We worked essentially from a blank sheet of paper, learning through trial and error. This may have been a blessing as well as a disadvantage as it obliged us to work out our own, original approach. Through that endeavor, I felt that there had to be a way for us to formulate a sustainable solution to the problem of speculation – some kind of a model that could indefinitely navigate the unpredictable market price fluctuations and sail along with trends to generate trading profits with some consistency. For a time I became obsessed with finding such a solution, and it was in one of those mysterious flashes of inspiration that the model which I chose to pursue emerged whole and with clarity.

Nature's speculation

Life can only be understood backwards; it must be lived forwards

Søren Kierkegaard

One evening while enjoying a wildlife documentary program, it occurred to me that if there were a sustainable solution to the problem of uncertainty and risk, it would have been worked out in some form in nature. Upon reflection, I realized that every form of life on Earth is in essence an embodiment of a strategy of survival. In natural life, species compete for energy and resources. Every individual animal is endowed with a physical body (repository of inner resources) and a set of behaviors whose primary objective is to enable that animal's survival and procreation. To produce offspring, the animal must take in more resources than it expends in the course of living; its activities have to be profitable in terms of sustenance, else it would perish. The existence of each species is

proof positive that its survival strategies are successful. Take spiders for instance. Their strategy is to build webs. A spider's body is designed to do this. She may not know that food will get caught in her net, but this is how she secures her nourishment which must be sufficiently abundant to recover the resources that went into the building and maintenance of her web and also to bring forth her offspring.

I further realized that nature faces uncertainty in a similar way that we did with our quantitative trading strategies. Namely, nature generates her models without knowing how long they would be viable. The design of a species is based on the environment experienced through its evolutionary past and every life form tends to be adapted to its present familiar habitat. However, habitats eventually change and species must adapt or go extinct. When we consider that over 90% of all species that have inhabited the Earth ultimately went extinct, it becomes clear that all of nature's models are fallible and that their fallibility makes part of life's design. In this sense, nature's designs are speculative and every model is a guess based on the known environment. Indeed, nature does not sustain life through permanent, immutable models, but by making its models flexible, constantly generating new adaptations and new species that can thrive for a time even as others go extinct. This process also doesn't tend to reinvent the proverbial wheel: old solutions are continuously reused as components to assemble new solutions¹.

Consider mitochondria. They are among the most important and oldest building-blocks of all multi-cellular life on Earth. These organelles live in nearly all eukaryotic cells² where they produce the chemical energy needed for the cell's metabolism, division and motility. But mitochondria predate the life forms they power: they evolved some two billion years ago from bacteria which entered into symbiotic relationship with early eukaryotic cells. Mitochondria have their own cell membranes and DNA which is distinct from the DNA of the organisms they help keep alive.

Nature's more recent invention are neurons which evolved as messengers in more complex organisms. Life forms like plants or sponges use chemical messaging. For example, when a giraffe eats leaves off of an Acacia tree, the tree releases signalling chemicals through its branches, which triggers the production of bitter-tasting and toxic tannins that deter giraffes. The disadvantage of chemical communication is that it is slow, so a giraffe can inflict some damage on a tree before its defensive mechanisms are fully armed. In the course of evolution, some species

¹ This is why humans share about 7% of genetic material with E. coli bacteria, 90% with mice and 98% with chimpanzees.

² Eukaryotic cells are those that contain a distinct membrane-bound nucleus.

evolved cells that could generate electrochemical spikes, enabling signals to travel much faster. Ultimately, such cells evolved to become neurons, specialized in messaging and information processing in organisms which could now respond to external stimuli almost instantly and in more intelligent and varied ways. Nature's inventions like mitochondria, neurons, chlorophyll, eyes, vertebrae, and countless others subsequently gave rise to a vast variety of new life forms which incorporated these solutions as building blocks in their own design.

As far as we can tell, nature does not have its chosen species. Likely, it is indifferent to the relative success of any individual model it creates, but the same is not true for the models themselves. Each species strives to maximize its reproductive success. Its growth is checked only by limited resources and competition from other species. Depending on how you choose to define it, the most successful models are those that manage to achieve the greatest longevity or greatest biomass. In terms of biomass, humans and termites are among the most successful species – for the time being, that is³. In terms of longevity, termites could turn out to be more robust – they've been around for some 150 million years. We only joined the party less than 200,000 years ago and for all we know, we could end up extinct before another 200,000 years lapse – the blink of an eye in evolutionary terms.

Speculation in the life of predators

Nature has ... some sort of arithmetical-geometrical coordinate system, because nature has all kinds of models. What we experience of nature is in models, and all of nature's models are so beautiful.

R. Buckminster Fuller

If we zoom in from the level of life on Earth to the level of individual agents, the survival strategies in nature that bear the most similarities to the activities of market speculators are those of predators. To live, predators must hunt – an activity that includes elements of speculation. Like trading, predation requires knowledge, skills, judgment and decision-making. It also entails risk and uncertainty. A predator can't be sure where

³ Of course, having a very large biomass means that you represent the most abundant potential food source for species that figure out how to crash your defenses and use you as lunch.

its next meal is coming from. Each hunt is an investment of resources; it involves the risk of injury and loss of energy expended in failed hunts, which tend to be more frequent than successful ones. To survive and procreate, predators must consistently generate a positive return on this investment. Too much of a losing streak could turn out to be fatal.

When pondering these issues, I tended to envisage the large cats hunting on the African savannahs and got quite excited when one day I came across a book titled, “The Serengeti Lion: A Study of Predator-Prey Relations” by George B. Schaller. Schaller spent several years in the Serengeti National Park in Tanzania during the 1970s, observing the activities of lions and other predators and fastidiously recording the details of hundreds of hunts. We have all seen wildlife television programs showing lions and cheetahs hunting, but Schaller’s work offers a much richer account of the life of predatory cats including their hunting behavior, which I condensed in the following section.

The anatomy of a hunt

Lions prefer to hunt at night, especially when the moon is not bright. Because most of the animals they hunt can easily outrun them, lions must take every advantage of external factors like darkness, dense vegetation or the vicinity of water. While hunting, they rely on sight, hearing, and smell in the order of decreasing importance. Lions see much potential prey in the course of a day and evaluate the likelihood of catching any that appear vulnerable. “Most are given a glance,” writes Schaller, “some merit a closer look, a few elicit hunting movements, and only a very few are actually pursued.”

Lions use several distinct methods of hunting, which include ambushes, drives, runs and stalks. On occasion, lions make unexpected kills when a sick or injured animal stumbles upon them. The most common strategy is stalking, where lions attempt to approach their prey undetected. To conserve energy, lions are extremely selective about engaging in the actual chase and generally don’t charge unless they’ve been able to approach their prey undetected to within about 30 meters or less. The decision to attack also depends on the lion’s judgment of her own fitness as well as that of the prey: chases after young animals are generally longer than those in pursuit of adults. If a chase is failing, the lion is quick to abandon the attempt and only seldom pursues the prey for more than 200 meters.

The risk of injury is another important concern. To avoid violent impact, prey is almost never attacked from the front, and when making a kill, a lion is careful to position her body where its victim’s horns or

thrashing hooves cannot reach her. Still, accidents do happen and Schaller reports seeing lions with broken jaws on several occasions. Such an injury is usually fatal for the predator. A lions' success at hunting depends on a variety of environmental factors and the method of hunting. Overall, running by a single lion is successful only about 8% of the time. When stalking or ambushing, a single lion kills on about one in six attempts, but if two lions hunt together they succeed once in about three hunts. Clearly, even though most of lions' hunts will fail, their success rates are sufficient for them to survive and procreate.

Decision-making in predators and speculators

One component of a predator's hunting that we can not observe, but which is clearly operative in every healthy animal's brain, is the decision-making process that directs her predatory behavior. This is a sophisticated and highly complex mechanism, but for our present interest, I'll only discuss those elements that parallel the speculative activities of traders. As we have seen, lions pass much time watching their environment for an opportunity to catch prey.

When actively hunting, a lion keeps track of a variety of factors to determine when to launch an attack. The size of her prey must be large enough to justify the expenditure of energy, but must also not be too large for her to tackle safely. She must also make a judgment about an animal's state of fitness and focus on the most vulnerable individuals.

She must also take her own fitness, speed and endurance into account, as well as a myriad of environmental factors. She may only charge when she is highly confident that her hunt can be successful. At that point the decision to launch the attack is made and she charges with full force. Her decision-making doesn't stop there however; the lion must conserve energy and abort her hunt as soon as her confidence in making a successful kill drops below some threshold. Then the process starts over.

In terms of decision-making, a lion's predatory behavior is similar to a trader's speculative behavior. The speculator spends much of his day scanning news, analyses and commentary about securities markets in order to identify attractive investment opportunities. Some opportunities or trade ideas may catch his attention and he then studies them more closely. When he is very confident that he can make a profitable trade of it, he buys or sells some quantity of the asset in question and assumes the risk in holding it. From that point on, he monitors his position to make sure it's unfolding as expected. But at this stage of the hunt, the behavior of speculative traders differs sharply from that of nature's predators. Predators are masters of conserving resources and cutting their losses. They can always

afford to abandon failed attempts because their survival depends on the cumulative result of the total of their hunts rather than on the outcome of any individual attempt. By contrast, speculators tend to treat each transaction as a departure from the status quo and are burdened with a hardwired loss aversion bias. If markets go against them, rather than cutting their losses, traders tend to gamble with them and escalate risk hoping that things will turn in their favor. This doesn't always happen, and most speculators end up losing. Many squander all of the resources at their disposal and eliminate themselves from the pool of market participants.

Nature's risk management

Looking deep into nature also gave us the solution to the problem of risk. Risk is not the same thing as uncertainty. Uncertainty means that we simply cannot predict the future. Uncertainty also can't be quantified in a meaningful way. By contrast, risk can be quantified and measured. In simplest terms, risk tells us how much we can lose if we bet the wrong way. If we make small bets, we risk small losses and if we make large bets, we can lose big. Nature has resolved the problem of risk to life on Earth through fragmentation of risk and diversification of species and individual agents.

This principle is appropriately encapsulated in the maxim, "no tree grows to the sky." While every species strives to grow, this is not done by infinite growth of individuals but by their multiplication at a certain – probably optimal – size. Thus, lions grow to about 115 kg for females and about 180 kg for males. If they are successful as predators, they will raise many litters of cubs. When fully grown, younger lions will establish new prides and spread as widely across the Earth's surface as they can. Competition for habitat and resources is the main business of every species on Earth and their action has spread life throughout the biosphere.

This diversification of life and its constant renewal as mature generations beget young generations has also enabled life to be perpetually adaptable. As conditions in a habitat change, life adapts by varying the genetic expression of species in their successive generations. Thus, even with perishable individuals and extinguishable species, nature has been able to sustain life for over three billion years and will probably continue to do so indefinitely as long as the conditions on the planet allow it.

The challenges encountered by natural life seem compatible to those we must address in financial speculation. For me, this realization made the idea of emulating nature to build a sustainable solution to the problem of speculation irresistibly compelling as it gave us coherent answers to the problems of uncertainty, risk, growth and adaptability. In this sense, we

could tackle the problem of uncertainty at the level of individual autonomous agents, which for our purposes would consist of a variety of systematic trend following strategies. Each strategy would come equipped with a set of rules determining its speculative conduct and a “physical resource” – an amount of money with which to take a predetermined quantity of risk. Using a multitude of such strategies would enable us to supplant the uncertainty of market price fluctuations with a more predictable risk class. Risk could be controlled by dividing the investment portfolio among a large number of such strategies, each in charge of a small fragment of the total portfolio risk.

Thus, if any one strategy failed, performance could still be sustained by the positive results of other strategies. The growth of an investment portfolio would be based on continuous addition of individual trading strategies. This would also introduce a degree of adaptability to the portfolios as new trading strategies would always be “educated” on the historical price data up to the present.

Although this approach seemed impossibly ambitious at the outset, its parallels with natural life struck me as compelling and beautiful, which made building a model based on these ideas the most stimulating objective my team and I could possibly set for ourselves. From the time we formulated this objective conceptually, it took us less than four months to build a functioning prototype of the core model.

Chapter 10: Building the I-System

What's needed is a sound intellectual framework for making decisions and the ability to keep emotions from corroding that framework.

Warren Buffet

Remember, all decisions are made on the basis of models. The assumptions in a person's head are not actual systems, but assumptions about actual systems. You do not have a family or city or corporation in your head. You have mental models — often poorly and incompletely defined models — of these real-life systems. The heart of the matter is your relative degree of confidence in each of these models.

Jay W. Forrester

Our quest for a sustainable solution to the problem of uncertainty was guided by the insights discussed in the previous chapters of this book, most importantly that:

- markets are complex systems that defy our ability to identify cause-and-effect relationships between observed phenomena
- available information is inadequate to construct a true picture of the economy
- analyzing economic fundamentals doesn't enable us to predict future outcomes
- prices may be the only true and timely form of market information
- price history gives us valid means to analyze and interpret a market
- prices move in trends and trends can be exploited profitably
- natural world offers compelling solutions to the problems of risk and uncertainty

Outlining the solution space

Around mid-1999 my collaboration with Dr. Gorazd Medić intensified as we sensed that we were perhaps onto something interesting. We spent the whole summer of that year working close to 16-hour days, seven days a

week, making progress at a pace that at times seemed like magic. Most days we would begin by defining the problem to be tackled that day.

Usually we had no idea how we would solve it, but almost without exception we'd have a working solution by the end of the day, leaving the office exhausted but exuberant at having added to our model something we didn't know how to do the day before. I must say that much credit for this is due to Gorazd's relentless determination at problem solving.

Many days, he'd continue hacking away late into the evening until he was satisfied with the results, often long after my batteries went dead for the day. The objective of our efforts was to build a sustainable solution to the problem of uncertainty based on principles we encounter in the natural world. To do this, we needed to outline a suitable solution space: a sort of knowledge framework within which we could formulate a multitude of different, autonomous trend-following strategies without having to reinvent the wheel with each new strategy.

Our knowledge framework needed to include those elements of chart analysis that enabled a competent analyst or trader to form a judgment about price fluctuations, identify trading opportunities and take decisions in the face of uncertainty. We named our model I-System, with "I" entailing three distinct but related meanings: (1) I as intelligent; (2) I (eye) as the visual organ; and (3) I as the *self*, defined under the computational theory of the mind as the executive (decision-making) process which directs a person's behavior.


Price charts – the external environment

The first element of chart analysis is the chart itself: it is a time series consisting of market prices. This element reflects the outside environment that we can observe and in which we ultimately transact our trades. With regard to futures markets we had to address a small complication in constructing historical price charts. Unlike stocks or bonds, futures contracts have an expiration date, past which we can no longer trade them on the futures exchange. At that point, we must either roll our positions out of the expiring contract and into the next one, or we must transact the actual physical commodity¹.

To construct a long term historical price chart of a futures market, we have to join together a sequence of futures contracts. By default, these so-called continuation price charts are constructed by adjoining contracts upon expiry: the price of the current contract is plotted on the chart until

¹ If our position after the contract expiry is *long*, we must accept delivery of the specified quantity of the commodity in question; if our position is *short*, we must supply such quantity of commodity to a designated receiver.

its last trading day, after which the price quotations for the next contract are plotted in continuity. This is adequate for visual analysis, but such a chart does not accurately reflect the prices we would actually trade.




Namely, as a contract nears its expiration date, the trading volume and open interest begin to decline sharply at some point and may become quite thin during the last few days of trading. As open interest and volumes thin out, the bid-ask spreads tend to widen and traders find it more difficult to trade out of their positions. For this reason, traders prefer to roll out of expiring contracts well in advance of their last trading day. That in turn means that the price curve we would effectively be trading won't exactly match the continuation chart constructed on contract expiry. The difference might appear small visually, but it could be a significant consideration in formulating systematic trading strategies.

We determine a strategy's effectiveness through backtesting, and it is critical that the backtest simulation correspond as precisely as possible to the way we would actually trade in a given market. Other wise, simulation results could be unrealistic and provide a distorted indication of future trading performance. This was one of the lessons we had to learn the hard way. At first we assumed that the difference between the default continuation charts and the price curve we'd actually be trading was

negligible and wouldn't meaningfully affect the validity of our backtests. We were wrong on this account but we resolved this problem in 2006 with a new version of our model which allowed us to construct price charts by joining successive contracts at the point in time when we would actually execute the roll-over trade.

Because roll-overs can involve complicated calendar algorithms, we simplified our solution by defining for each contract the calendar day when we would retire the contract and roll our position to the next one. We would select that day arbitrarily to precede the point in time at which the expiring contract's volume and open interest started to drop sharply. Exhibit 2 offers an illustration.

Exhibit 2: rolling over futures contracts



Financial and commodity futures markets trade in a number of delivery months. Some, like Brent Crude futures and other NYMEX energy markets trade in all 12 monthly deliveries (FGHJKMNOPQUVXZ). Others trade in only four or five. CME currency futures trade in March, June, September and December (HMUZ) deliveries, as in the above example.

With this small adjustment to our model we were now sure that the external environment on which we would base our trading strategies would accurately reflect the environment in which we would be trading.

Recognizing price trends

The next problem to resolve was significantly more complex, but also infinitely more interesting: how do we recognize price trends? For a

human analyst, the task is relatively easy – our brain is equipped with sophisticated organs of visual analysis and algorithms that do the job effortlessly. But to build a computer model that could do the same thing seemed difficult in the extreme. Part of the problem is that we can't describe a trend with a crisp definition that could distinguish all trends from all non-trends. This is because a trend is a *fuzzy* concept where some patterns look more like trends than others. Modelling fuzzy concepts requires fuzzy logic and neural networks.

Fuzzy logic is a multivalent alternative to conventional logic. Conventional logic is bivalent. It allows for statements to be either true or false. Fuzzy logic allows for degrees of truth, providing a better approximation of the way human cognition works. Consider for example the statement, "Mary is old." By conventional logic, Mary can either be old, or not old. To decide, we must adopt a convention such as, "all persons aged 70 or more are old." In other words, every person of that age group is a member of the "old people" set. Mary's membership in that set depends on where her age falls relative to the convention. If she is 70, "Mary is old," is true. If she is 69, she is not old. This logic isn't hard to grasp, but that's not how we decide such issues in reality.

Our mental representation for old people is a fuzzy set, where membership is a matter of degree – if she is 80, "Mary is old," is more true than if she is 60. Fuzzy logic is an essential ingredient of intelligent decision making as it is innate to the way intelligent systems work. Indeed, most of the categories that make up our mental representation of the world are fuzzy sets, where some objects and events are regarded as better examples of a category than others. Similarly, certain patterns in a price chart will look more like trends than others. For our purposes, the key criterion we were interested in was a system's trading performance: a useful test of a trend needed to identify the price moves we could exploit profitably. Adopting fuzzy logic to resolve the problem of recognizing price trends proved to be the critical conceptual breakthrough we needed. Still, we struggled to design a model that would systematically calculate the actual solution to this problem. Recognizing patterns that looked like trends in the past was not very useful for our purposes.

Trends are obvious in hindsight, but decisions have to be made in the present. So, for every point in time, we needed to have a systematic answer to the question: do current price fluctuations constitute a trend? Most of the time, this question cannot be answered with the simple *yes* or *no*. Instead, we can only have a judgment expressed with some degree of confidence: if a trend is very obvious, we might be certain in our judgment, but on most occasions, our confidence will fluctuate between an utter lack of conviction and absolute certainty. In mathematical terms, we

could assign those levels of confidence 0 and 1, with zero representing ambivalence and 1 denoting certainty. With this conceptual solution outlined, we proceeded to design the actual model that would generate trend confidence values in response to market price fluctuations.

Gorazd and I sat down and formulated a set of tentative trend definitions. We came up with five around which we could write mathematical algorithms. The simplest one relied on a moving average (MAV): the trend is up when the price is above the moving average and down when it is below the MAV. Using two different MAVs at the same time allowed us another way to evaluate a trend. The third definition involved relative positions of local extremes in the price charts: successively higher peaks and higher troughs indicate an uptrend, and successively lower peaks and lower troughs indicate a downtrend. As technical analysis makes much use of trendlines, our fourth definition involved the use of lines projected through local price extremes. The slope of trendlines and their relationship to the price would give us an additional way to judge the prevailing trend.

Six different trend evaluation processes	
Single moving average	The relationship between the current price and an N-periods moving average (MAV) could give us a way to recognize the current trend.
Two moving averages	The relationship between the current price, a shorter-term MAV and a longer-term MAV offers another way to evaluate trend.
Local extremes	In uptrends, we have successively higher peaks and higher troughs; in downtrends, we have successively lower troughs and lower peaks. Analyzing the relative positions of these local extremes was another way to evaluate the current trend.
Trendlines	Chart analysis heavily relies on the use of trendlines to establish trends; the way human analysts project trendlines through price extremes could be approximated in a mathematical algorithm to have another method of judging price trends.
Global extremes	When prices reach new long-term highs or lows, we may have near-certainty that we are observing an uptrend or downtrend.
Global experts	Periodic surveying of market experts or economists (or even non-expert participants) could give us an additional way to evaluate trends.

Finally, long-term price extremes were another way to evaluate trends: if the current price is the highest (or lowest) price attained over the previous year or two or five, then we were looking at an uptrend (or downtrend).

We also envisioned another, sixth trend evaluation process that would depend on human expertise; namely, we could survey a number of market analysts, economists, or even non-expert participants to simply tell us what they thought the price trend was in a given market and how strongly they felt about their judgment, and incorporate their expertise in our system's final trend judgment.

With the five price-based trend definitions written down on paper, we still found it far from obvious how to actually build our algorithms. For example, calculating an N-period MAV for a set of data is simple enough, but that only gives us two values to work with at each point in time: the MAV and the current price. Calculating trend confidence involved *interpreting* the relationship between them. In interpreting various values and relationships derived from a time series, it is particularly important to apply clear thinking and proper understanding of the problem.

In developing novel solutions, small conceptual errors can easily lead off course with some tendency toward overcomplicating simple things. We succumbed to this already with the moving average algorithm. For example if the current price is equal to the MAV, perhaps we have no trend; if the price is some small distance above the MAV, presumably we have a weak trend; if that distance is large we might have a strong trend, etc. The slope of the MAV curve could further qualify the trend judgment, as well as the rate of change of that slope. Also, at times when price fluctuates closely around the MAV, frequent crisscrossing of the MAV and the price curve tends to produce a series of losing trades, so we also thought about ways of adjusting the trend judgment during such periods, or filtering them out somehow.

Although all these ideas were technically feasible, we found ourselves drifting deeper and deeper into the complexity of trying to intelligently interpret the simple relationship between two bits of data, and this approach didn't quite feel right. The problem, as we eventually realized, was that we were trying to write intelligence into algorithms where intelligence didn't belong. Our breakthrough came with the realization that intelligence couldn't be packed into one super-algorithm. Instead, intelligent behaviour might emerge from the interaction of a set of simpler algorithms, each specialized in carrying out a limited task. This, in essence, is what we call expert systems. Such systems consist of a number of simpler algorithms or *experts*, each specialized at solving a part of the problem efficiently and reliably.

Through an ordered interaction of such algorithms, expert systems can produce intelligent solutions to complex problems, which is the way intelligence emerges in nature. Consider the following illustration involving my dog, Maya. One day, while walking her along a mountain

path I was tossing her a tennis ball for play. At one point, the ball bounced sideways and dropped off the path into a deep ravine. Fearing that Maya might jump after the ball into the ravine, a moment of terror overcame me, but of course she chased the ball to the edge of the path and then stopped to decide what to do next. As in any intelligent creature, Maya's brain can direct a range of behaviors. One of them is the chasing instinct: *if a small object is speeding away from me, I chase after it and try to grab it*. Maya also has self-preservation instincts which tell her: *if I come to a deep ravine, I stop*. Now, if all this behavior had to be managed by a single algorithm, it would have to be very complex, constantly accounting for a huge range of possible occurrences at every step. This is unlikely to be the way intelligence works in nature.

Rather, natural intelligence stems from a multitude of different algorithms, each capable of engaging the clutch of the creature's behavior depending upon the circumstances. If a ball is bouncing away from Maya on a plain flat field with no perceptible risk to her safety, the self-preservation instincts might remain on stand-by and she could just focus on chasing the ball. As soon as any potential danger is perceived, self-preservation instincts can override the chasing instincts and the animal aborts the chase to reassess how to proceed.

Still other algorithms may alert her to thirst or tiredness or the presence of other animals, fire, unusual smells or noises and so on. A single linear algorithm keeping track of all these possibilities at all times would most probably be hopelessly cumbersome and complex. Instead, my dog's intelligence must be the result of an ordered interaction among many simpler algorithms which in turn consist of even more basic algorithms in a hierarchical system of experts, each in charge of a simple task enabling Maya to respond intelligently to unpredictable events in her surroundings.


With respect to our problem of building an intelligent trend recognition model, nature's solutions yet again led me to the conclusion that we had to resist the temptation to cram too much intelligence and complexity into individual algorithms. Instead, we needed to allow our experts to be dumb: efficient and reliable at solving one limited problem and letting other parts of the system take care of the rest. In this way, intelligence might emerge out of the interaction of many different problem-solving parts. In the context of our MAV vs. the current price problem, there were only three distinct possibilities at any moment in time: the price could be above the MAV, equal to the MAV, or below the MAV.

Limiting ourselves to only these three possibilities made the problem of interpreting these relationships very straightforward: if the price is above the MAV, we were likely looking at an uptrend. If it was below, the trend was likely down. If the MAV and the price were the same, we could

either say that our trend confidence was zero, or we could leave it unchanged from its previous value (i.e. the same value as with the preceding price quotation).

Assigning numerical values to these relationships was a matter of convention. In terms of the continuum between -1 (certainty of a downtrend) and 1 (certainty of an uptrend), the MAV algorithm could only give us a limited degree of confidence about the price trend, so the values had to fall short of certainty at 1 and -1. We decided to assign trend confidence the value of 0.5 when the current price was above the MAV, and -0.5 if it fell below. When the price was equal to the MAV, we would keep the last calculated value unchanged. This was simple enough for a manageable algorithm. We left open the question of the algorithm's main parameter: the MAV algorithm could carry out its task on the basis of a 3-day MAV or a 300-day MAV, and we could determine which one worked best depending on the market and the time-frame.

Exhibit 3: trend confidence calculation based on the current price and a simple moving average (MAV)

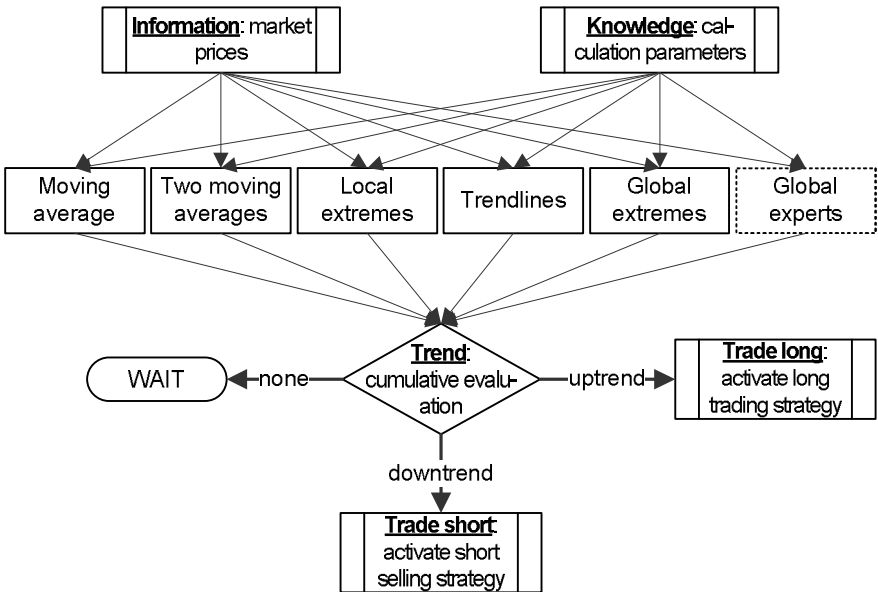


CME Yen futures (daily) with the 16-day moving average and the trend confidence curve based on the relationship between the "current" price (last closing price in the above example) and the 44-day MAV. The trend confidence value is plotted against the scale on the right, ranging between -1 and 1.

BUILDING THE I-SYSTEM

With this we completed the job description of the first expert in our system. As exhibit 3 shows, the result of this algorithm’s work does not appear particularly intelligent, but we expected to attain more intelligent results by defining the job descriptions of the rest of the experts in our model. We did this by following the logic similar to that used with the single MAV value. Once we completed all the trend confidence algorithms, our next care was to allow for their weighing on the premise that some might add more value than others, and should be weighted accordingly. Schematically, the resulting expert system resembled a neural network consisting of six parallel processes, each of which produces an evaluation of the market trend in the present.


Exhibit 4: I-System trend determination neural network



I-System’s trend determination network is an expert system, consisting of six parallel processes, each of which produces an evaluation of market trends in the present. Of the six, five are based on the historical prices and one allows for the input from human experts to influence the system’s judgment. An executive process (cumulative evaluation) consults each of the experts, weighs their “opinions” and arrives at a judgment about the prevailing trend.

An “executive” process consults each of the experts, weighs their opinions and arrives at the system’s final judgment about the trend. When fully operational, this model generated a more intelligent looking trend confidence function. Here’s an example.

Exhibit 5: I-System's intelligent trend confidence function



CME Japanese yen futures (daily) with the trend confidence function based on five trend evaluation processes defined in the I-System's trend judgment neural network. The trend confidence value is plotted against the scale on the right, ranging between -1 and 1.

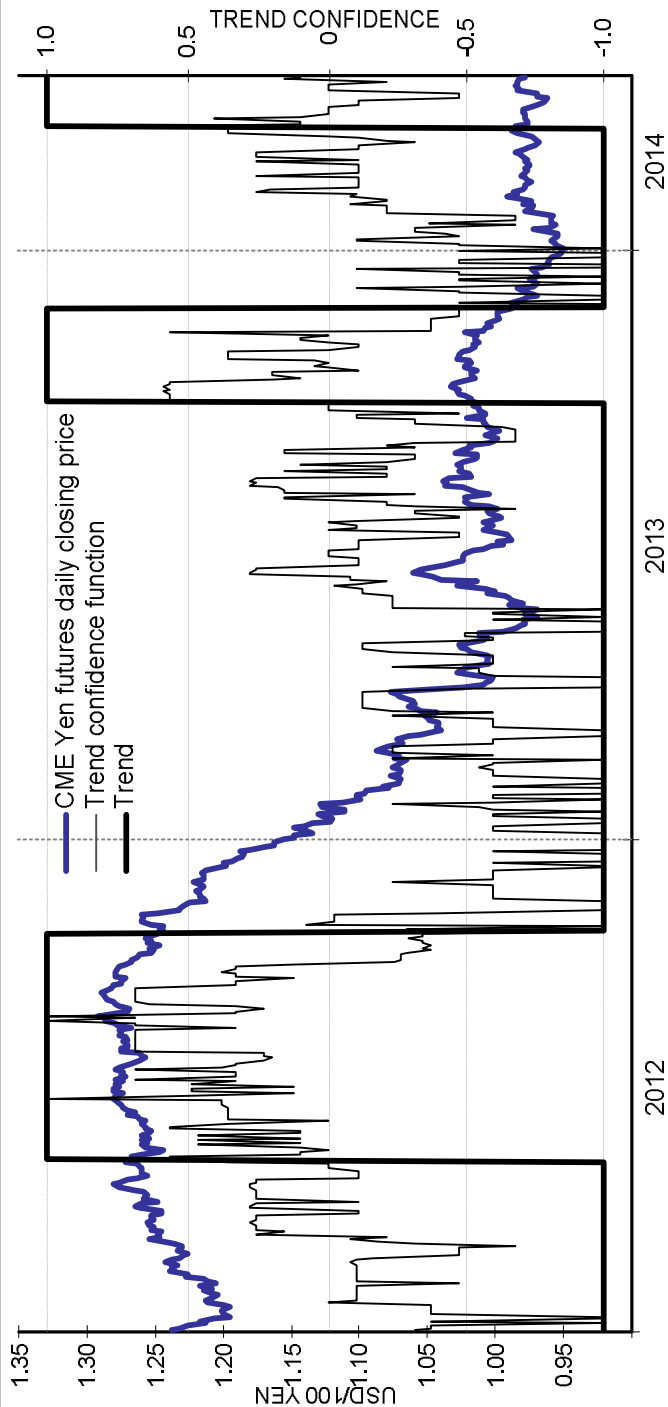
We were satisfied that the trend confidence function roughly resembled what human trend judgment might look like, but with one important advantage: humans can't quantify their judgment accurately, let alone be sure how to act on it in any consistent way. Generating even a rough approximation of human judgment through an artificial neural network had one important advantage: it gave us a numerically exact measure of its confidence at any time and enabled us to calculate with precision the level of confidence that justified taking risks on trades.

Would the trading results be best when we were 100% certain in our judgment? Or should we trade on lesser confidence? As we discovered, we obtained the best trading results with confidence thresholds below 0.5 in uptrends, and above -0.5 in downtrends, which to us were surprisingly low figures. For a human trader, taking risks with low confidence in his judgment would psychologically be very difficult to do. At the same time, it is clear that catching a trend early precludes waiting to be certain about it. With I-System's trend confidence function, we were now able to demonstrate this mathematically, as setting confidence thresholds close to 1 invariably gave poor trading results.

Exhibit 6 illustrates I-System's trend judgment for US Dollar/Japanese Yen futures. Using threshold levels of 0.5 and -0.5, trend judgment switches from calling an uptrend (1) to downtrend (-1) without pausing at zero. With threshold levels defined differently we could also have an indeterminate judgment (0) which would imply that we should refrain from trading altogether².

² As we discovered through millions of backtest simulations, trading strategies tend to perform best over time when they are always in the market, either on the long or on the short side.


Exhibit 6: I-System's intelligent trend confidence with the derived trend definition function



CME Japanese yen futures (daily) with the trend confidence function and the final trend judgment. In the above example, confidence thresholds of 0.5 and -0.5 are used: when trend confidence rises above 0.5, the price is deemed to trend up; when it drops below -0.5, it's trending down. In this case, trend judgment favors either uptrend or downtrend. Threshold levels can also be set in a way that allows for periods where the trend is indeterminate (and trend judgment stays at zero). The trend confidence value is plotted against the scale on the right, ranging between -1 and 1.


I-System’s capability to accommodate different parameter values in calculating trend confidence enabled us to define price trends in more than one way for any market. Typically, we can define a price trend as a shorter-term event, a medium-term event, or a longer-term event, as exhibits 7a, and 7b illustrate using the example of the Gold price chart.

Exhibit 7a: long-term trend definition



Very long-term trend strategies switch sides only seldom and are capable of taking advantage of major sustained trends and “ride out” significant corrections and extended periods of price consolidation. These strategies tend to perform very well over the long term, but using them entails enduring significant draw-downs and months or years of flat performance (as in exhibit 7a above, during 2006/2007, 2008, and 2011-2013 periods).

Exhibit 7B: short-term trend definition

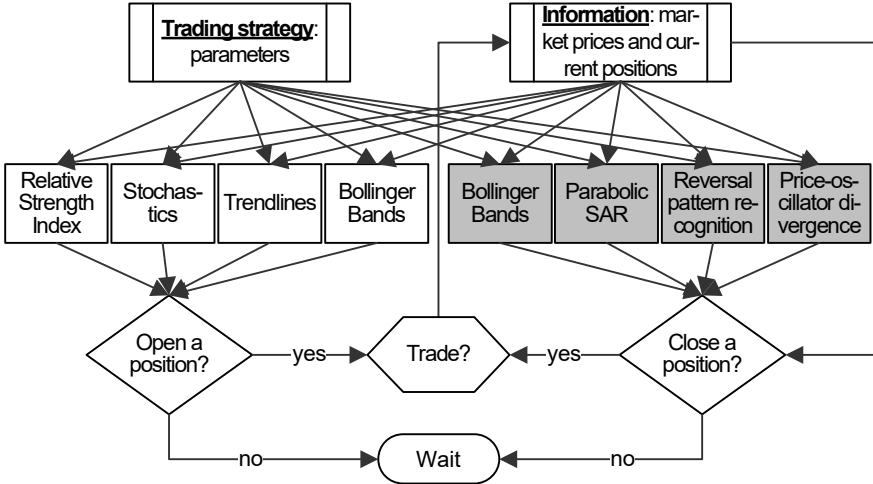


Medium and shorter-term trend strategies will be quicker to switch sides amid trend reversals, generating a less volatile performance curve. However, such strategies can get badly whipsawed during periods of sideways fluctuations as the strategy in exhibit 7b endured during the 2013-2015 period. In whichever way we ultimately choose to define a strategy's trend-judgment function, the key criterion is always the strategy's ability to generate trading profits. In this sense, the model proved effective – at least in backtest simulations.

Determining trade entry and exit signals

Once we were able to determine price trends, we thought we could improve trading results by adding a set of algorithms that would define additional entry and exit signals for trades, including stop-loss trades when there were significant corrections in a trend as well as profit-taking trades where trends had made a significant advance. Here again we borrowed a number of studies from technical analysis including Stochastics, Relative Strength Index (RSI), Parabolic Stop-and-Reverse (SAR), Bollinger Bands, trendlines as well as algorithms that performed reversal pattern recognition (double tops and bottoms and head-and-shoulder patterns). Thus, in an up-trend, entry signals would produce *buy* decisions and exit signals would produce *sell* decisions. Exit signals could be either stop-loss or profit-taking trades. In a down-trend, entry signals generated sell decisions and exit signals the decisions to buy. Exhibit 8 illustrates this part of the I-System model:

Exhibit 8: I-System trading signals generating network



Depending on the system’s judgment about the price trend, a trading strategy is activated. A series of modules defines signals to enter and exit trading positions.

Most of the algorithms defining entry and exit signals were relatively simple to formulate as there was no need to interpret any algorithm’s results. Simply, detecting an opportune entry or exit point in price fluctuations would trigger a trading signal. For example, where one of the oscillator studies (Stochastics or RSI) showed that the market was *oversold* in an uptrend, this gave us a *buy* signal. Similarly, if the price pulled back to a trendline, this also gave us the signal to buy.

In downtrends, the same algorithms would generate *sell* decisions if the market was *overbought* or the price rose to a downward sloping trendline. The only process that involved some complexity were the pattern recognition modules which consisted of several algorithms whose job was to identify double top and double bottom formations and head-and-shoulder patterns in both topping and bottoming markets.

Strategy risk profile

The two networks of algorithms described above solved the problems of when to trade and in what direction to take risks. But we still needed to work out *how much* risk a strategy should take with each trade. Whether in real trading or in backtest simulations, each trading strategy had to start with a risk budget: an amount of cash it could use to buy and sell the securities in question. Trading in futures markets simplified our problem

somewhat: to buy or short sell a futures contract, we only need to post the initial margin rather than the full monetary value of the underlying asset.

For example, to trade one contract of Brent Crude Oil futures, the initial margin requirement is about \$5,000³ so our risk budget needed to be a multiple of this figure, large enough to accommodate any losing streak a strategy was likely to experience. In the case of Brent futures, an adequate risk budget would be about \$25,000 per contract. Another element we needed to take into account were the broker's commissions payable with each trade. This information was easily available from our futures brokers.


These three elements – the strategy's risk budget, the initial margin requirement, and the broker's commission were sufficient to define a strategy's basic risk profile. However, once a trading strategy starts doing its thing, more questions arise. If a strategy is successful, it will accumulate trading gains well in excess of its initial risk budget. If so, do we keep increasing the position size in proportion to the available cash? This approach may seem logical but for one problem: if we keep increasing our position size as our account balance grows, we'll also take bigger losses when the market goes against our positions. Although there are many ways to approach this problem, we decided to resolve the dilemma simply by adopting risk management lessons from natural life.

In nature, every life form grows to a certain size beyond which it can only grow by making further copies of itself. When hunting, predators always take similar risks and target similar-size prey, regardless of whether their recent hunts have been very successful or not. We intended that each trading strategy formulated through the I-System would be treated similarly as an individual predator in nature: it would be given an initial risk budget, it would always make the same size trade, and it would grow by funding new trading strategies similar, but not identical to itself. Any individual trading strategy only needed to generate positive trading gains over time. To our great encouragement, our backtests indicated that they were indeed successful in this regard.

As the example in exhibit 9 shows, a trend following strategy can be a remarkably successful way to speculate provided that price trends materialize. Through 1997/98, the strategy reverses direction three times, failing to profit from a major downward price move. Consequently, the result in 1997 is slightly negative, and the strategy makes the bulk of its gains between 1998 and 2000 as the oil price vaulted from \$10 a barrel to nearly \$35. On the whole, this was all very encouraging to us.

³ This amount is set by the futures clearing house and can vary over time. Initial margin requirements tend to correspond to between 2% and 10% of the monetary value of the underlying assets bought or sold.

Exhibit 9: Performance of a typical trend-following strategy



A trend following strategy on Brent crude oil futures. The strategy starts on the 1st January 1997 with an initial risk budget of \$25,000. By mid-2003, it has generated more than 100% in gains.

But the picture also reveals another point – a critical one for a would-be trend follower to ponder: trend following strategies work when markets trend. When they don't, or when major trends reverse, they are almost certain to generate negative, or at best, flat returns.

If a trader implemented the above strategy in mid-2000, he would have to wait nearly two and a half years before making any trading gains, and it takes a great deal of discipline to keep doing something that fails to meet your expectations for that long. More seasoned practitioners understand that trend following is a sequence of feast or famine, and that in addition to a valid model, long-term success requires much perseverance and the utmost discipline.

To recapitulate, the construction of customized, valid price charts, trend recognition, trade entry and exit signals and the strategy risk profile completed I-System's solution space. Trading with our model would be determined by a multitude of individual strategies defined within that space. In total, 70 different parameters define the speculative behavior of I-System strategies, enabling us to educate a very large number of intelligent virtual traders specialized in almost any market.

Each set of parameters defines a decision-making strategy that performs the work of a trained market analyst or trader in a rigorous, disciplined and quantitative way, free from unhelpful psychological biases, emotions or distractions. Each strategy's objective is to generate value from favorable market moves while limiting losses from adverse fluctuations. Once implemented, trading strategies perform their function autonomously, reducing the complex job of analyzing markets to simple, actionable decisions. Generally, trading strategies differ from one another along two key attributes:

- **Trend cycle** – whether the trend is defined as a long-term, short-term, or medium-term event.
- **Time in the market** – the proportion of time that a strategy spends in trading positions. A strategy might be in the market most of the time, or might enter and exit trading positions very selectively, passing more time waiting for the right triggers.

The experience of having evaluated several million backtest simulations has shown that in most markets, strategies that (1) use longer-term trends and that (2) spend the most time in trading positions, tend to perform best over the long term. At the same time however, such strategies suffer the heaviest losses when major trends reverse because they will recognize the

trend reversal only when the prices have moved in the opposite direction for some time.

Strategies that switch their trend judgment more quickly perform better when major trends reverse, but fall behind during long periods of price consolidation or range-bound trading, as they interpret larger price corrections as trend reversals and repeatedly take positions on the wrong side of the subsequent price move. We hoped that our model would ultimately turn out to be a holy grail of sorts. By allowing us to rely on a quantitative decision making process, I-System might enable us to replace the uncertainty about future price fluctuations with a more predictable risk class: a large set of virtual decision-makers, each in charge of a small fraction of an investment portfolio's risk budget.

By emulating nature's model of risk management in this way, we hoped that we would be able to master the uncertainty inherent in market speculation. Although we were a long way away from proving the effectiveness of I-System in live trading, we were encouraged by the results of our backtests. There was however, one very significant fly in this whole ointment of ours. Namely, as we gradually put together the I-System, our ideas evolved somewhat and we felt the need to adapt the model here and there. We also discovered errors in our code that needed to be corrected. But as the program grew bigger and more complex, correcting errors and making any changes to it became more and more difficult and there was a tendency to introduce one or more new errors with every correction or change we effected.

By the time I-System's prototype was complete – this was in late August 1999, – the need to overhaul the model in a more robust way was glaringly obvious and of the utmost priority if our endeavor was to have a long term future.

Chapter 11: Building the I-System, again.

Peace of mind isn't at all superficial to technical work. It's the whole thing. That which produces it is good work and that which destroys it is bad work.

Robert M. Pirsig

All human error is impatience, a premature renunciation of method...

Franz Kafka

In January 2014, a gentleman – let's call him Arnold – presented himself at the Monaco offices of my current employer, Altana Wealth. He was soliciting funds to complete an ambitious quantitative investing model. During our meeting, we learned that his team had been working on this model since 1993, that he personally invested over 16 million British Pounds in its development, and that he needed further funds – about 500,000 euros – for his team to complete the software program and make it operational. This man was clearly not stupid, and his 12-person team included two PhDs, four masters-level scientists and several software developers. Nevertheless, after more than 20 years of continuous work and a fortune spent on research and development their model was still not operational. To the uninitiated, this may seem quite incredible, but I was not very surprised at Arnold's problems.

Over the years I've come across several similar cases where the software development process became bogged down in its own complexity and ultimately completely stalled without achieving completion. In fact, a very significant percentage of all software projects ultimately fail to attain their objectives. This is due to the complex, but manageable challenges inherent in systems engineering. I was about to learn this lesson soon after Gorazd Medić and I completed the prototype version of the I-System.

In the summer of 1999, our software seemed to function beautifully, but it was very fragile and difficult to maintain. Any change to it carried the risk of introducing new errors and instead of implementing it to start trading, I felt compelled to ask my boss for further funds in order to hire professional software programmers and build a more robust version of the

model. By this time however, our endeavors went quite off the company script, and I had a hard time persuading my boss to continue supporting the project. He specifically wanted us to produce a model that would generate price forecasts so that we could make high probability bets in energy and currency markets – an objective that no longer made good sense to me. Ultimately however, I managed to secure a very small budget to hire a software programmer and finish the job however best I could.

Having studied software programming during my high school days in Croatia, I knew a good many people in the software community there and I contacted a few of them to inquire about whom I should hire. I intended to find the very best programmers in the country and soon I had a list with two names on it. One of them was unavailable, but I was able to meet with the other gentleman: Boris Brec. I explained to Boris what Gorazd and I had been doing and what I would need him to do. Boris found the idea intriguing, but he politely explained that he was very weary of working with dilettantes and told me that he would be very reluctant to take up the project. I had actually been warned in advance that Boris would almost certainly decline to work with me, but I tend not to take no for an answer easily. After our initial meeting I went to see Boris at his office several times over the following days (he was working at the IT department of the Croatian national utility company, HEP).

During that time, I noticed an interesting thing about him: he was very relaxed and appeared to have all the time in the world to chat. As I later understood, this was because his programs required very little maintenance and tinkering so he enjoyed much leisurely time at the office. However, our chats were frequently interrupted by his colleagues who would invariably step into his office stressed and exasperated about being unable to solve some programming problem they were working on. In every case – and I must have witnessed a dozen or so – it took Boris mere minutes to identify the problem and suggest the solution for his colleagues who would then rush off happy and relieved, thanking him and dismayed that they haven't seen the solution themselves. This only made me more determined that Boris was just the man I needed to build an industrial-strength version of the I-System.

After several days of talks and much coffee, Boris said he would consider taking on my project on the condition that I study up on the subject of software engineering under his guidance, which I accepted. He supplied me with study materials – four university textbooks on subjects covering systems analysis, software design, and process diagrams, as well as a number of papers and document templates produced by various software engineering institutes. Fully convinced that I was talking to an

authority, I seized upon this opportunity and returned to Monaco with my stack of study materials.

Computer science is no more about computers than astronomy is about telescopes.

Edsger W. Dijkstra

I can't say that my reading assignment was boring in any sense, but it was very technical and I took almost a full year to work through all the materials. One thing I understood early on was the difference between software programming and software *engineering*. The best way to put it is that a software programmer is to a software engineer as a construction worker is to an architect. You could get away with hiring a construction worker to build a small house, but you couldn't hope to accomplish a complex high-rise building without hiring a competent architect. The same is true for software projects. While most quants can do a decent job of programming, software engineering involves a very different set of skills in which most quants have no training.

Among other things, software engineering focuses on the *process* or *methodology* used in building software systems, in which the actual programming is merely one of the last stages. The quality of the ultimate product is largely determined by the quality of the process applied in a system's development and maintenance. In contrast to our approach with the I-System prototype, which consisted of going from an idea straight to coding, best practices in system engineering require that a project advance through a number of distinct stages in the project life-cycle. In general, these are:

1. user requirements
2. software (and hardware) requirements
3. software architecture
4. software programming instructions
5. production
6. transfer
7. maintenance


At the first stage, the future user of a software system must clearly articulate all the functions that the program should fulfil and how it should fulfil them. The user must document these requirements and produce the

“user requirements document,” which sets the foundation for the subsequent phases of the process. Producing this document forces the user to think in a clear and structured way about the processes and functionalities that the software package must fulfil and to articulate them in a comprehensible way. It also forces the user to make countless decisions that must be made to remove any ambiguities a software developer is likely to encounter. Building any system involves many decisions, and most of these must be made by the user and not the software developer. Defining the user requirements also imposes a scope on the development project so that new ideas which tend to emerge during the software’s development don’t end up sidetracking the project and dissipating time and resources on work that wasn’t part of the original plan. The documentation of user requirements consists of process flow diagrams and text defining the software’s functions, explaining the procedures, and specifying the data involved.

Once completed, the user requirements document forms the basis on which software and hardware requirements are formulated, then the software architecture, and so forth, so that each stage’s outputs are the inputs for the next stage. In my case, the first stage was learning about the process and methods of systems development and about my own role in it as the user. After I had finished my reading assignments it was clear that my next task was drafting the user requirements document. This stage involved overcoming a good deal of reluctance on my part: I had already built the model which worked and I was eager to trade and start generating some concrete results. Going back to the drawing board and spelling out the whole system on paper felt like homework from hell.

It was clearly going to take a great deal of time and effort on my part. Unfortunately, I also knew that if my project was going to have a long-term future, this work was absolutely essential and that nobody else could do it in my stead. Boris helped me by drawing my first, top-level or context diagram:

System context diagram



Top-level, context diagram – my job was to decompose the system part down to the last individual process making part of the system. Data provider, refers to the source of historical and real-time price data which the system requires to perform its functions.

He further explained that I would need to break that diagram down to its most basic elements in such a way that I would have two, maximum three arrows pointing to each process and one arrow pointing out to the next one. So I got busy, bought myself a nice thick notebook and started charting out the process in pencil and drafting my requirements.

At first I found the process incredibly frustrating and difficult, which tends to happen when you have to structure and articulate your own mind's tacit knowledge so that it could be intelligible to others. It took me a full year to complete the user requirements document, which comprised 66 pages of process diagrams, descriptions, formulae, tables and a data dictionary. When I was done, I turned it over to Boris to study. Because of the model's complexity, Boris felt that he needed to get more closely acquainted with all of its processes and algorithms, so in the summer of 2002 he, Gorazd and I got together in Monaco for a few weeks to work through the entire system.

During this time, we ended up re-building the program in Visual Basic language so that we could test it and make sure that everything was working correctly. On this occasion, we also needed to decide which development platform to select in building the system. After some debate and research, we chose Microsoft's SQL data-base and C# programming language. With this, Boris returned to Croatia, defined the system's architecture and produced the "Software Architecture" document. With that, he was able to complete the project with only periodic feedback from Gorazd and myself. The much anticipated product was finally complete in September of 2003, fully four years after we built the prototype. By this time, we had built three versions of the model: the prototype, the 2002 Visual Basic version, and the 2003 C# version. This enabled us to test the system thoroughly and weed out any errors by running the three versions in parallel and comparing their outputs.

The testing phase of software development inevitably calls attention to many things that can be improved in the product, and this was the case with I-System. After about two years of testing, we upgraded the system which resulted in a new version in 2006. One of the key changes was giving the system flexibility in building historical time series by rolling futures contracts over to correspond to the way we would actually trade them¹. Boris was also able to vastly increase the model's speed of computation: the 2003 version of the model took almost 4 seconds to complete a single backtest simulation over a 20-year history of daily price quotes. With the 2006 version, he had reduced that number to about 90 milliseconds, making the new version more than 40 times faster than the

¹ See previous chapter for a more detailed explanation of this.

previous one. This was an important feature for a model capable of expressing a nearly infinite variety of trading strategies. When you need to run tens of thousands of backtest simulations, the difference between 4 seconds and 90 milliseconds is very significant. Finally, the new version also included quality enhancements in the model's user interface which made the process of formulating and testing trading strategies much easier and more intuitive.

As Gorazd and I gradually put together the first version of the I-System in 1999, we became keenly aware of the enormous challenge of maintaining this complex piece of software. At the time, we agreed that we would need to bring the model to a state where we would not hesitate to put our grandparents' savings at risk in trading. At the time, we didn't expect that this would take fully seven years of concerted effort. Nonetheless, I firmly believe that this investment of time and resources was essential. The alternative – rushing from idea to programming, failing to document the requirements and define the project's scope and development plan – was perfectly exemplified by the unfortunate case of Arnold, mentioned at the beginning of this chapter, who spent over 20 years and 16 million pounds without managing to complete his model.

By adhering to best practices in software engineering we obtained a superbly stable and robust product that has functioned continuously since 2006 (nearly ten years at the time of this writing) with virtually no glitches or further maintenance requirements. Also, by building three different versions of the software and testing them in parallel for an extended period of time we were able to remove all doubt that our model correctly fulfilled its intended purpose. The relevance of this achievement could hardly be overstated as I would later learn through the experience of using the I-System. As Baruch Spinoza put it, "*all things excellent are as difficult as they are rare.*" Thanks largely to the master software engineer, Boris Brec, we were able to build an excellent software product.

Now, all that was left to do was to put the system to test in the real world. This, for a host of reasons, proved to be much more difficult than I anticipated, as we'll see in the next chapter.

The curious case of Lockheed Martin's F-35 fighter jet

In the 1990s, the U.S. defense establishment under the administration of President Bill Clinton conjured up a plan to build a new generation Joint Strike Fighter, the F-35. The objective, which justified building the most expensive weapons system in history, was for the US military to achieve superiority in aerial combat and to control the skies in any military confrontation.

By 2014, the project was seven years behind schedule, more than \$160 billion over budget, and appeared hopelessly off track with a distinct chance that it would ultimately have to be dramatically simplified or scrapped altogether. Operation of the aircraft depended on a massive software package – a morass of 24 million lines of code – which was bogged down in development and deemed unreliable by the Pentagon. On any given day, more than half of the F-35s were liable to be down for repairs or maintenance. The aircraft also turned out to have such surprisingly basic defaults as not being unable to fly at night.

How did such powerful organizations with almost unlimited budgets and many decades of accumulated know-how and experience in building military aircraft manage to produce such a disaster? It appears that the fiasco started with a poorly drafted user requirements document. According to Winslow Wheeler², director of the Straus Military Reform Project³, the F-35 project was set up to fail at a great cost.

From the outset, the F-35's design was based on contradictory attributes: the Pentagon wanted a short take-off and vertical landing (STOVL) aircraft to also be supersonic (these two design characteristics apparently don't go together). Further, they wanted a multi-role aircraft, piling on additional contradictory characteristics of an air to air fighter and an air to ground bomber. Someone's wish list included the stealth quality for the F-35, making the aerodynamic design awkward and the whole system more complex by an order of magnitude. Finally, each military service – the Air Force, the Marine Corps, and the Navy – added their own wish lists for the design resulting in a hugely complex, high cost project doomed to deliver poor performance (performing many roles, each of them poorly).

In this case, it is quite apparent that the F-35 project was ruined because it was controlled by an unwieldy bureaucracy and individuals

² “Here's What 60 Minutes Didn't Tell You About the F-35,” PogoBlog, 19 February 2014.

³ Straus Military Reform Project is a part of the military think tank Project at the Center for Defense Information at the Project on Government Oversight

BUILDING THE I-SYSTEM, AGAIN

long on vision but short on specific knowledge of flying combat aircraft. In Wheeler's words, "*Technologists who consider combat lessons an afterthought control the beginning design, and advocates in the industry, Congress, and the Pentagon seek to commit the entire government to the program by spending billions and billions before any empirical data becomes available from testing to show what the actual cost and performance are.*"

The lessons of this case are multiple. For one thing, throwing large amounts of money at a problem can't fix the consequences of an ill-defined and flawed process. To achieve a successful project, it must be defined and controlled by the most knowledgeable and experienced individuals available. These individuals must then be granted adequate resources and conditions to work unhindered by politics and other irrelevant considerations. Unless these conditions are created for the developers to work under, the project is unlikely to achieve a full measure of success. It might also fail altogether.

Chapter 12: Learning to fly

Q: But how do you know that the model you have created is right?

A: There is no proof that Einstein's theory is right. There is no proof that Ohm's law in electricity or Boyle's law in gasses are right. There is only an experimental demonstration that such laws are useful for specific, limited purposes. There is no way of proving that a model or law or theory representing the real world is right. ... The heart of the matter is your relative degree of confidence in each of these models.

Jay W. Forrester¹

From this point on, the personal aspect of this story gets more difficult to omit because this project became my life and vice versa. As I was feverishly working with Boris Brec and Gorazd Medić on the new version of the I-System, I did my utmost to persuade my boss to continue to support our project. My idea was to put it to a trading “forward test” with no money, then to start trading with real money, gradually adding capital and ultimately to establish a proper hedge fund that could be offered to outside investors. Although I knew close to nothing about creating and running a hedge fund, I put together a decent-looking business plan, which to my mind was rather compelling. The last thing I expected from writing that plan was that it would ultimately get me sacked, but after some discussions, that's exactly what happened: my boss was not interested in diversifying the business away from oil trading and gave me the choice to either focus on the firm's core activities or to take my project elsewhere.

By this time, I had lost my passion for Greenoil's business but had full faith in the future of my project with the I-System. The choice of a “safe” job on a sinking ship was not really a choice, so I happily jumped off on the life-boat of my savings and a small severance package I got from Greenoil. I reckoned, I would figure it out from there.

¹ In an interview with McKinsey Quarterly, 1992.

However, once I started reaching into the world of hedge funds, I realized I was next to invisible there. My professional background at a small oil trading outfit turned out to be quite a liability and my initial attempts to find a new home for the project attracted very little interest. Hedge fund firms tended to be spun off from major financial institutions like Goldman Sachs, Deutsche Bank or UBS, or from other well-established hedge funds where experienced managers with demonstrable performance track records set out to create their own firms, usually with significant support from their employers and investors. After leaving Greenoil, all I had was a model and a story I thought would get investors interested. But I found only a handful of opportunities to even present the story and without a track record my arguments had little credibility.

It was clear that I needed to generate a track record one way or another. In February 2004, I set up a \$2 million account with Refco Simulated Trading Services. I allocated the funds to 14 different trading strategies in as many different futures markets and started to trade through Refco's online trading platform. The money I used was virtual, but I executed trades against real bid-ask prices in the futures markets. The results were encouraging and by the end of 2004 my portfolio generated a gross return of 31.06%. During the same period, various indices of commodity funds, or CTAs² reported returns ranging from 1% to 6%³. At the time I was tempted to attribute this outperformance to the sheer awesomeness of the I-System, but the truth was rather more nuanced.

A significant part of my outperformance was down to dumb luck. For one thing, my portfolio was leveraged about ten times while the average CTA tended to be leveraged at most 3 to 3.5 times. Furthermore, by trading in only 14 futures markets, my portfolio was not terribly well diversified. Large CTAs routinely trade in as many as 50 to 100 different markets. For this reason, their performance tends to be less volatile. Less diversified portfolios tend to be more volatile and their performance depends on what happens in the markets where they trade. For example, if you only traded one market and that market experienced strong price trends, you could achieve very high returns during that period.


A good case in point was the BlueGold Capital Management which focused solely on trading energy derivatives. In 2008, BlueGold generated a net return to investors of 209% and another 55% in 2009. This was during a time when oil prices staged exceptionally favorable trends.

² CTAs are commodity trading advisors

³ Edhec CTA Global index returned 4.64% for 2004; Credit Suisse First Boston (CSFB) Invx Managed Futures index 1.04%, CSFB Hedg Managed Futures 5.97%, and Morgan Stanley Capital International (MSCI) Directional Trading index 3.00%.

However, if you traded these same markets from around April 2011 through mid-2014, as a trend follower you were stuck in an extended losing streak. Exhibit 1 provides an illustration based on a typical I-System trend following strategy:

Exhibit 1: Markets move in trends – but not always...



This chart shows the performance of a trend following strategy trading Brent crude oil futures. The strategy begins with an initial risk budget of \$50,000 and trades a single 1,000-barrel contract on the IPE (International Petroleum Exchange). From 2007 through mid-2008, oil rallied to about \$146/bbl. Over the following six months, it collapsed to \$40/bbl. In 2009, the trend reversed again and the price doubled from \$40 to \$80/bbl. These were very favorable winds for trend sailors. Panels A, B, C, and D show four distinct phases in this period.

A: strongly trending price offered very high returns on trend following. The above strategy generated a return of over 300% from January 2007 through December of 2009, with only one major drawdown in the wake of the trend's reversal in mid-2008. This draw-down may appear small, but it amounted to about 60% of the strategy's risk budget.

B: price consolidates around the \$80/bbl level. When prices fluctuate in a sideways range, trend-following strategies tend to generate losses.

C: In late 2010, the price broke out of its range and advanced another 50% from the \$80/bbl level, enabling further trading gains.

D: The period from April 2011 through June 2014 represents an exceptionally unfavorable environment for trend following. If you used a trend following strategy in crude oil markets at that time, you would have experienced more than three years of negative performance. Although these losses appear minor compared to the preceding returns, in the case of the above strategy they amount to almost \$80 per barrel of Brent crude oil, or \$80,000 per contract – more than 150% of the strategy's risk budget. Accordingly, an undiversified trend follower – even with the world's finest trend following model – was liable to sustain a total loss.

In fact, BlueGold shut down about one year after the oil price peaked in April of 2011, after losing about 34% in 2011 and continuing to lose more money in 2012. Thus, an undiversified portfolio can generate very strong results as well as very steep losses, depending on market conditions. The more diversified a portfolio is, the more closely its performance will reflect the relative quality of the manager's trading strategies and risk management skills. It will also tend to have more limited performance as losses in markets that fail to trend offset the gains from trending markets. Because I only selected 14 markets for my Refco trading account, my outperformance was fortuitous to a significant degree.

As I was beginning to work these things out, I realized that I still had a lot to learn about the business of managing investment portfolios, as opposed to building models or just trading. I needed to upgrade my skills at formulating trading strategies, constructing diversified portfolios and managing risk. I also began to appreciate that for all the blood, sweat and tears we had put into the building of our model, I-System was only a tool. It was a fine and superbly engineered tool – but still just a tool. It was a solution to the problem of speculation as much as an airplane is a solution to the problem of air travel. An airplane is a magnificent feat of engineering, but flying one in a safe and comfortable way still requires much skill and experience.

In 2004, I was only just beginning to acquire the skills and experience of managing an investment portfolio. Among other things, that experience would teach me that in addition to a good model, asset management requires steadfast discipline and strong conviction. Occasionally, as the panels B and D in Exhibit 1 show, it also requires great perseverance and nerves of steel.

Formulating robust strategies

There was another way I was lucky with my 2004 track record. My trading strategies turned out pretty good even though I did not formulate them with all the care and attention I would later learn to apply to the process. Namely, at first I only looked at a single attribute in evaluating trading strategies: their absolute performance over a period of time, without much regard for *how* they achieved that performance. Around that time, I had the good fortune to meet one Jan Haraldson, Monaco-based futures trader who had been managing a diversified portfolio of systematic trend following strategies since the early 1990s. Although not a household name, Jan is probably one of the world's very best trend followers and I had the privilege of periodically meeting with him for lunch and discussing our respective models and strategies. Through these

conversations, he gave me the mentorship which I very much needed at the time. When it comes to the technical minutiae involved with systematic trading, there are certain things you can only learn by experience. The opportunity to learn such lessons from the rich experience of one of the best trend followers anywhere was invaluable.

Jan suggested that I should be much more discerning about my strategies and rather than just looking at the end result, to analyze carefully how and why a given strategy performed as it did. For example, for a given kind of trading signal to be considered valid, it had to be systemic to the market in question. If some study, say a moving average, stochastics, or Bollinger Bands, only generated a handful of very profitable trades over a ten or twenty year time period, such a signal might not recur for a very long time in the future. Or it might never occur again, meaning that the strategy that depended on such a signal couldn't replicate its past performance.

To formulate more reliable strategies, every trading signal used had to show some regularity of occurrence. With strategies based on daily price history, this would mean at least 15 or 20 occurrences over a ten year period. Also, a robust trading strategy should generate results in as even a manner as possible. In every market, price fluctuation dynamics change somewhat over time, and a good strategy should sustain reasonable performance throughout. A strategy that generates the bulk of its gains during a few short intervals while running flat or negative over most of the price history is unlikely to perform well in the future. At the time, such insights had not yet matured in my mind.

For instance, pondering the changing nature of markets, at first I thought that it would perhaps be wiser to use shorter lookback periods to formulate strategies so that they would be better adapted to the current market environment. But upon reflection, this wasn't such a good idea. Formulated over a shorter time interval, a strategy was more likely to be ill-adapted once market dynamics changed again – a change we would only recognize after the fact, perhaps after suffering a long series of losses. Then you might formulate a new strategy, fit over the most recent price history, but this one might also fail if that environment changed, and soon you're back to wasting energy (and money) like a dog chasing his tail. This leads us to the key reason why it is critical for trading strategies to be as robust as possible – meaning, as suitable to *all* market environments.

It all hinges on your confidence

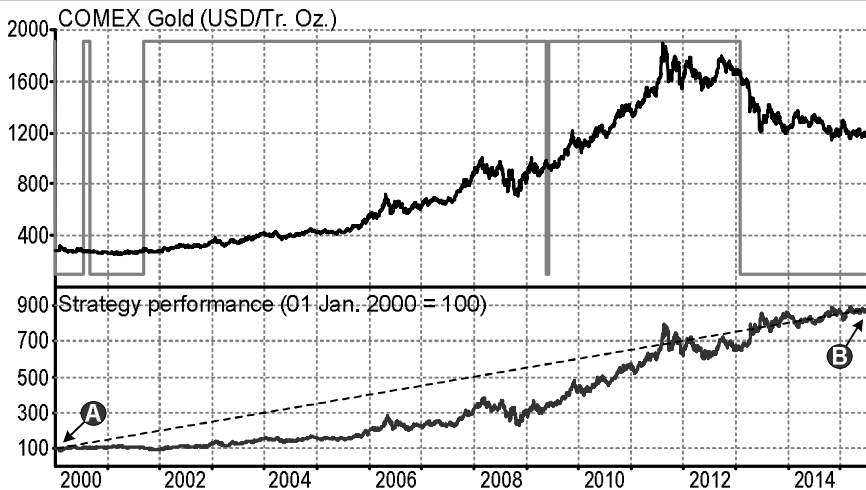
Importantly, this reason has more to do with psychology than it does with performance itself: a robust strategy gives us a high degree of confidence

in using it. This confidence will determine how we cope with adverse market conditions and with the *experience* of sustaining losses over an extended period when – not if – they happen.

Expectancy and the psychology of trading

Expectancy is the answer to the question, “*what happens if I continue doing this?*” This deceptively simple question is of central importance in speculation. With regard to systematic trading, “doing this” implies using a trading strategy. We always formulate strategies to have a positive expectancy, and establish this through backtesting. Backtesting measures how a given set of rules would have performed in the past. Strategies that appear to generate the highest trading gains with the lowest volatility of returns (the smoothest growth of profits over time) are the ones we prefer to implement in the face of an unknowable future.

Exhibit 2: A good strategy has positive expectancy. A long-term trend following strategy, illustrated below generated strong trading gains over time.



This strategy seems superb between points A and B. However, if we implemented it in the late 2011 we would have to endure nearly two years of negative performance.

The limitation of backtesting is that it compresses time into a snapshot of history. Examining a strategy’s performance over a long period of time (as between points A and B in Exhibit 2) can’t convey the day-to-day *experience* of making and losing money. This experience has powerful psychological effects that can influence the results of an investment management process. For example, if you implemented the strategy shown

in exhibit 2 in the second half of 2011, your initial experience with it would have been almost a full year of negative performance.

Watching your losses mount for that long can make the urge to change something almost irresistible. You might consider abandoning the strategy or replacing it with one that looked better at that time. However, in succumbing to this urge you could abandon a strategy that's near the end of its losing streak, depriving yourself of the gains that would follow. Even worse, in replacing the strategy, you might start using one whose losing streak was only about to begin.

Psychologically, we are all hardwired with the loss aversion bias. We are also hardwired to expect that the future will resemble the recent past. A few months of losses could convince us that "doing this" leads straight to ruin. To stick with a trading strategy through a losing streak, the trader must have full confidence in it, as well as a high degree of conviction in the correctness of his model. Otherwise, he's liable to alter course, tinker with the strategy or replace it. Worse yet, he might abandon his risk management discipline and start gambling with his losses.

The corollary of this lesson is that even with well-formulated, positive expectancy strategies, achieving high investment returns over time requires being able to tolerate extended losing streaks without losing composure and altering course. This may well be the hardest and the most important lesson to master in investment management.

Designing a diversified investment portfolio

Another skill I needed to develop was the construction of risk-balanced, diversified portfolios. Systematic trend followers aren't too particular about the markets in which they trade, so they tend to seek the greatest and most balanced achievable diversification for their portfolios. Futures markets offer trend followers around 100 viable markets in six groups: energy, metals, agricultural commodities, equities, treasuries (interest rates), and currencies.

All these markets differ in terms of price volatility. Consider for example, the contrast between one of the most volatile markets (coffee) and one of the least volatile ones (2-year U.S. Treasury Note). Over the ten year period from 2004 to 2014, the average and largest daily price changes in coffee futures were 1.49% and 13.85%, respectively. Over the same period, the average and largest daily price changes for 2-year Notes were 0.06% and 1.05%. Thus, to make similar-sized bets in both markets, we would want to make relatively small bets in coffee futures and much larger bets in 2-year Note futures.

To work out a balanced risk exposure across many markets and determine the position limits for a given portfolio we need a meaningful way of measuring risk. One of the useful methods to achieve this is the so called value at risk model, or VaR. VaR uses statistical analysis of historical price fluctuations to estimate the extent of likely losses from exposure in some market. There are several ways of calculating the VaR, but the most common one looks at the statistical distribution of 1, 3, or 5-day price changes over some lookback period.

Assuming normal frequency distribution, it calculates the potential losses at a 95% or 99% interval of statistical confidence. In plain English, for a given exposure size, the 5-day, 99% confidence VaR quantifies the risk of loss associated with the largest 1% of 5-day price moves, which is a useful way to quantify the volatility of market price fluctuations.

One of the weaknesses of VaR is that it doesn't tell you what happens beyond the 1, 3, or 5-day periods. Namely, the price in a certain market might move strongly against your position, but if you keep it unchanged and the price continues moving against it, your losses could end up much larger than your VaR estimate. For this reason, I also looked at the catastrophic loss scenario – the size of losses resulting from the single largest 5-day price change in each market's price history, as well as each trading strategy's draw-down history.

I had the chance to put all these elements into practice in 2007. From the time I left Greenoil up until this time I had been paying my bills by using the I-System to provide trading decision support for the Monaco-based Galaxy Energy Group⁴. As that experience went rather well, some time before 2007, Galaxy's shareholders proposed to finance my project of establishing a hedge fund.

The idea was that they would seed my fund with a substantial amount of capital provided that I first generated a satisfactory track record with a smaller, \$1 million portfolio. This time we were talking real, not virtual money. The first step in my portfolio design model was to select as many markets as possible given the size of the portfolio.

One million dollars is insufficient to fund a well diversified futures portfolio, but I was able to include 21 futures markets covering all six market groups. I then divided the portfolio among these groups and markets and assigned a risk budget for each market, aiming to achieve a roughly balanced risk profile, as illustrated in exhibit 3.

⁴ Subsequently named Berkshire Management., then G.E.G. Group

MASTERING UNCERTAINTY IN COMMODITIES TRADING

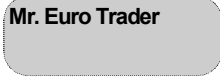


Exhibit 3: The first step in portfolio construction is to allocate a risk budget to each selected market

Initial investment	Market groups	Markets																											
<div style="border: 1px solid black; padding: 10px; width: 100%; height: 100%;"> Total initial equity (\$ 1,000,000) </div>	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="text-align: center; padding: 5px;">Currencies (\$150,000)</td> </tr> <tr> <td style="text-align: center; padding: 5px;">Energy (\$115,000)</td> </tr> <tr> <td style="text-align: center; padding: 5px;">Metals (\$175,000)</td> </tr> <tr> <td style="text-align: center; padding: 5px;">Agricultural commodities (\$245,000)</td> </tr> <tr> <td style="text-align: center; padding: 5px;">Equity indices (\$185,000)</td> </tr> <tr> <td style="text-align: center; padding: 5px;">Treasuries (130,000)</td> </tr> </table>	Currencies (\$150,000)	Energy (\$115,000)	Metals (\$175,000)	Agricultural commodities (\$245,000)	Equity indices (\$185,000)	Treasuries (130,000)	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="padding: 2px 5px;">Yen (50,000)</td></tr> <tr><td style="padding: 2px 5px;">Euro, (50,000)</td></tr> <tr><td style="padding: 2px 5px;">Brit. Pound (50,000)</td></tr> <tr><td style="padding: 2px 5px;">Brent crude (35,000)</td></tr> <tr><td style="padding: 2px 5px;">Gas oil (35,000)</td></tr> <tr><td style="padding: 2px 5px;">Gasoline (45,000)</td></tr> <tr><td style="padding: 2px 5px;">Gold (40,000)</td></tr> <tr><td style="padding: 2px 5px;">Silver (30,000)</td></tr> <tr><td style="padding: 2px 5px;">Copper (50,000)</td></tr> <tr><td style="padding: 2px 5px;">Palladium (55,000)</td></tr> <tr><td style="padding: 2px 5px;">Soybeans (35,000)</td></tr> <tr><td style="padding: 2px 5px;">Coffee (35,000)</td></tr> <tr><td style="padding: 2px 5px;">Cocoa (43,500)</td></tr> <tr><td style="padding: 2px 5px;">Cotton (43,500)</td></tr> <tr><td style="padding: 2px 5px;">Sugar (45,000)</td></tr> <tr><td style="padding: 2px 5px;">Orange juice (43,000)</td></tr> <tr><td style="padding: 2px 5px;">S&P500 (80,000)</td></tr> <tr><td style="padding: 2px 5px;">Nasdaq (55,000)</td></tr> <tr><td style="padding: 2px 5px;">Nikkei 225 (50,000)</td></tr> <tr><td style="padding: 2px 5px;">2-yr. T-Note (70,000)</td></tr> <tr><td style="padding: 2px 5px;">30-yr. T-Bond (60,000)</td></tr> </table>	Yen (50,000)	Euro, (50,000)	Brit. Pound (50,000)	Brent crude (35,000)	Gas oil (35,000)	Gasoline (45,000)	Gold (40,000)	Silver (30,000)	Copper (50,000)	Palladium (55,000)	Soybeans (35,000)	Coffee (35,000)	Cocoa (43,500)	Cotton (43,500)	Sugar (45,000)	Orange juice (43,000)	S&P500 (80,000)	Nasdaq (55,000)	Nikkei 225 (50,000)	2-yr. T-Note (70,000)	30-yr. T-Bond (60,000)
Currencies (\$150,000)																													
Energy (\$115,000)																													
Metals (\$175,000)																													
Agricultural commodities (\$245,000)																													
Equity indices (\$185,000)																													
Treasuries (130,000)																													
Yen (50,000)																													
Euro, (50,000)																													
Brit. Pound (50,000)																													
Brent crude (35,000)																													
Gas oil (35,000)																													
Gasoline (45,000)																													
Gold (40,000)																													
Silver (30,000)																													
Copper (50,000)																													
Palladium (55,000)																													
Soybeans (35,000)																													
Coffee (35,000)																													
Cocoa (43,500)																													
Cotton (43,500)																													
Sugar (45,000)																													
Orange juice (43,000)																													
S&P500 (80,000)																													
Nasdaq (55,000)																													
Nikkei 225 (50,000)																													
2-yr. T-Note (70,000)																													
30-yr. T-Bond (60,000)																													
The initial assets under management (AUM) constitute the main determining attribute of an investment portfolio.	The first step is to allocate them roughly among the market groups we intend to include in our portfolio.	Further, we allocate smaller amounts to individual markets, creating virtual trading accounts for each market.																											

Under this model of portfolio construction, I treated the money allocated for each market as a virtual trading account to fund one or more trading strategies. By coupling these trading accounts with strategies, I conceptually created virtual, autonomous trading agents, each representing an independent unit of speculative behavior, equivalent to a human trader who has a certain amount of money to invest in the market in which he is specialized. Exhibit 4 below illustrates the idea.

LEARNING TO FLY

Exhibit 4: Constructing virtual autonomous trading agents

TRADING STRATEGY	RISK BUDGET	AUTONOMOUS AGENT
Trading strategies consist of rules about how we make <i>buy</i> or <i>sell</i> decisions. One or several strategies are formulated for each selected market.	A strategy needs a risk budget to do its work: a sum of money it can use to trade. Each strategy's risk budget is treated as a virtual trading account at its own disposal.	A strategy and its risk budget make an autonomous trading agent – an independent unit of speculative behavior.
	+ 	= 


With this, I determined the position limits in each market by observing three criteria: value-at-risk, “catastrophic loss scenario,” and the maximum drawdown per strategy:

A	Value-at-Risk (5-day, 99%)	The most severe 1% of 5-day price changes in each traded market corresponds to a gain or loss of no less than 0.2% and no more than 0.8% of the notional amount of assets under management (AUM), for a total value-at-risk of no less than 6% and no more than 12% of total assets of \$1,000,000.
B	Catastrophic loss scenario	The most severe 5-day price change in the twenty year history of each market should result in a gain or loss of no more than 100% of each strategy's risk budget.
C	Largest drawdown	The largest drawdown per strategy, a result of a series of losing trades (over the past twenty years) should not exceed 100% of the strategy's risk budget.

Applying these criteria to the selected 21 markets (and allowing for specific exceptions) resulted in a reasonably balanced risk exposure, as illustrated in Exhibit 5.

MASTERING UNCERTAINTY IN COMMODITIES TRADING

Exhibit 5: Portfolio composition by Value-at-Risk (5-day, 99% confidence)



Market	Risk budget	Value-at-Risk		Catastrophic loss		Largest drawdown		Position limit
	USD	USD	% of AUM	USD	% of risk budget	USD	% of risk budget	
Yen/USD	50,000	3,733	0.37	30,325	61	37,500	75	2
Eur/USD	50,000	2,702	0.27	11,438	23	37,500	75	1
GBP/USD	50,000	3,800	0.38	8,900	18	20,000	40	2
Brent	30,000	3,023	0.30	7,240	24	12,900	43	1
Gas oil	37,500	3,126	0.31	7,275	19	17,625	47	1
Nat. gas	47,500	6,972	0.70	22,306	47	32,300	68	1
Gold	40,000	4,724	0.47	13,580	34	24,800	62	2
Silver	30,000	2,798	0.28	9,610	32	27,000	90	1
Copper	50,000	3,611	0.36	9,225	18	23,000	46	2
Palladium	55,000	4,282	0.43	42,260	77	20,350	37	2
Soybeans	35,000	4,583	0.46	15,000	43	22,400	64	2
Coffee	35,000	5,679	0.57	63,225	181	24,150	69	2
Cocoa	42,500	2,615	0.26	12,450	29	13,175	31	3
Cotton	45,000	2,590	0.26	31,720	70	13,950	31	2
Sugar	45,000	5,748	0.57	16,330	36	18,900	42	6
Org. juice	42,500	4,786	0.48	21,938	52	14,875	35	5
S&P 500	80,000	3,195	0.32	24,225	30	28,800	36	3
Nasdaq	55,000	1,747	0.17	44,780	81	28,050	51	2
Nikkei	50,000	5,003	0.50	20,950	42	23,000	46	2
2yr Note	70,000	4,615	0.46	28,969	41	39,200	56	9
30yr Bond	60,000	6,157	0.62	23,375	39	28,200	47	4
Totals	1,000,000	86,489	8.65	465,120	46.51	486,675	48.67	


The following table summarizes what exactly these numbers mean:

LEARNING TO FLY

A	8.56%	Value at Risk: If the largest 1% of 5-day price moves occurred at the same time, and we were fully exposed on the wrong side in each market, the portfolio would lose 8.65% in such a 5-day period.
B	46.51%	Catastrophic loss: If a price swing equal to the single largest historical 5-day price change exactly coincided in each of the 21 markets traded and our strategies held the “wrong” exposure in every one of them, we could expect to lose about 46.5%.
C	48.67%	Largest drawdown: If each of the 21 strategies experienced a draw-down equal to its most severe losing streak (in its simulated behavior since 1992), and such draw-downs bottomed out on the same day for every strategy, the portfolio would lose about \$490,000 (49%).

In this way I had a precisely defined and reasonably well diversified investment portfolio. The above values, which represent plausible, but improbable events provided us a framework of what to expect in case of an unusually disruptive adverse dislocation in the markets. More importantly however, the benefit of having a numerically defined portfolio enabled us to backtest the whole portfolio and derive a fairly comprehensive set of statistics about its speculative performance. For example, measuring the portfolio’s daily profits and losses gave us a rather realistic idea about the portfolio’s risk profile.

Exhibit 6: Distribution of daily fluctuations in portfolio value




99% of observations fall in the +/- 1.4% (\$14,000 per \$1,000,000 investment).

For a total of 1,429 simulated observations from January 2001 to July 2006, the largest recorded daily loss was just under \$36,000 while the largest daily gain was \$27,500. The average daily value change was

\$1,590 with 99% of observations falling within the +/- \$14,000 interval. Backtesting also enabled us to analyze the portfolio's rolling 12-month returns, which gave us a way to form our expectations for the first 12 months' results once the portfolio went live.

The curve shown in exhibit 7 represents a series of snapshots of 12-month performance relative to the initial investment simulated from 2001 through mid-2006. Targeting a net performance⁵ of around 24%, by this measure we could expect the actual performance during the first year of actual trading to fall between 14% and 40%.

Exhibit 7: Rolling 12-month return on assets (ROA)



Based on performance simulation from 2001 through mid-2006, the diversified portfolio detailed in Exhibit 5 was expected to generate for investors a net return of between 13% and 40% - provided that the futures markets included in the portfolio continued to fluctuate similarly as they have done in the past.

Such statistical measurements serve a very important purpose. They enable us to ascertain whether the actual results of our trading activities conform to the expected values. If they do, this is an important confirmation that our model and our investment management process is functioning as intended. If not, they signal that something needs to be adjusted. In this way, the quantitative approach to asset management provides a uniquely solid foundation for high confidence investment management. After I prepared and presented my business plan and the portfolio to Galaxy's shareholders, we formed a new venture called Galstar Derivatives Trading. By the end of March 2007 an offshore legal

⁵ Because hedge funds usually take out management fees and performance fees, the above simulation included a monthly deduction of 1/12th of 2% in management fees, and a 20% quarterly deduction on the portfolio's gross gains.

LEARNING TO FLY

entity was set up, along with a trading account with Fimat (later renamed as NewEdge) one of the leading futures brokerages, and in April 2007 I started trading. I have to say, in spite of having done my homework quite meticulously, I was nervous about going live. Nonetheless, Galstar took off to a good start and began generating encouraging results, very much within the bounds of my optimistic expectations. Galstar's results are detailed in the following chapter.

Chapter 13: So far so good

Don't be too timid and squeamish about your actions. All life is an experiment.

Ralph Waldo Emerson

I'm not afraid of storms, for I'm learning to sail my ship.

Aeschylus

I have used the I-System continuously since 2004, invariably to my satisfaction. Which is not to say that everything I did with it was a success. As with most – probably all – trend following systems, the results tended to follow a succession of feast and famine periods, depending on the conditions in the markets I traded. My first track record with a diversified investment portfolio was excellent as I discussed in the previous chapter. In part, this was thanks to favorable trends and not excluding some luck. Nevertheless, on the back of that track record, I was able to engage a small group of investors willing to support me and in February of 2005 I launched a hedge fund with \$2.3 million in assets under management.

Biosphere Strategic Capital Fund

The fund, which I named Biosphere Strategic Capital Fund was an offshore, Cayman Islands-based vehicle, with the advisory office in London under the umbrella of an FSA¹ authorized company named PCE Investors Limited². Unfortunately, that fund failed in a few short months. There were several reasons for its failure, exacerbated by three unfortunate decisions on my part. The first one was launching an undercapitalized fund. Even in 2005, when life was infinitely simpler for start-up hedge fund managers than it is today (in 2015), it was not a good idea to launch a

¹ The FSA (Financial Services Authority) at the time was the regulatory agency covering the financial services industry. In the aftermath of the 2008 crisis, the FSA became the FCA (Financial Conduct Authority).

² That's quite a sausage of a sentence, but it reflects the peculiarity of the legal structure customary for European-based hedge fund firms. The hedge fund itself is usually an offshore entity based in the Cayman Islands or a similar jurisdiction while the actual work of research and trading happens at an onshore advisory firm hired by the fund to manage the fund's assets. In addition, the offshore fund normally appoints a third party administrator to do accounting for the fund, calculate its net asset value and coordinate investments and redemptions.

hedge fund with less than \$10 million in assets under management (AUM). Before launch, I was only able to come up with \$2.3 million in hard cash, but during my pre-launch marketing effort I had obtained commitments from several other investors which amounted to between \$16 and \$40 million. These investors couldn't invest on day one, but indicated they would invest once there was at least a few months' worth of live trading track record.

Eager to get going, I decided to launch with what I had, and believed further investments would materialize soon enough. However, during the first three months after the fund's launch, commodity markets were rather unkind to trend followers and the Biosphere fund started with drawdowns. These were made worse by another of my decisions: because my setup and running costs were quite high relative to the fund's AUM (I needed to generate gross returns of almost 7% to just break even), I felt I needed to target high investment returns. This entailed higher leverage and made my drawdowns worse than they needed to be. During the first three months of operation, my fund sustained a 32% drawdown³.

Another unfortunate decision contributed to these losses. Namely, in starting to trade my investment portfolio, I decided to open trading positions only with new trading signals. In other words, if on the first day of trading I was supposed to be long this commodity and short that one based on signals that had occurred some time in the past, I ignored those positions and only traded on new signals. This decision turned out particularly unfortunate because the only market segment that was trending favorably from the get-go were the energy derivatives which advanced some 20% from February through May 2005. But because these were "old" positions – trades based on signals that preceded the fund's launch, I sat out a good bull market in energy futures without benefiting from the trend⁴. All this was enough for my largest investor. After three months of trading he wanted out, and that spelled *game over* for the fund. Once he pulled the rug out from under the venture, I was obliged to reimburse the remaining investors and shut the fund down. Without a doubt, this was one of the most excruciating experiences in my life. In time however, I was able to make peace with it.

Like a race car driver could crash his car on the race track, I crashed my first hedge fund. I still had a well functioning tool in my hands and discouragement was out of the question. At the end of the day, my debacle

³ Part of this drawdown included operating costs and fund liquidation costs

⁴ Later, Jan Haraldson would instruct me that I was wrong to wait for new trading signals: the *curve* I was trading was the cumulative performance curve of my entire portfolio. *When* each trading signal was generated was irrelevant once the portfolio was live. In other words, I should have opened *all* of my trading positions on the fund's first day of trading.

was a good learning experience – nothing sharpens your focus quite so keenly as the pain of a total and humiliating defeat.

Galstar Derivatives Trading

There are some things which cannot be learned quickly, and time, which is all we have, must be paid heavily for their acquiring. They are the very simplest things and because it takes a man's life to know them the little new that each man gets from life is very costly and the only heritage he has to leave.

Ernest Hemingway

I was very fortunate to have an alternative source of income at that time. Namely, in the spring of 2004, I proposed the I-System to a Monaco-based oil trading firm Galaxy Energy Group for decision support in trading and risk management. This contract enabled me to pay my bills for a time, but also ultimately led to my next reincarnation in the investment management universe. After more than two fairly successful years of using I-System strategies in oil trading, Galaxy's shareholders offered to finance my next hedge fund project with a much more substantial seed investment.


The first step was to generate one or two years' worth of track record with a smaller amount of money, audit it, and then move on to set up the proper hedge fund venture. After a few months of preparations, we set up a new offshore entity named Galstar Derivatives Trading (GDT) and in April 2007 I started trading the diversified portfolio described in the last chapter as a managed account⁵.

As I was not keen on repeating my 2005 experience with the Biosphere fund, Galstar's portfolio was much less aggressive. I also opened all of my trading positions on the first day of trading. All in all, Galstar was a success, generating a very decent track record that outperformed many of the leading managed futures funds. Exhibit 1 shows Galstar's results from inception through November 2013: Over the entire period, Galstar generated an annualized compound rate of return of 7.67%, with the worst drawdown reaching -23.92%.

⁵ A managed account is a simple arrangement where an entity simply grants an advisor the power of attorney over an account. The advisor then invests the funds in the account.

SO FAR SO GOOD

Exhibit 1: Galstar Derivatives Trading Ltd. vs. Dow Jones Credit Suisse Blue Chip Managed Futures funds (April 2007 – November 2013) – value added monthly index net of 2% management fees and 20% performance fees.



GDT and DJCS Blue Chip funds diverged strongly in 2013 because the blue chips are all large, multi-billion dollar funds holding the bulk of their exposure in treasuries (interest rate) markets which experienced a very sharp downward correction during April and May of 2013. Galstar benefited by holding a relatively small exposure to these markets, and no exposure to Japanese Government Bonds futures which sustained a very abrupt correction during this period.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Year	Cum.
2007				0.84	0.82	2.43	-1.11	-0.89	4.96	1.34	2.12	1.35	12.33	12.33
2008	5.27	10.96	-5.64	-3.98	1.48	2.56	-4.93	-1.01	2.16	12.94	4.93	1.29	27.12	42.79
2009	1.54	-0.65	-2.36	-0.53	3.09	-8.21	0.75	2.68	4.32	0.29	3.27	-2.50	1.03	44.26
2010	-5.66	-0.20	0.54	0.00	-6.73	0.72	-3.04	0.80	7.77	9.74	0.80	12.67	16.76	68.44
2011	0.05	4.66	0.29	6.85	-7.91	-6.62	5.18	-0.36	-4.09	-4.61	0.14	2.11	-5.40	59.34
2012	-1.21	0.97	-1.74	-1.99	-1.40	-2.27	0.13	1.73	0.21	-3.80	0.59	-1.29	-9.75	43.80
2013	0.96	1.38	2.25	3.21	3.61	0.95	-0.36	-3.63	1.33	0.61	2.93		13.93	63.70

At a risk of sounding like a parent fawning over his child, I find these results remarkable. Not because they are spectacular in the absolute sense, but because they are completely the result of a set of quantitative strategies. To achieve this performance, I strictly did not need to know anything at all about the economic conditions affecting the markets I traded or about geopolitical factors, monetary or fiscal policies of various governments, crop reports or anything else – I merely implemented a set of trend following strategies which produced positive performance net of commissions, management fees (2% of AUM per annum) and the performance fee (20% of gross returns).

In spite of the very encouraging results attained by Galstar, unrelated developments frustrated the project's ultimate goal of establishing a full hedge fund structure⁶. A massive global financial crisis broke out in 2008 and some of Galstar's shareholders suffered significant losses in their core business as well as personally. Not only were they unable to come up with the promised seed investment of \$10 million, but some of them were obliged to redeem their investment in Galstar in 2009. Fortunately, their

⁶ Unlike Galstar, which was closed to new investors, a regulated hedge fund would enable us to offer the asset management service to outside investors.

share was taken up by other investors, and in this way I was at least able to secure the continuity of the track record.

I carried on looking for investors to seed the new fund, but in the aftermath of the financial crisis, investor appetite for start-up hedge funds fell to about zero. In addition to losses sustained in global stock market collapse, investor confidence was run into the ground as a result of another related calamity. Namely, Madoff Securities, one of the largest and most trusted money managers in the world run by Bernard Madoff in New York, failed in 2008 and turned out to be nothing but a colossal Ponzi scheme through which Madoff and his accomplices managed to defraud investors out of some \$50 billion over the years.

This event poured cold water on the investor community globally and the seeding market⁷ for start-up funds froze almost entirely. Nevertheless, I kept reaching out to investors and one of the firms I contacted was Trafalgar Asset Managers run by Lee Robinson. While my first attempt to engage Lee didn't get any traction, several months later Trafalgar Asset Managers set up an office in Monaco, which was literally on the same block where I lived, as I discovered one day walking my dog past their office entrance. I recalled that this was Lee Robinson's firm, and he was one of the people I needed to get in front of.

I rang up Trafalgar and in a few days we set up a meeting. Lee looked at my performance track record and gave me an open minded hearing, but was unconvinced by my credentials in the industry and didn't think he could do much with my fund. So, that was that as far as business went, but being new in Monaco, Lee asked me if I knew any people he could join to play football or squash from time to time. This was very lucky as I needed a squash partner too, and from that time on we would get together weekly to play squash and usually chit chat about work, children, markets and the general situation in the world after the financial crisis. I discovered that Lee and I shared very similar worldviews and a rather pessimistic outlook on how the persisting global economic imbalances would be resolved.

Altana Inflation Trends Fund

Lee and I both believed that the financial crisis would ultimately lead to high inflation or hyperinflation of the US dollar and its ultimate demise as the world's reserve currency. In early 2011, Lee told me he would unwind Trafalgar Asset Managers and set up a new firm whose strategy would be focused on protecting investor wealth from the risks that were on the

⁷ Seeding market for hedge funds is the equivalent of venture capital market for startup technology firms. Seed investors typically allocate very substantial investments to new or start-up hedge fund teams in exchange for profit sharing or an equity stake in the firm.

horizon – primarily inflation – and that my trend following model would be an interesting part of that strategy.

Over the following months Lee took a closer look at the I-System and my track record, and together we deliberated a bit about the best way to proceed. His view, with which I fully agreed, was that inflation was one of the greatest risks to investor wealth. Our chief worry was not the normal kind of inflation that results from economic growth, but the malignant kind resulting from steady erosion in the purchasing power of the currency which could push prices of goods like food, energy and metals into uncharted territory. History shows that once inflation accelerates, the process can be self-reinforcing and take as much as a decade or more to run its course.

This erosion of a currency’s purchasing power tends to cause dramatic loss of investor wealth. Since 1960, over two thirds of the world’s market economies have suffered episodes of inflation which exceeded 25% in at least one year. On average, investors lost 53% of purchasing power during such episodes⁸.

Exhibit 2: Inflation can cause massive losses in real wealth

Inflation experience	10-yr inflation rate (annualized)	60/40 stock/bond portfolio real return (annualized)	Decline in real portfolio value
USA (1972 - 1982)	9%	-3.5%	-65%
UK (1910 - 1920)	11%	-9.3%	-86%
Japan (1946 - 1956)	23%	-3.3%	-52%

Source: Alliance Bernstein, “Deflating inflation – redefining the inflation-resistant portfolio,” April 2010.


As the above figures show, even moderately elevated inflation can cause very significant destruction of wealth. One of the logical consequences of the erosion of a currency’s purchasing power is the sustained rise in prices of commodities which can form strong trends. These can in turn be exploited through trend following, making that investment strategy the logical hedge against inflation. At any rate, this was the intuitive idea we were exploring at the time. Subsequently, we also found much empirical evidence to support our intuition. In its winter 2011 issue, the Journal of Wealth Management published a research paper which analyzed the performance of CTA investments from 1980 to 2011, along with other

⁸ Fischer, Stanley. “*Modern Hyper- and High Inflation*,” National Bureau of Economic Research working paper #8930

asset classes including equities, bonds, commodities and gold.⁹ Its authors conclude that, “*managed futures outperform other asset classes,*” and that, “*No other asset class presents itself as a viable inflation hedge.*”

Financial services firm Alliance Bernstein reached a similar conclusion by examining the inflation hedging effectiveness of asset classes like commodity futures, commodity stocks, equities, treasuries, treasury inflation-protected bonds (TIPS), precious metals futures, gold bullion, REITs and farmland. Alliance Bernstein’s report stated that, “*the investment that ranks best by far in terms of inflation beta is commodity futures.*”¹⁰ Inflation beta refers to an asset’s inflation sensitivity.

Exhibit 3: Asset class inflation beta (1965 – 2009)



Alliance Bernstein research indicates that commodity futures have the highest inflation beta – historically, for a 1% rise in inflation, commodity futures rose 6.5%.

In effect, a tail-event hedge¹¹

Thus, in order to provide a meaningful hedge against inflation for investors, an investment product would need to maintain strong exposure to commodity prices, and it would have to be relatively aggressive. It needs to be aggressive because investors typically allocate a rather small portion of their investable capital to managed futures products – typically

⁹ Twomey, J., J. Foran and C. Brosnan, “Assessing Managed Futures as an Inflation Hedge Within a Multi-Asset Framework.” *Journal of Wealth Management*, Winter 2011.

¹⁰ Alliance Bernstein, “Deflating Inflation: Redefining the Inflation-Resistant Portfolio.” April 2010.

¹¹ The term “tail event” refers to events of large magnitude (like stock market crashes) that occur very infrequently over time.

2-5%. Thus, if in a given year inflation jumped to 5%, an investor's 5% allocation to such a fund would need to earn a 100% return to offset the investor's loss of purchasing power caused by inflation. We also decided to allow our trading strategies to take short trades so as not to exclude the deflationary trends in our trading process. Since we could not predict when inflation might take off, nor dismiss the possibility of deflation in the nearer term, our fund would trade both on the long and short sides of price trends.

Agreeing broadly on the new fund's key investment objectives, we then selected a set of 28 commodity and financial markets to trade and I proceeded to design a portfolio that would have about 70% of risk focused on key commodity markets and the remaining 30% evenly split between equity futures and long-dated treasuries, deliberately excluding currency futures from the portfolio. The strong concentration on commodity futures would also differentiate the fund from the large, established CTA funds with several billion dollars under management, because these funds tend to hold no more than about 25% of their risk in commodities¹². For the very largest of them, exposure to commodity prices falls to well under 10% of their investment portfolios. We set up the fund's legal vehicle under Lee's new firm, Altana Wealth and named it Altana Inflation Trends Fund, or AITF.

Once everything was in place, Lee seeded it with a \$10 million investment and the fund launched in November 2011. Over the ensuing months and years we saw neither of the two tail events (stock market crash or an acceleration of inflation) came to pass with underwhelming results for the portfolio. Nonetheless I do expect that if either of the two events for which the portfolio had been conceived does ultimately materialize, the portfolio will yield the intended results. In every case where one of the markets included in AITF portfolio experienced a significant price readjustment, by virtue of trend following we were on the right side of the move. Thus, when prices gold and silver collapsed in the first half of 2013 our strategies all picked up the trend, capturing more than \$500 per ounce move in the price of gold. When oil prices collapsed in the second half of 2014 we likewise sailed with that trend from just above \$100/barrel to the very bottom of the trend below \$30/barrel. In early 2016 when equity prices corrected about 12%, AITF portfolio enjoyed a corresponding run-up of 19.6% net to investors. Again, these results were obtained strictly on the basis of simple disciplined trend following to the exclusion of all other inputs.


¹² According to the NewEdge Trend Indicator, the typical large CTA holds 25% of risk exposure in commodity futures, 15% in equity index futures, 30% in currencies and 30% in treasuries.

But apart from a handful of these market moves, commodity prices in general spent most of the time by far trading in a very unusual horizontal range, rendering the period from 2011 through 2016 a very difficult environment for trend followers. Indeed, by mid-2014 more than 150 CTAs went out of business, including a number of large funds with decades of successful trading behind them.¹³ We have been able to persevere and maintain our discipline through this adversity because we never lost conviction in two core ideas that shaped the fund's portfolio: (1) high inflation and possibly hyperinflation remains the most likely outcome of the current economic imbalances, ultimately making very significant price readjustments in commodity markets inevitable, and (2) these price readjustments will form major price trends over a multi-year period.

At risk of sounding like a proud parent fawning over his child, I must say that I remained very pleased with I-System's reliability as a trend-following auto-pilot. In every case where price trends *did* unfold, without exception AITF generated positive returns on those moves. Exhibits 6, 7, and 8 provide illustrations of some of those trends and our strategies' corresponding trading performance.

¹³ Madison, Marriage. "Trend-following hedge funds' future in doubt." Financial Times, 8th September 2014.


Exhibit 6: COMEX Silver (daily) with AITF daily gross performance per strategy [Nov. 2011 – Sept. 2014]



The eight Silver trading strategies in the AITF portfolio performed rather well during the 33-month period shown above. Some moves could definitely be characterized as trends, and with few exceptions, AITF strategies caught the moves and generated profits, which peaked in mid-2013, coinciding with the trough in the price of silver. The ensuing trendless period resulted in downward performance for all strategies.

When looking at a chart like the one presented in exhibit 6, some people will ask me why we use all these strategies when one of them (the one approaching a 350% return) has clearly proven to be a superior performer? My answer is that I would do that if I had a way of knowing in advance which strategy would perform best in the future. But I genuinely have no way of knowing this, so using a good variety of diverse strategies seems like the next best idea.

Exhibit 7: CSCE Coffee (daily) – daily gross performance per strategy [Nov. 2011 – Sept. 2014]




From November 2011 through August 2014, Coffee prices trended rather steadily. AITF's strategies handled the sporadic corrections and the early-2014 reversal as well as it can be expected from a trend following model and generated strong positive performance for almost full three years.

Coffee prices peaked just above \$3.00/lb in early 2011, subsequently falling to \$1.01/lb in 2013. This decline formed a trend that represented a nearly ideal trend following environment. After this trend bottomed in 2013, coffee prices bounced back strongly, almost doubling in the first half of 2014.

After profiting from the 2-year downtrend, AITF strategies reversed from short exposure to long early enough into the move to also profit from the 2014 rally to \$ 2.15. The ensuing correction and period of price consolidation between \$1.60 and \$2.00 led to negative performance and moderate drawdowns.

Exhibit 8: S&P 500 eMini futures (daily) – daily gross performance per strategy [Nov. 2011 – Sept. 2014]



The nearly linear ascent of the S&P 500 index made it the trend-follower's dream and all strategies gained sailing along with almost uninterrupted long exposure.

The nearly linear rally in the S&P500 futures hardly merits a comment. AITF strategies continued trading this market on the long side throughout, generating strong positive returns. Such favorable events represent the *feast* periods for trend followers, allowing us to sail also through inevitable *famines* of trend reversals and periods of sideways price consolidation.

Future applications

I-System's peculiarity is that it represents a knowledge framework within which we can formulate an almost infinite number of different trading strategies. Generally, the larger the portfolio of strategies and the greater diversity of uncorrelated markets they trade in, the more reliable and the less risky the investment process should become. A high concentration in a few fast-moving commodity markets will always get the most bang for the buck when conditions are favorable (i.e. prices trend strongly) but as we have seen, this is not always the case. Trends end, correct and reverse, and prices can spend a long time consolidating in a sideways range, so the feast periods are invariably followed by periods of trendless famines which are difficult for the traders and upsetting for investors. For this reason, most trend followers prefer to manage large, well-diversified portfolios with low drawdowns and low volatility of returns. This is also what investors prefer so it is a good combination for all sides.

For I-System, this would be a rather small and logical evolution. However, I believe that our model would grow into its full potential by supporting very large, diverse portfolios of global equities, implementing thousands of trading strategies. This would be a somewhat bigger challenge involving an incremental evolution of our technology's capabilities as it would likely entail a hybrid approach combining the key attributes of trend following and momentum investing, that could hopefully deliver moderate but steady returns on invested capital over decades, again without having to reinvent the wheel.

From today's perspective, I cannot predict whether we will be fortunate enough and capable enough to realize this potential, but it is human nature to always seek to transcend one's circumstances and reach for the stars, so we strive forward and press on with the present experiment. This story goes on.

Chapter 14: Trends and corporations

The responsible decisions in organized economic life are price decisions; others can be reduced to routine.

Frank Knight

When a management with reputation for brilliance tackles a business with reputation for poor fundamental economics, it is the reputation of the business that stays intact.

Warren Buffett

In 2010, AngloGold Ashanti, the world's third largest gold mining company accumulated \$2.47 billion in losses by hedging its exposure to the price of gold. Hedging involves the use of derivatives like futures, options or swaps to offset one's exposure to the price of some commodity, currency, or interest rate.

Thus, if a gold mining firm is concerned about the gold prices falling, they can lock in the current selling prices well into the future by selling futures or buying put options on gold. Such a bet would generate losses if the gold price rose. AngloGold Ashanti's hedge, put in place by the firm's then CEO Bobby Godsell, locked the firm into forward gold sales at an average price of less than \$450 per troy ounce during the time when the price of gold rose more than three-fold, reaching \$1,400 in 2010. Rather than enjoying record profits from the high price of gold, AngloGold Ashanti had to issue \$1.4 billion in new stock shares and convertible bonds to keep itself in business, diluting its shareholders in the process.

This was not a unique stroke of bad luck at one company. In 2013, Barrick Gold, the world's largest gold mining corporation posted a quarterly loss of \$8.6 billion when gold prices crashed from nearly \$1,700 per ounce to just over \$1,200. Barrick Gold's error was the opposite of AngloGold Ashanti's – they *did not* hedge their exposure to gold price and consequently suffered when the gold price collapsed.

Similar stories recur frequently in the corporate world, underscoring the fact that price risk can have a disproportionate impact on profitability for any business with significant exposure to commodity prices, whether it be gold, oil, copper, coffee, or some other commodity. The same is true for exposure to foreign currency and interest rates. This is not a terribly controversial assertion, but the extent to which this source of risk can affect a firm's profitability, shareholder value and competitive advantage is not sufficiently well appreciated.


Hedging and profitability among oil and gas producers

In 2006, I conducted a case study examining the hedging practices of nine independent North American oil & gas producers: Anadarko Petroleum, Apache Corporation, Burlington Resources, Canadian Natural Resources, Devon Energy, Encana, Kerr McGee, Talisman Energy and Unocal. Among these one firm, Kerr McGee (KMG), stood out as a particularly aggressive hedger. In 2003, KMG management decided to hedge 70% of their 2004 crude oil production, fixing the company's selling price at \$27.69 for their North American production. The other eight firms hedged on average only 25% of their 2004 production¹. As the oil price continued to rise through 2004, the year's average reached \$41.41 on the New York Mercantile Exchange.

Kerr McGee's management almost certainly pursued what they believed to be the prudent course of action. Their decision to aggressively hedge the firm's production may have had something to do with the long-term oil price forecasts produced at the time by the leading oil research institutes which unanimously predicted that the oil price would hold in the low-20s through 2005 as we saw in chapter 3. Whatever their reasons, the decision proved very unfortunate: by limiting their exposure to the favorable trend in oil price, Kerr McGee missed the opportunity to earn an extra \$13.72 in revenues per barrel. With 52 million barrels of crude oil produced in 2004, Kerr McGee's hedging deprived the firm of a \$500 million profit windfall.

¹ Marsh, Joe. "Crude Protection," Energy Risk (p. 30-32), April 2004.


Exhibit 1: KerrMcGee's \$500 million hedging decision



In 2003, Kerr McGee hedged 70% of their projected 2004 oil production, setting a price cap of \$27.69 – a decision that deprived the firm of a \$500 million windfall from the rising trend in oil prices, pushing the average in 2004 to \$41.41 per barrel

As bad as that was, Kerr McGee's underperformance was punished much more severely in the stock market: between October 2003 and August 2005, propelled by rising oil prices, shares of the oil and gas companies included in this case study appreciated on average by 250%. Over the same period, Kerr McGee underperformed its peers in the stock market by more than 30%, resulting in a full \$3.5 billion shortfall in the firm's market capitalization.

Exhibit 2: A rising trend in oil prices conferred value to oil and gas producers



Kerr McGee (KMG) share price compared to an index of other oil and gas producers. The eight firms comprising the index hedged on average only 25% of their production. Greater exposure to the oil price trend led to an average annualized return to shareholders of 63.3% for the period, while KMG, hedging 70% of its exposure generated "only" 44.6%.

Had Kerr McGee performed in line with its peers, the firm's market cap would have reached \$17.3 billion by August 2005, a full \$3.5 billion more than its then market cap of \$13.8 billion (at \$85.64 per share) in August 2005. In other words, Kerr McGee's hedging decision contributed to a destruction of shareholder value amounting to \$3.5 billion².

Market trends and value creation

Commodity prices perpetually fluctuate in the international markets. In the short-term, these fluctuations may appear news driven or even random, but major changes in market prices occur as trends which may span several months or even several years' time. For example, the price of gold rose for over ten years, from under \$260 per troy ounce in 2001 to nearly \$1,900 per troy ounce in 2011. Crude oil similarly rose from \$10 per barrel in 1998 to \$140 per barrel in 2008. U.S. ten-year Treasury Note yields declined from over 10% in the early 1980s to less than 2% in 2012.

Trends such as these undeniably represent an important source of opportunity for any business with significant commodity price risk. The same is true for foreign currency or interest rate risk. For such businesses, harnessing market trends is not at all incidental to value creation. To the contrary, it may be vital, as the previously cited examples of oil producers and gold mining firms illustrate. In a constantly changing economic environment, such cases are the rule rather than exceptions.

We have already considered evidence of this phenomenon in the McKinsey study I cited in chapter 6: through an exhaustive review of the performance of 100 of the largest U.S. corporations during the 1993-2003 business cycle, McKinsey found that 90% of firms that were able to grow at a rate higher than the GDP were concentrated in only four industries – those that benefited from favorable market trends during the period. Stated otherwise, favorable market trends were the key driver of value creation for 90% of the outperforming firms. The study's authors observed that, *“What's striking for large growth-minded corporations is just how crucial it is to have this kind of favorable wind at their backs when they try to achieve strong growth.”*

For firms whose performance depends on commodity prices, the best way to avail themselves of some “favorable wind at their backs,” is by incorporating skilful hedging operations as a part of their business process. If done properly, hedging can add value by enabling firms to keep

² It is true that hedging may only have played a part in Kerr McGee's valuation. However, given the magnitude of oil price rise between 2003 and 2005 and the windfall profits it brought to oil producers, we can safely assume that exposure to oil price played the predominant role in the oil and gas firms' stock market valuations.

exposure to commodity prices when they move in the firm's favor and limit this exposure when prices move in the firm's disfavor.

For example, had AngloGold Ashanti left their exposure to the soaring price of gold unhedged, it would have generated very large profits instead of a crushing loss from hedging. Conversely, had Barrick Gold hedged its production in 2013 before gold prices collapsed, it would have avoided some of that year's losses. Of course, this is all perfectly clear after the fact, but who knew gold prices would be rising for ten years, or that they would drop by almost \$500 per ounce in 2013³? Nobody could have known it for sure, but it is important to recognize that it is not necessary to know how the future will unfold to be able to take advantage of major shifts in commodity prices. Because major price readjustments unfold as trends sustained over a period of time, firms can use trend following techniques to manage their exposure to commodity prices.

I think it is no exaggeration to say that most worthwhile trend following systems would have suggested keeping long exposure during most of gold's ten-year uptrend, and short exposure when prices fell from \$1,700 to \$1,200 per ounce in 2013.

Yes, but we do not wish to speculate...

Managing commodity price risk is precisely the problem that defined my professional career. While working in oil trading for Greenoil, because profit margins on oil trading became paper thin, our performance depended largely on whether we got the price swings right or wrong. We were not alone in this predicament either. At one point in the year 2000, one of our clients – a midsize central European oil distribution firm – approached us asking for help managing their own oil price exposure.

As I was already fully focused on this problem matter, I volunteered to take on the assignment and after a few months I produced a detailed proposal specifying exactly how this firm could hedge their exposure, including a set of trend following strategies, the necessary organizational arrangements complete with an operations manual. Although I put a great deal of effort into this endeavor, to my knowledge our client never implemented the proposed process. All the same, the upshot of that experience for me was the realization of how important the problem of

³ Truth be told, in advance of the 2013 collapse in the price of gold, a handful of institutions like Goldman Sachs and Societe Generale issued remarkably accurate forecasts predicting the fall. Deutsche Bank at the same time forecast the average 2013 price of gold at \$1,637 per ounce and \$1,810 for 2014, highlighting again the problem with forecasts: they tend to be wrong as frequently as they are right and as such can't remove the uncertainty against which market participants must make forward looking decisions.

managing commodity price risk is for so many firms, and how difficult it is for most of them to address it adequately. Upon leaving Greenoil in 2003, I devoted considerable energy to marketing risk management services and consulting to firms in the energy and transport industries.

While most of the managers I spoke to had a keen interest in the problem of hedging, very few felt sufficiently comfortable to actually tackle the problem. The refrain I kept hearing was, “yes, but we do not wish to speculate.” What the managers usually wanted was some way to sidestep this risk by removing the possibility of adverse events while keeping some exposure to favorable developments. While this is technically achievable, the cost of this type of headache-free hedging can be prohibitively high and ultimately unfeasible.

Managers’ reluctance to get their hands dirty with hedging was the result of the skill gap between what is required to run a business operation and what is needed to speculate in commodity markets. Operating firms select for skills that support its core operations like production, marketing, sales or finance. By contrast, hedging requires a totally different set of skills related to market speculation, and these are not normally cultivated in operative businesses. As a result, managers often regard hedging with reserve and firms tend to manage their commodity price risk either passively or in some crude form⁴. Given all that can go wrong in speculation, this cautious stance is not entirely without merit. At the very least, it can keep the firm safe from self-inflicted risk-related disruptions and enable it to perform broadly in line with its rivals. That, however, is *all* that it will do.

The passive approach will not enable outperformance or competitive edge that could be achieved through a suitable solution to the problem of commodity price risk and uncertainty. As Thomas Aquinas put it, “*If the highest aim of a captain were to preserve his ship, he would keep it in port forever.*” Of course, people do not build ships just to preserve them. Likewise, managers should not run their businesses just to avoid risk, especially where risk – if adequately managed – has strong profit potential. This is clearly the case with the hedging of commodity price risk. It follows that such potential should at the very least be made a priority issue to explore at any commodity firm.

⁴ Active hedging also carries extra career risk. A manager who assumes an active role in hedging may not get rewarded if the results turn out good. But if they disappoint, he is likely to be accused of speculating recklessly, and his career will suffer as a result. For this reason, many managers see little upside from taking an active approach to hedging.

Hedging and competitive advantage

Skillful hedging can facilitate significant competitive advantages: for instance, if the prices of key inputs are rising, a company can secure a lower cost of material by buying adequate quantities for future delivery. Conversely, if the selling prices for its products are declining, it can fix higher prices by taking short positions in the futures markets. In this way, hedging can significantly improve operating profits.

Based on the income statement of an average S&P 1500 company (and assuming constant sales volumes), a 1% improvement in the selling price would generate an 8% increase in operating profits. Conversely, a 1% drop in the cost of goods sold would lead to a 5.36% increase in operating profits. This impact is more than double that of a 1% increase in sales volume⁵. For commodity businesses where operating margins are typically very low, hedging can have a much greater impact on profitability. It can also provide the most powerful means for a firm to differentiate itself from competitors and gain a difficult to match advantage over them.

Consider for example the U.S. petroleum wholesale industry where the operating margins average at about 0.8%⁶. At a cost of around \$800 per metric ton of heating oil, a typical wholesaler could hope to earn a margin of about \$6.40 per ton of heating oil sold. But suppose a wholesaler were able to reduce their cost of merchandise by an average of 1%⁷. That improvement would add \$8 per metric ton to the firm's profit margin, raising it from \$6.40 to \$14.40 – a 125% improvement in operating profits. Given such dramatic value creation potential, managers at commodity firms should not undiscerningly proscribe all speculation. After all, taking risks is what firms must do to generate profits.

In “Risk, Uncertainty and Profit,” Frank Knight compellingly argued that generating excess value from markets should be one of management's main concerns in running a business: “... *the most fundamental determining fact in connection with organization is the meeting of uncertainty. The responsible decisions in organized economic life are price decisions; others can be reduced to routine.*”⁸

Of course, the question remains whether achieving and sustaining this kind of competitive advantage is realistic. In all likelihood, it is not realistic to expect that hedging can achieve either permanently reduced

⁵ Mam, Michael, Eric Roegner and Craig Zawada. “The Power of Pricing” McKinsey Quarterly, 2003, Number 1.

⁶ This is based on a 2003 FirstResearch survey of 13,828 companies engaged in the petroleum wholesale business. Since then, margins may have shrunk still further.

⁷ An oil wholesaler could use hedging to its advantage by buying forward quantities of heating oil when the prices are advancing and keeping their position unhedged when they are declining.

⁸ Frank Knight, “Risk, Uncertainty and Profit” p 317.

cost of goods sold, or permanently higher selling prices. But at times when commodity prices undergo significant changes and form major trends, gaining a definite and significant price advantage through hedging is entirely realistic. The next question then is, how should firms go about exploring this potential advantage? Quality answers, I believe, would not fail to emerge through a systematic approach to problem solving and some organizational engineering. That may sound like a mouthful, but it is mostly just common sense stuff.

Organizing to manage uncertainty

A purposeful, effective approach to managing uncertainty and risk requires an adequate organizational framework. For any organization, the questions of what risks are taken, in what measure, and how they are managed are strategic questions and must be decided at the board level. The implementation of these decisions must be owned by the firm's CEO.

As with all risk management, managing commodity price, currency, or interest rate risk should enable a firm to take risks in a controlled and purposeful fashion, accept occasional losses and communicate such losses to its stakeholders openly and transparently, without losing stakeholder confidence in the validity of the firm's strategic choices or the management's capability to achieve them. Without clarity and guidance from the company's board and the CEO, the firm may be vulnerable to serious risk-related disruptions, or failure to take advantage of favorable market events.

Running a formal audit of key areas of risk exposure – by business unit and by risk category – should form the foundation of a firm's risk management process. For each category of risk, alternative instruments and methods of risk management should be identified and their respective advantages and disadvantages thoroughly examined and documented. Having evaluated the pros and cons of the available alternatives, management can formulate specific objectives and strategies to be implemented in achieving those objectives.

Definitive risk management strategies should set forth the company's risk management methods and its appetite for risk. It should also set out the responsibilities for risk management throughout the organization. At that point, management should anticipate the necessary organizational adjustments, training and staffing requirements and it should undertake a thorough documentation of the management process, controls, restrictions and paperwork flow. Finally, the whole solution, once implemented will almost certainly need adjustments and maintenance. Constant monitoring and periodic reviews must remain an integral part of a firm's risk

management strategy. For this purpose, firms should establish an independent middle office staffed with a team of highly skilled risk professionals who regularly report on exposure and risk issues directly to senior management and the CEO. The challenge of developing and implementing this business process should be no more difficult than that of developing any other business project.

The particular allure of hedging activities is that firms do not need to bet the proverbial ranch on it. At first, firms can apply their new risk management process only to a smaller portion of their risk exposure – say, 5% or 10% of their hedging requirements – and add to that in subsequent periods as the firm, its staff and stakeholders grow more familiar and comfortable with the process and its impact on the firm's performance.

For this to happen, the communications aspect of the project within the organization may be as important as its operational execution: all parties involved should be offered the opportunity to question and understand the process and be periodically kept informed about its progress and results. While the challenges involved aren't slight, the objectives and their potential should go far to kindle managers' entrepreneurial spirits and be well worth their efforts.

Chapter 15: Speculation and society

Speculation is dealing with the uncertain conditions of the unknown future. Every human action is a speculation in that it is embedded in the flux of time.

Ludwig von Mises

Speculation is a controversial subject evoking both fascination in some people and scorn in others. At one extreme, speculation – particularly in the form of trading and hedge fund management – has become one of the most prestigious career venues for many intelligent and aspiring young men and women. This has something to do with the growing social stature of many hedge fund managers. A few decades ago, these speculators were rather anonymous figures working in an obscure and little understood profession. Today, some of them have become prominent figures in the public eye. Their opinions and expertise are sought after by journalists, investors and politicians and their affluence provokes admiration and envy. However, not everyone admires speculators.

At the other extreme, they are frequently demonized for the destructiveness of their activities and their parasitic relation to society at large. Both points of view have some merit, but most people adhere to such views as personal convictions that aren't usually open to debate. As with most professions, traders and investment managers count among themselves capable and incapable individuals, smart and not so smart ones, honest and dishonest ones, the scrupulous and unscrupulous and everything in between. Setting aside for the moment the question of whether speculators are nice or rotten, we would do better to discuss the issue of speculation as a human pursuit and its role in society. It is a multifaceted subject and reducing it to simplistic good or bad labels isn't conducive to an intelligent debate about it.

Speculation is an inevitable part of human economic activity. Many of our ordinary decisions in life have an element of speculation in them. Do I buy a home, or do I rent? Do I get a job after school or do I go to university? Should I stay in my job, or start my own business? Do I buy flood insurance or not? Shall I save up to buy a tractor in cash, or do I lease it without delay? To the extent that such decisions deal in the present with uncertain outcomes in the future, they are speculative. But these mundane examples of speculation are hardly controversial.

It is when we engage in financial transactions for profit that the uneasy aspects of speculation emerge. This can be investing our savings in stocks and bonds, or buying a second or third house in order to resell at a higher price. The desire to profit from such transactions is often seen as motivated by greed – not commonly regarded as a positive human virtue. But greed isn't the only reason we might feel inclined to speculate. In fact, the nature of the modern monetary system compels us to seek returns on our savings. Nearly everyone understands that the money they save up is steadily losing value over time and that they can't keep it tucked under a mattress. Passively saving for a rainy day or for retirement simply isn't an option. Instead, we need our savings to earn interest and grow so that they won't lose their value in real terms.

However, over the last thirty years or so, interest rates have declined quite dramatically and have hit such low levels since the 2008 financial crisis that savings accounts and safe investments like government bonds pay extremely low interest if any at all. In 2014, some European banks actually started charging depositors negative interest rates.

In such an environment, savers (and their pension plans) are obliged to consider riskier and riskier investments, so that even the majority of people who are not motivated by greed are pushed into speculation just to try to keep their wealth from evaporating. To be sure, a small but interesting minority of people are positively motivated by greed and get in the game not just to earn a reasonable return on their savings, but because they are fascinated by the “game” and wish to excel in it, master it, and strike it rich. Today we live in such interesting times that high-rolling speculators are overtly glorified while their actions are somehow understood to grease the wheels of progress for mankind.

Some highly ranked captains of finance even believe themselves to be doing god's work here on Earth. While that may be an exaggeration, some credible sounding arguments advance the idea that speculators do play a positive role in society. They do so in two ways. In business school, they teach that speculators benefit society because they provide liquidity in securities markets and free up risk capital for more constructive pursuits. The case is often made using the example of the farmer who may obtain a

good price for his crops weeks, months, or even years in advance of the actual harvest. Rather than waiting for his crop to mature, he can sell it to a speculator today, lock in an advantageous price, and use the capital to upgrade or expand his business operations.

But speculators tend to benefit society in another important way which is seldom mentioned in business school. Because most speculators by far ultimately lose much of their money through speculation they also unwittingly provide the productive sectors of the economy with considerable amounts of low-cost financial capital.

To be beneficial, speculation must be regulated

In nature, many otherwise beneficial systems and processes can become pathological if they became unbalanced or unrestrained. This is also true of speculators; by providing financial capital and liquidity in the economic system, they clearly play beneficial roles in society. However, in absence of an effective and vigorously enforced system of regulation, speculators – some of them at any rate – will seek to gain a systemic advantage in some sector of the economy. If they succeed, their activities will likely become pathological to the economy. To prevent this, regulation must keep speculators from colluding against the interests of society at large.

Importantly, no speculator should be allowed to play so large a role as to be able to materially influence the conduct of commerce in key commodities. Further, because speculation in commodities markets is a zero-sum game¹ it is imperative that all participants transact on a level playing field, so that none have privileged access at the expense of others. In other words, speculators can be beneficial to society to the extent that the markets in which they speculate are transparent, competitive, and strictly regulated.

Through much of the modern history of commodities markets in the Western world, this has generally been the case. However, starting in the 1990s, there have been concerted efforts to relax market regulations and render their enforcement ineffective. This has enabled some of the largest financial and trading organizations to gain an upper hand and increase their influence over many key commodity markets. Over the years, these organizations have used their political leverage to further weaken regulation, gaining increasing control over many markets. Furthermore, privileged access to nearly limitless ultra low cost credit has allowed them to buy out vast holdings in the energy and metals production and wholesale distribution infrastructure. This has given them a dangerous

¹ Meaning that what is gained by one side in the transaction is lost by the other

degree of control over the flow of these commodities through the economy and ultimately over their availability and price.

Market regulatory agencies – in many cases staffed and overseen by these same dominant firms’ alumni – have either looked the other way or approved these transactions. This has gradually brought society to a point where everyone else’s ability to participate in a fair, transparent market depends on voluntary self-restraint by these organizations. But these organizations are not accountable to the society, but only to their shareholders who demand profits, not fairness. Hoping that they will prove good stewards of the economy to everyone’s benefit is naïve and unrealistic. Unless there is a strong and resolute pushback against uncontrolled encroachment on markets by large speculators, these developments will erode the integrity of markets leading to crises and far reaching disruptions benefiting few at the expense of everyone else. Thus, an activity which can and should be beneficial to society has been allowed to grow into a pathological element which will prove detrimental to Western societies and their economies. As a relatively small-scale professional speculator, observing these events is quite disconcerting.

Efficient, transparent securities markets are the hunting grounds where I have honed my skills and where I earn my bread and butter. The destruction of these structures makes me – professionally speaking – an endangered species.

Chapter 16: Advice for the young at heart

Success is not final, failure is not fatal: it is the courage to continue that counts.

Winston Churchill

Success lies in perseverance; ceaseless, restless perseverance!

Baron Manfred Von Richtoven

Imagine this: you are the manager of a hedge fund with \$1 billion in assets under management and in addition to earning an annual 1-2% fee on those assets (\$10-20 million), you also get a 20% performance fee charged against the fund's gross returns. If you generated a 10% gross return in a given year (\$100 million), your earnings would rise to some \$30 or \$40 million. If your fund had \$10 billion in assets under management your compensation would be between \$300 and \$400 million. This would make your occupation exceptionally lucrative by any standards. Indeed, many hedge fund managers might tell you that among professions, hedge fund management offers the highest return on intellectual capital. Although I'm not entirely sure that this is true¹, hedge fund management certainly does offer very high returns on intellectual capital.

But money isn't the only allure of this occupation – its very nature is deeply appealing at many levels. Hedge fund managers often refer to their

¹ The reason why I'm not sure this is entirely true is that many writers, musicians and film makers have also earned wealth comparable to that of hedge fund managers, arguably also from a kind of "intellectual capital"

work as “the game,” and managing investment funds *is* a bit of a game. You are the player, and global markets are your playing field. By paying attention to the financial, social, economic, geopolitical and even environmental conditions in the world, understanding the historical patterns of events and identifying the right investment opportunities you might be able to enrich yourself while financially rewarding your investors. The game as such also appeals to the ego because your success implies that you have understood the complexities of your world and mastered a domain where only a very few people may qualify as your equals in this respect.

It is little wonder that becoming a hedge fund manager has become the foremost professional aspiration for many smart and ambitious young people. Unfortunately for most of them, the odds of ever becoming a hedge fund manager are quite slim. The odds of being a successful hedge fund manager are much slimmer still. So how should the young person cursed with this ambition advance towards his or her goal? As there is no simple recipe for success, I can only offer you a few bits of general advice.

Formal education

It would be difficult to affirm which field of formal education would best prepare you for a career in hedge funds. I think that the study of natural sciences like mathematics or physics should prove useful, but perhaps even more so studies like mechanical or software engineering, as they provide substantial training in both mathematics and in practical problem solving. My preference for these fields of study is largely due to my preference for a systematic, quantitative approach to investing. But many other areas could also prove valuable.

For example, a student of biochemistry could be valuable to an investment fund specializing in the biotech industry; a naval architect might have an important role in funds investing in shipyards or other sectors of the maritime industry; a student of fine arts could work at a fund investing in art; a military officer could be a valuable asset in defense sector investing, and so on. Examples and possibilities are as varied as the hedge fund industry itself, so an advanced education in any field could conceivably gain your entry into the investment management industry. This is probably true for economics and finance as well, but I tend to regard these fields as probably the least useful of all.

To a novice, it might seem that economics would equip you with the proper understanding of the economy and capital markets, but this is hardly the case, as I tried to elaborate in chapter 3. Worldwide, the dominant schools of economic thought tend to be ideologically biased and

they have regressed either into an elaborate apoloia of the entrenched economic and monetary systems, or have drifted toward extraneous intellectual contortionism producing papers with such incomprehensible titles as, “*The Two-Period Rational Inattention Model: Accelerations and Analyses*,” or “*Continuous Time Extraction of a Nonstationary Signal with Illustrations in Continuous Low-pass and Band-pass Filtering*”.² I wonder if such endeavors aren’t a waste of talent that could have achieved something really interesting in a different domain.

Informal education

A university degree in any subject only formalizes your education up to that point – it should by no means mark the end of your education. On the contrary, to become a successful speculator, you will need to commit to a lifetime of learning. Most investment managers spend the bulk of their work hours reading everything and anything that may be directly or indirectly relevant to their work. This includes economics (I do highly recommend studying economics, but without the imperative of earning a degree in it), finance, mathematics, biology, history, sociology, anthropology, psychology, philosophy, geology, climate science, and many other subjects.

I think of reading as mining for treasures – even if a book doesn’t offer direct advice about how to make money in securities markets (best avoid such books), a great variety of books hide nuggets of gold – unique and valuable insights that can shape and refine your understanding of the world, spark original ideas, or free you from misconceptions you might have held unknowingly. I’ve encountered such gems in texts on unlikely subjects such as quantum physics, neurology and even linguistics. In sum, any study or practice that sharpens your mind and contributes to the development of your character will improve your odds of being a successful speculator. That will also help you be more resilient and composed through the inevitable ups and downs of your career. One caveat though: prioritize reading the masters. There are too many books out there, so choose the most authoritative authors.

Keep a journal

As reading becomes one of your main day-to-day occupations, you will find yourself digesting an enormous amount of information, facts, stories, ideas, and theories. Many of them will intrigue you, but the day you

² These examples were provided by James Grant as cited in Dr. Marc Faber’s January 2011 “The Gloom, Boom & Doom Report.”

encounter them you might not know how to use them or explore them further. Write down in a notebook whatever seems relevant, or just intrigues you. Use a paper notebook – do not type stuff or speak into your computer or smart phone. As you might discover, some of the facts or ideas in your notes might encounter other facts or ideas you will learn in the future and spark something original and interesting in your mind. You'll also come to discover that you will have forgotten many of the things you wrote down, and that without that reminder in your journal, the interesting and potentially life-changing gem might become inaccessible, lost in the mushrooming hay-stack of information you will have processed through your mind.

Career path

If you are to become a hedge fund manager, perhaps the best start professionally would be to join a well established hedge fund firm or the asset management side of a major bank. Working for a well known organization will likely make a positive difference on your curriculum vitae in the future. Once you start working there, make yourself useful and most importantly, get yourself the best mentor you can. What I just wrote may seem like a major accomplishment in and of itself for a young person, but you must never think it impossible.

To be sure, you should be prepared for rejection, but you do not need to accept it. If you reach out and fail, give it a bit of time and then try again – and again – until you succeed. Your reaching out will not go unnoticed and your perseverance will help you stand out from the crowd. Over time anyone's circumstances may change and those changes might well create an opening for you.

Your job today is to continue working on yourself and advancing until such an opening comes your way. The better prepared you are at that time, the better you'll be able to take advantage of it, so never stop advancing, even if you can see no light at the end of the tunnel. And always continue reaching out, remembering that even the wealthiest, most charismatic of hedge fund managers are still only human beings who once upon a time were in your position. They *will* be sympathetic to your endeavors.

Mental posture

In a world where uncertainty plays so great a part as it does in our progressive private-property society, the virtue of truthfulness becomes the very pearl of character.

Frank Knight

A man's character is his guardian divinity.

Heraclitus

Your career in hedge funds might turn out to be smooth sailing to success. More likely however, success tends to arrive as a culmination of many years of concerted work. Along the way, you'll have good months and bad months, good years and bad years, some breakthroughs and some major flops. As the adage goes, hedge fund management is a marathon, not a fashion show, so cultivate your personal resilience, perseverance and humility. When things get ugly, these attributes will determine whether you have a future with investors or not.

During good times, try not to let your success go to your head. As every season has its stars, you might also have your season of stardom. But that too shall pass, and whatever wave lifted you high on its crest is likely to draw you back to the mean. This is generally how things are and while striving to better yourself, you should make yourself at home at that mean level. Keep in mind also, that your career is only a part of your life. Try to maintain a healthy balance between work, family and social life. I've seen it more than once that when a person's home life falls apart, their professional life soon follows suit.

Hygiene

Keep clean – always. By this I don't mean that you must wash regularly. What I mean is that you must not get involved in shady or illegal business at any point in your career. Besides being a successful speculator, you'll also need to market yourself to investors as a trustworthy steward of their assets. People and institutions who allocate tens or hundreds of millions of dollars to hedge funds usually conduct background checks on their

managers and these can be remarkably thorough. Involvement in any shady business will disqualify you from the game – possibly for good. When you market yourself, use truth as your sole currency of presentation.

At times, you might feel tempted to embellish something about yourself and your accomplishments. While you don't need to advertise your weaknesses and shortcomings, attempting to deceive prospective investors will prove counter productive – sooner or later. Lies somehow always end up sticking out like a sore thumb in your story, and once you say a lie, you can't unsay it. If you get caught in one lie, people will rightfully doubt everything else you tell them – especially the best parts of your story. This will nullify the value of your true qualities and achievements. As much as possible, avoid doing things that might turn out to be embarrassing – not only professionally, but in general. The ubiquity of social networks on the internet today makes it likely that such experiences will follow you for a long time like a “kick me” sign on your back. While you don't need to be a saint or invisible, it is no exaggeration to say that a clean reputation is your most precious personal asset so be doubly judicious about your business as well as social conduct.

Go systematic

Again, I personally prefer a systematic approach to investing. Discretionary decision making in speculation is a daunting challenge and the human mind – no matter how brilliant – may simply not be up to the task. For all the information and statistics we can digest about the markets, we can't hope to grasp their complexity in anything more than approximate terms. When we connect the dots, there's a chance that we connect the wrong dots and reach mistaken conclusions. No matter how hard we try to be right it is unrealistic to expect that we can accurately navigate a process that eclipses our ability to comprehend it by orders of magnitude. Of course, whether by chance or by design, we can always find managers with the winning hand.

Every season has its stars who, through a combination of smarts, ability and luck come up winners at any given time. But the story doesn't end there because managers face an additional challenge – one that's more important than markets. In addition to trying to understand markets, the manager must deal with himself or herself. A winning run is likely to have psychological effects on a manager – he may start to believe that he has a special gift, that he has mastered the game, and that he can do no wrong. Overconfidence is not a helpful trait in speculative trading. The manager's thinking can become entrenched and lazy and his risk management discipline sloppy. Conversely, if the manager suffers a losing streak, his

confidence is bound to suffer, he might be unsure about his judgment, hesitate about his moves and end up passing up good investments.

The burden of coping with two complex worlds – the external world of economics, finance, politics, assets, legal environments, quarterly results, and the internal world of knowledge, judgment, conviction, confidence and emotional states – is probably more than one man or woman can handle day in and day out and remain on a winning streak for very long. A systematic investment strategy can greatly unburden the decision maker.

A systematic strategy can help the manager focus his attention on a limited set of parameters and impose critical discipline on his decision making and risk management. Importantly, the manager can backtest a systematic strategy and measure its performance objectively. If a strategy's actual performance is meaningfully different from what was expected, the manager can investigate the source of discrepancy and refine the strategy with only limited losses. This valuable feedback loop is not realistically available to managers who process all the inputs in their heads before taking and executing speculative decisions. Empirical evidence confirms that systematic hedge funds are more resilient than those based on discretionary decision making. Analyzing a large sample of CTA funds between 1994 and 2009, Julia Arnold of the Imperial College in London found that systematic CTAs have a higher median survival horizon than discretionary CTAs: 12 years vs. 8 years.³ In other words, going systematic could extend the longevity of your career in hedge funds by 50%! So by all means, go systematic. And as you do so, learn about systems engineering and scrupulously follow best practices as discussed in chapter 11. It could make all the difference for you.

Don't trade

Even if they aren't dreaming of becoming hedge fund managers, too many capable men and women want to try their hands in trading as a personal challenge, to make extra returns from their savings, or even just for fun. My strong advice to them: don't. That adventure is very likely to turn into a colossal waste of your talents, time, and ultimately a lot of your hard earned cash. Not so long ago, I came across an amusing looking chart, "The learning curve of professions⁴," that even as a caricature fairly portrays what you might expect from a stint with trading. You'll learn the ropes quickly and your learning curve will flatten, as will probably your bank account and a few other aspects of your life.

³ Arnold, Julia: "Survival of Commodity Trading Advisors: Systematic vs. Discretionary CTAs" Imperial College London, June 2012

⁴ <http://www.tradingmemes.com/meme/learning-curve-professions-trading/>

deserves mention about here was Jesse Livermore. In 1929, he correctly predicted that the U.S. economy would experience a depression and that the stock market would collapse. Trading on his macro convictions, he became a star and hugely wealthy at the time when most investors took massive losses. However, by 1932 Livermore was declared bankrupt, and in 1940 he scribbled, "I'm a LOSER!" on the wall of his hotel room before putting a gun to his head. Jesse Livermore was neither the first nor the last successful speculator who got trampled by markets and ultimately chose to end his life.

Cultivate discipline

Your trading will not always be uniformly successful. When you're on a winning streak, you'll feel vindicated and enthusiastic and you'll regard your profits as proof that you have been right or that your model works. Losses will lead you to doubt your understanding of the markets, second-guess your convictions and suspect your models. Coupled with loss aversion, these doubts could make the urge to "do something" hard to resist. Be extra judicious about your actions at such times.

When you are in a rut, doing nothing might be better than thrashing around trying to dig yourself out of the hole. Strategy drift is one of the terminal diseases of investment funds. The cure to this disease is steadfast discipline and perseverance. This is all the more difficult if you happen to only trade your own money. When you trade for other investors, you do so on the understanding that you will faithfully implement a certain strategy on their behalf – that would be the strategy you sold them as their investment manager. Being accountable to others is a very effective way of maintaining discipline. Over time, investors will ask you how and why you made or lost money and you'll need to produce periodic reports and newsletters summarizing your activities.

Knowing that you are answerable to people who entrusted you with their hard-earned money isn't pleasant, but it is a good way to stick to the pre-set strategy. If you only trade for yourself, you can freely drift from idea to idea and you may end up doing things you wouldn't do with other people's money, like chasing after that lucky trade that could change everything or escalating risk far beyond what you'd do if there ever were a chance of your having to explain yourself. So especially when you're in a tough spot, cultivate your discipline and stick with it.

Money isn't everything

The learning curve of professions conveys another bit of wisdom: make sure you are cultivating other valuable skills and contacts that could

cushion your fall if things go badly with trading. If you can work as a math or economics professor, tennis coach or scuba-diving instructor, try to keep the wheels under those options greased. Hopefully you'll never need to use them other than for pleasure, but if you do, having that lifeline could make a huge difference in your life.

The unfortunate fact of the matter is that as a trader your profession is of little value to the society in which you live. Other skills are more transferable. If his restaurant fails, a chef can cook at a different restaurant or a school cafeteria and his experience – even failure – might be valued at another establishment. If you fail as a trader, your experience might only serve to make you look foolish. So by all means, do have a plan B and keep it in good repair. This too counts as risk management, the critical discipline you will have to cultivate if you choose to make speculation your profession.

Advice for organizations and nations

At risk of sounding immodest, I will say this: in building the I-System, my team and I have achieved something remarkable and very difficult, even though our departure point was essentially a blank slate. None of us had any special formation or experience in the subject matter. At the outset, our project appeared almost impossibly ambitious but as we put ourselves to work, problem after problem got solved and our understanding of what we needed to do evolved as the solution took shape. The fact that you may have a complex problem on your hands and that you may have no idea how to solve it – whatever the problem be – should not make you believe that a solution is out of your reach and that your only option is to buy it from some global corporation for an ungodly sum of money.

In most cases, you are likely to obtain a better result if you pool together a small team of talented people and give them resources and manoeuvring space to build a quality solution. They'll probably build it for a fraction of what the off-the-shelf product would cost you and that's only the first of your advantages.⁶

These young people will also have acquired a deep and accurate understanding both of the problem matter and of different ways to mitigate it. This will make them far better decision makers in the future and far

⁶ If you happen to be one of those people working on a complex problem be aware that you have one important disadvantage compared to large corporations, and it has nothing to do with the quality of your work: large corporations can bribe the decision-makers in your organization. This is an unfortunate reality of our world, so make sure as many people in your organization know and understand what you're doing. Communicate your work prolifically to make it difficult for your management to "buy from IBM..."

more committed to your organization. Instead of making yourself hostage to external providers, you'll be able to chart your own course without the need to negotiate interminable licensing agreements and turn over huge sums of money at every turn. In addition, by building your own solutions you might obtain new marketable products or spinoff projects if your solutions can be sold to other clients. You will also have cultivated a generation of leaders able to mentor new talent with the confidence of knowing that any problem can be met with a quality solution, and this in itself has an inestimable economic value.

Any practically solvable problem has no chance of remaining unsolved if you are determined to tackle it. Which leads me back to the profound wisdom of Johann Wolfgang Goethe, for I can think of no more appropriate words with which to end this book:

“Until one is committed, there is hesitancy, the chance to draw back — concerning all acts of initiative (and creation), there is one elementary truth that ignorance of which kills countless ideas and splendid plans: that the moment one definitely commits oneself, then Providence moves too.

All sorts of things occur to help one that would never otherwise have occurred. A whole stream of events issues from the decision, raising in one's favor all manner of unforeseen incidents and meetings and material assistance, which no man could have dreamed would have come his way.

Whatever you can do, or dream you can, begin it. Boldness has genius, power, and magic in it. Begin it now.”

Thank you.

If you read my book through to the end, you have done me an honor and I thank you from the heart. I've put over five years of my life into writing it, which usually happened in my kitchen, late into the night after tucking my children to bed. It does not take five years to write two hundred pages – I spent the bulk of that time refining the text to make it as readable, interesting and as free of errors, fluff and superfluous words as I knew how, so that reading it might be pleasant, profitable, and enlightening to the reader.

I've also spent more than half a year preparing the book for publishing. Namely, I worked on this book under the illusion that when finished, I'd be turning it over to a publisher of business books who would then produce the final product, market, distribute it, and do everything else that publishers do. Alas, I discovered that I was virtually invisible to traditional publishers: after several months of sending proposals to various publishers and awaiting their answer, I received only one reply: a no.

So at the end of this volume I will ask you, dear reader for a small favor; if you enjoyed this book, please take a moment to give it an honest review on Amazon because reader reviews are the most valuable currency for new authors. Thank you again.

