
VOLATILITY TRADING

Founded in 1807, John Wiley & Sons is the oldest independent publishing company in the United States. With offices in North America, Europe, Australia, and Asia, Wiley is globally committed to developing and marketing print and electronic products and services for our customers' professional and personal knowledge and understanding.

The Wiley Trading series features books by traders who have survived the market's ever changing temperament and have prospered—some by reinventing systems, others by getting back to basics. Whether a novice trader, professional, or somewhere in-between, these books will provide the advice and strategies needed to prosper today and well into the future.

For a list of available titles, visit our website at www.WileyFinance.com.



VOLATILITY TRADING

SECOND EDITION

Euan Sinclair

WILEY

Copyright © 2013 by Euan Sinclair. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

First edition published by John Wiley & Sons, Inc. in 2008.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Sinclair, Euan, 1969–

Volatility trading / Euan Sinclair. — Second edition.

pages cm. — (Wiley trading series)

Includes bibliographical references and index.

ISBN 978-1-118-34713-3 (cloth/website); ISBN 978-1-118-42044-7 (ebk);

ISBN 978-1-118-57487-4 (ebk); ISBN 978-1-118-41672-3 (ebk)

1. Options (Finance). 2. Hedging (Finance). 3. Futures. 4. Financial futures. I. Title. HG6024.A3S5623 2013

332.64'5—dc23

2012047062

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

To Ann—

Sometimes a trader wins much more than he deserves.

CONTENTS

	Acknowledgments	xi	
	Introduction to the Second Edition	xiii	
CHAPTER 1	Option Pricing	1	
	The Black-Scholes-Merton Model	1	<hr/>
	Modeling Assumptions	7	vii <hr/>
	Conclusion	11	
	Summary	11	
CHAPTER 2	Volatility Measurement	13	
	Defining and Measuring Volatility	13	
	Definition of Volatility	14	
	Alternative Volatility Estimators	20	
	Using Higher-Frequency Data	29	
	Summary	33	
CHAPTER 3	Stylized Facts about Returns and Volatility	35	
	Definition of a Stylized Fact	35	
	Volatility Is Not Constant	36	
	Characteristics of the Return Distribution	40	
	Volume and Volatility	43	
	Distribution of Volatility	45	
	Summary	46	

CHAPTER 4	Volatility Forecasting	49
	Absence of Transaction Costs	50
	Perfect Information Flow	50
	Agreement about the Price Implications of Information	50
	Maximum Likelihood Estimation	54
	Volatility Forecasting Using Fundamental Information	60
	The Variance Premium	62
	Summary	65
CHAPTER 5	Implied Volatility Dynamics	67
	Volatility Level Dynamics	70
	The Smile and the Underlying	80
	Smile Dynamics	82
	Term Structure Dynamics	90
	Summary	91
CHAPTER 6	Hedging	93
	Ad Hoc Hedging Methods	95
	Utility-Based Methods	96
	Estimation of Transaction Costs	109
	Aggregation of Options on Different Underlyings	113
	Summary	115
CHAPTER 7	Distribution of Hedged Option Positions	117
	Discrete Hedging and Path Dependency	117
	Volatility Dependency	123
	Summary	129
CHAPTER 8	Money Management	131
	Ad Hoc Sizing Schemes	131
	The Kelly Criterion	133
	Time for Kelly to Dominate	143
	Effect of Parameter Mis-Estimation	144

	What is Bankroll?	146
	Alternatives to Kelly	148
	Summary	161
CHAPTER 9	Trade Evaluation	163
	General Planning Procedures	164
	Risk-Adjusted Performance Measures	171
	Setting Goals	178
	Persistence of Performance	180
	Relative Persistence	180
	Summary	184
CHAPTER 10	Psychology	187
	Self-Attribution Bias	191
	Overconfidence	193
	The Availability Heuristic	197
	Short-Term Thinking	199
	Loss Aversion	199
	Conservatism and Representativeness	201
	Confirmation Bias	203
	Hindsight Bias	206
	Anchoring and Adjustment	207
	The Narrative Fallacy	208
	Prospect Theory	209
	Summary	212
CHAPTER 11	Generating Returns through Volatility	213
	The Variance Premium	214
	Reasons for the Variance Premium	220
	Summary	222
CHAPTER 12	The VIX	223
	The VIX Index	224
	VIX Futures	225
	Volatility ETNs	227

	Other VIX Trades	229
	Summary	230
CHAPTER 13	Leveraged ETFs	231
	Leveraged ETFs as a Trade-Sizing Problem	234
	A Long-Short Trading Strategy	234
	Options on Leveraged ETFs	235
	Summary	237
CHAPTER 14	Life Cycle of a Trade	239
	Pretrade Analysis	239
	Posttrade Analysis	245
	Summary	247
CHAPTER 15	Conclusion	249
	Summary	252
	Resources	253
	Directly Applicable Books	253
	Thought-Provoking Books	256
	Useful Websites	257
	References	261
	About the Website	273
	About the Author	279
	Index	281

ACKNOWLEDGMENTS

A project like this invariably could not be accomplished without many forms of help and support from many people. Although they had no direct involvement in the writing, I need to first thank my parents, who made me appreciate education and taught me that boredom is a mortal sin. I must also thank my wife for her patience as I wrote this instead of doing various chores around the house.

xi

For more direct professional help, the regulars at www.nuclearphynance.com have been very generous with their time and knowledge. Their help with finding various research papers has also been invaluable.

Finally, Chris Merrill, Jitu Patel, and Erin Callahan gave direct help with proofreading and technical assistance. Thank you.

Since the first edition of *Volatility Trading* was published, the volatility markets have changed. One might think that the big change was the exceptionally volatile period at the end of 2008, followed by the long, slow decline of volatility since then. Although this has certainly made trading challenging, markets have behaved like this before and no doubt will again. This kind of change is always happening. More interesting has been the trend for exchanges to list volatility-dependent products (futures, exchange-traded funds [ETFs], and exchange-traded notes [ETNs]). This gives us new ways to make volatility trades and relative value bets. In this edition, we look at the Volatility Index (VIX) futures and ETNs, as well as leveraged ETFs.

Another major change has been the increasing effects of high frequency trading. Traders who are aggressively market-making equity options, trying to be on every bid and offer, will have to take such algorithms into account. These strategies are not covered in this book. My focus is on positional trading of volatility, and these trades generally have holding periods of days to months. These timescales are far from high frequency.

■ About This Book

The concepts behind this type of trading are largely unchanged, but my emphasis has probably changed slightly. Although academics have continued to work on forecasting volatility, I think the improvements from the perspective of an option trader are now marginal. For example, the profitability of equity volatility long short portfolios has declined significantly

over the past five years. This would probably not be what I would choose to focus on now. So although I have slightly expanded the section on volatility forecasting, I think it is now more important than ever to look for situations where the odds are in our favor whether we can accurately measure and forecast volatility. To this end, I have added chapters on stylized facts of markets and the variance premium.

I have also expanded the chapter on psychology. The usefulness of behavioral psychology to traders is still a matter of debate. Traders themselves seem to be greatly polarized, either being true believers or completely dismissive. I make no apologies for being a believer. I know that studying this has helped me.

I see trading as a process and I've tried to write this book to be consistent with this view. It was intended to be read in the order presented. However, most chapters are self-contained and there is probably no great harm in jumping around from section to section. An exceptionally impatient reader could even gain much of the benefit of the material by just looking at the summary at the end of each chapter. Some more tangential material, mainly mathematical in nature, has been relegated to sidebars. These can probably be skipped entirely without losing continuity.

This book is about trading volatility. More specifically, it is about using options to make trades that are primarily dependent on the range of the underlying instrument rather than its direction.

Before discussing technicalities, I give a brief description of my trading philosophy. In trading, as in most things, it is necessary to have general guiding principles in order to succeed. Not everyone need agree on the specific philosophy, but its existence is essential. For example, it is possible to be a successful stock market investor by focusing on value-style investing, buying stocks with low price-to-earnings or price-to-book ratios. It is also possible to be a successful growth investor, buying stocks in companies that have rapidly expanding earnings. It is *not* possible to succeed consistently by randomly acquiring stocks and hoping that things just work out.

I am a trader. I am not a mathematician, financial engineer, or philosopher. My success is measured in profits. The tools I use and develop need only be useful. They need not be consistent, provable, profound, or even true. My approach to trading is mathematical, but I am no more interested in mathematics than a mechanic is interested in his tools. However, a certain level of knowledge, familiarity, and even respect is needed to get the most out of these tools.

There is no attempt here to give a list of trading rules. Sorry, but markets constantly evolve and rules rapidly become obsolete. What do not become

obsolete are general principles. These are what I attempt to provide. This approach isn't as easy to digest as a list of magic rules, but I do not claim markets are easy to beat, either. The specifics of any trade are always different, but general guidelines can always point us somewhat in the right direction. Some latitude in strategy is desirable and adaptability is essential, but there are also a number of things that have to be firmly in place in order to succeed. Picasso and Braque may have broken many rules, but they could certainly paint very well technically before they did so. Similarly, before you start adjusting strategies, make sure you have a good grasp of the fundamental aspects on which all trades need to be based: edge, variance, and appropriate size.

Certain old-school traders have used arguments like: "Trading is about humans. Your models can't capture the human element." This generally seems to be said in a defensive manner. Maybe *their* models can't capture the human element, but ours will capture at least part of it. Most of the reluctance of such traders to embrace quantitative techniques can be attributed to defensiveness and aversion to change. It probably isn't due to any deep aversion to quantification. After all, in the same way that traditional baseball people hate the new statistical analyses but are fine with batting average and earned run averages, many traditional option traders denigrate quantitative analysis while being perfectly happy with the Black-Scholes-Merton paradigm and the concept of implied volatility. They are probably just unwilling to admit that they need to continue to learn and are worried that their skills are becoming obsolete. They should be. We all should be. This is a continually evolving process.

However, when successful traders say something like this, we need to consider that they may be partly correct. Some traders do indeed have finely honed intuition, generally called *feel* when applied to market sense. Intuition exists and can even be developed, but generally not quickly. Also, just because some traders have *feel* does not mean all, or even many, do. The approach we develop based on mathematics and measurement can be systematically learned. Given that it can be learned, what excuse is there not to learn it? Further, while a logic-based trader may never be able to develop effective intuition, an intuitive trader can always benefit from logic-based reasoning.

Although markets are designed and populated by human traders, with their typical human emotions and foibles, there is no justification for using this as a reason to avoid quantification and measurement. Baseball is also a game played by humans, and batting average is a useful way to measure the quality of a hitter. Similarly, before making a trade we need to be able

to somehow quantify the level of risk we will be incurring and the amount of edge we expect to gather. This is exactly what mathematics is good for. Estimating return and risk (however we define it) is purely a mathematical task. If something cannot be measured it cannot be managed. Further, if the human element is going to be important to our success, it will need to have measurable effects. The markets may indeed be driven by animal spirits, but I will remain thoroughly agnostic until they turn into poltergeists and start to actually throw prices around.

Pragmatism must always be our guiding philosophy. When I have had to choose between including something because I have found it useful, or omitting it because I could not prove it, I have tended to err on the side of inclusion. Successful trading is based on making correct decisions under conditions of uncertainty and incomplete information. There will always be things that we suspect are true but cannot prove. Waiting for proof may well mean waiting until the methods are no longer useful.

There are almost certainly other ways to trade options successfully. What I offer is *a* way, not *the* way. It is very much a data-driven, statistically oriented approach that should be applied over a wide range of products. But even traders who focus on one or two markets should be able to find some things useful and directly applicable. Traders who do not trade options should nevertheless find most aspects of the book useful as well.

The companion website for this book includes spreadsheets that illustrate some of the ideas presented in the book, along with the inevitable errata.

■ The Trading Process

Trading can be broken down into three main areas: finding profitable trades, managing risk and bankroll, and psychology. There is little to be gained by arguing over their respective levels of importance. Most traders will be more proficient in one of the three aspects than the others, but all three must be present for a trading operation to be a success.

When trading options, finding an edge involves forecasting volatility and understanding how volatility determines the market price of options. This means that we need a model for translating between price space and volatility space. Over the past 40 years, traders and financial engineers have proposed a number of option pricing models of varying complexity. We choose to use the Black-Scholes-Merton (BSM) methodology. Traders have learned to think in BSM terms. As a trader once said to me, “I want a model that a lot of guys have blown up using.” He meant that a good model

was one whose weaknesses were well known and had been discovered by someone else's misfortune, rather than one whose weaknesses have yet to be discovered. (Ironically, the same trader later blew up using the BSM paradigm. So it goes.)

There is a misconception that more complex models are better. But it doesn't matter how complex the model is. If a trader sells an option for 5.0 and buys it back for 3.0, he makes two ticks no matter what model he is using. A model is just a way to formulate our thoughts and translate between our volatility forecasts and the option prices. If traders are comfortable with a stochastic volatility model, they are more than welcome to use it. However, I have found the BSM framework is robust enough to be modified so that it is more representative of reality while still remaining simple and intuitive.

Although most stock and commodity options are U.S. style and hence not technically priced using this equation,¹ knowledge of the derivation of the pricing equation is necessary to get any feel for options. In our derivation we will emphasize the elements that we later hope to profit from. The market thinks in Black-Scholes terms, and to trade against it we need to understand what it really means. In this way trading is like a debate: In order to sensibly disagree with someone, we need to at least understand what they are saying.

Our derivation of this model in Chapter 1 is very informal. It directly proceeds from the starting point of holding a directionally neutral portfolio and shows how adjusting this dynamically leads to the BSM equation. It also makes clear the direct dependence of the equation on the range of the underlying and how this is proxied by volatility of returns. We also emphasize all the approximations and assumptions that are needed to arrive at the BSM equation. The rest of this book shows in detail how to deal with and profitably trade these inadequacies.

The largest source of edge in option trading is in trading our estimate of future volatility against the option market's estimate. Before we can forecast volatility we need to be able to measure it. In Chapter 2, we look at methods of historical volatility measurement including close-to-close volatility, Parkinson volatility, Rogers-Satchell volatility, Garman-Klass volatility, and Yang-Zhang volatility. We discuss the efficiency and bias of each estimator and also how each is perturbed by different aspects

¹There are several BSM equations. The differential equation certainly applies to U.S. options. The closed-form solutions to this equation, which are also called BSM equations, do not.

of real markets, such as fat tails in the return distribution, trends, and microstructure noise. We discuss different frequencies of measurement.

Next we look at how volatility actually behaves by examining its time series properties and how it relates to the returns of the underlying. We see that volatility clusters, exhibits mean *reversion* and *seasonality* and that it displays persistent correlation with returns and volume.

Then we try to forecast the volatility that will be present over the lifetime of the trade. We look at simple moving window forecasts, exponentially weighted moving averages, and various members of the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) family. But for trading we need more than a point estimate of future volatility. We need some estimate of the possible range of volatilities so we can make sensible judgments about the risk/reward characteristics of prospective trades. To find this we examine the construction and sampling properties of volatility cones.

Although our focus is to look for situations where implied volatility is at variance with our forecast of realized volatility, dynamics of the implied volatility surface are also interesting and important. An understanding can help our trade execution and timing. In Chapter 5, we look at normal shapes of the volatility surface both over time and by strike. We examine implied skewness and its sources, including credit, actual skewness of returns, put buying as static hedges, call buying as takeover hedging, and index skewness from implied correlation. We extend the Black-Scholes paradigm to include skewness and kurtosis and provide several rules of thumb for comparing volatilities across time and underlying product.

To profit from our forecast of volatility we need to hedge, so that our risk is actually realized volatility. Hedging removes the risks that we do not wish to take. We wish to accept risks that we believe to be mispriced and eliminate—or at least mitigate—our exposure to other sources of risk. With the simple, exchange-traded options that we generally consider, these unwanted risks are the drift of the underlying and movement in interest rates. Hedging is costly but it reduces risk. So when exactly should we hedge? In Chapter 6, we examine how to optimally solve this risk/reward issue. We also look at how to aggregate our positions to further reduce the need for hedging.

Once our position is hedged, what can we expect to happen? Chapter 7 examines the profit-and-loss distribution of a discretely hedged position and shows how this changes depending on the volatility we choose to use for delta estimation and the particular path taken by the underlying.

This completes the first stage in the trade process: finding a trade with a positive expectation. Now we need to look at how portfolio management choices can affect our success.

Chapter 8 demonstrates how different choices for trade sizing can dramatically affect returns. We introduce Kelly betting and compare it with other schemes, such as fixed-sized trading and proportional sizing. We also note how the sizing decisions affect risk by looking at risk of ruin and drawdowns. This is initially done for the simplest possible case, a trade that has a binary outcome. This is a long way from being even a partially realistic model of reality, but it is necessary to start with such simple examples, as even traders with many years of experience seem to have little if any idea of the implications of trade sizing. They obviously are aware that it is better to play a game with a positive edge many times so as to take advantage of the law of large numbers, but rarely take their understanding beyond this level. Futures traders seem to know more about this than options traders. Gamblers know even more. Most of the research in this area has been done by gamblers, particularly blackjack players. (Generally, it seems that the more complex the financial product, the less complex the actual trading process, ranging from the very complex blackjack strategy and sizing decisions to the trading of structured derivatives, where most of the edge lies in pricing and sales.)

Volatility trading is not binary in outcome. We need to extend Kelly to deal with situations that have a continuum of outcomes. (We really just need to extend the generally used version of the Kelly criterion. The Kelly criterion itself is far more general than the version that is often presented.) Further, volatility is a mean-reverting process. We must again extend our sizing methodology to account for this and show by simulations how this leads to familiar (to market makers) and simple scaling rules.

We also present some alternatives to the Kelly paradigm that may be more applicable to trading situations where the long run is of less interest than the short term. Traders should be aware of these methods. People who allocate capital to traders should be even more aware of them. Generally traders and trading firms will have somewhat different sets of priorities here.

To distinguish our results from chance we need to keep careful track of the results of our trades. In particular we need to be aware of much more than total profit and loss. This is particularly important in evaluating the efficacy of a new trade. Chapter 9 examines a number of measures including win/loss ratios; drawdowns; Sharpe, Calmar, Sortino ratios, and K-ratios; and the omega measure. Unfortunately, this type of record keeping and posttrade analysis is often left undone. I believe that this is the most important aspect of trading, and the most often overlooked. How can you improve if you don't really know what your results are?

Psychology is often mentioned in trading books. But successful trading is emphatically not due to being confident, reading the market, or having

no fear (although I was recently told that this was why a particular trader was good). This book does not go into this self-help style of psychology. Most of the psychological topics dealt with in books for amateurs or semiprofessionals can be addressed by sound money management rules and a sensible means of finding and measuring edge. However, knowledge of behavioral finance can be useful to even experienced traders. In Chapter 10, we look through some of the cognitive and emotional biases option traders will most often experience, both from a defensive and offensive viewpoint: things to watch out for to avoid hurting our own trading, and things to look for in the market that can be profitably exploited. Most sources of edge exist because of some behavioral aspect of psychology.

Before we can successfully trade volatility we must be aware of the variance premium: the tendency for index-implied volatility to be higher than the subsequent realized volatility. This can either be traded as a stand-alone strategy or used as the basis for a relative value idea like dispersion trading. Even if we don't want to be short volatility we need to understand the effect and be aware of the hurdles long volatility strategies need to clear.

Most of this book is about trading options. These are the most widely available and liquid volatility products. However, there are other volatility-dependent exchange-listed products available. VIX futures provide a clean, simple way to bet on implied volatility but they also present some unique problems. We examine these futures and a basis trading strategy that uses them. We next look at implementing a similar trade with some volatility ETNs, specifically VXX and VXZ, the iPath short and medium-term S&P 500 implied volatility ETNs

An interesting and commonly misunderstood way to gain exposure to realized volatility is through trading leveraged ETFs. We look at their options like characteristics and also the options that are listed on these. Options on UVXY (the double leveraged VIX ETN issued by ProShares) might possibly have exposure to more types of volatility than any other product. They depend on the implied volatility of UVXY, the realized volatility of VXX, the term structure of the VIX futures, the realized volatility of the VIX and the implied and realized volatilities of the S&P 500. There are often good trades to be found inside this Russian doll of volatility.

Finally, we examine one relative value volatility trade in detail, going through its complete life cycle from conception to expiration.

VOLATILITY TRADING

Option Pricing

It is possible to trade options without any valuation model. For example, traders might buy a call option because they think the underlying will rally further past the strike than the price they have paid. This is the simplest, most direct use of options. At a level of complexity only slightly greater than this we can trade volatility without a model. Traders might sell a straddle because they think the underlying will expire closer to the strike than the value of the straddle. There are an enormous number of option positions like this where traders can attempt to profit from their opinion of the future distribution of the underlying. However, if we want to express an opinion based on the behavior of the underlying before expiration, we will need a model.

A model is a framework we can use to compare options of different maturities, underlyings, and strikes. We do not insist that it is in any sense true or even a particularly accurate reflection of the real world. As options are highly leveraged, nonlinear, time-dependent bets on the underlying their prices change quickly. The major goal of a pricing model is to translate these prices into a more slowly moving system.

A model that perfectly captures all aspects of a financial market is probably unobtainable. Further, even if it existed it would be too complex to calibrate and use. So we need to somewhat simplify the world in order to model it. Still, with any model we must be aware of the simplifying assumptions that are being used and the range of applicability.

■ The Black-Scholes-Merton Model

We will present an analysis of the Black-Scholes-Merton (BSM) equation. The BSM formalism becomes the conceptual framework for an options

trader: In the same way that we hear our thoughts in English, an experienced derivatives trader thinks in the BSM language. This is an important difference between the models used by traders and the models used in a hard science such as physics. Models in physics aim to make statements about the world that are at least in some sense true, and then use the model to make predictions. The degree of truth needn't be consistent between all models. There are some successful theories that are based on highly simplified phenomenological models. An example would be Rutherford's model of the atom, which assumes that electrons orbit the nucleus like planets orbit the sun. This contains some truth: The atom consists of electrons and nuclear particles, but the planetary model isn't an accurate depiction of atomic structure.

Trading models are fundamentally different. The BSM model isn't good because it is an accurate representation of reality. It is actually fairly poor in this regard, with most of the model's assumptions being gross oversimplifications. It is a good model because the weaknesses are well understood and the model gives results that are intuitively sensible. The model fits its purpose. It is useful. It makes as little sense to say it is correct or incorrect as to say that German is incorrect and French is correct.

The standard derivation of the BSM equation can be found in any number of places (for example, Hull 2005). Although good derivations carefully lead us through the mathematics and financial assumptions they don't generally make it obvious what to do as a trader. We must always remember that our goal is to identify and profit from mispriced options. How does the BSM formalism help us do this?

Here we approach the problem backward. We start from the assumption that a trader holds a delta-hedged portfolio consisting of a call option and Δ units of short stock. We then apply our knowledge of option dynamics to derive the BSM equation.

That this portfolio is delta hedged should be obvious to option traders. Actually, traders knew about delta hedging long before BSM (for an interesting history, refer to Haug 2007a). But even if this is the first derivation of BSM the reader has seen this shouldn't be a remarkable fact. A call (put) option gains (declines) in value as the underlying rises. So in principle we can offset this directional risk with a position in the underlying. This should be obvious. The details of exactly how much of the underlying to hold are certainly not obvious.

Even before we make any assumptions about the distribution of the underlying's returns, we can state a number of the properties that an option must possess. These should be financially obvious.

- A call (put) becomes more valuable as the underlying rises (falls), as it has more chance of becoming intrinsically valuable.
- The value of a call (put) can never be more than the value of the underlying (strike).
- An option loses value as time passes, as it has less time to become intrinsically valuable.
- An option must have positive dependence on uncertainty. If the underlying had no risk there would be no need to pay for a product that only has value in certain states of the world. Options only have value because we are uncertain about the future, so it follows that the more uncertain we are the more valuable the options will be.
- An option loses value as rates increase. Because we have to borrow money to pay for options, as rates increase our financing costs increase, ignoring for now any rate effects on the underlying.
- Dividends (and storage or borrowing costs) have different effects on calls and puts. The holder of an option does not receive the dividend. This means that a dividend lowers the effective value of the underlying stock for the purposes of option valuation. So a dividend increases the value of a put and lowers the value of a call.

As we have said, even before the invention of the BSM formalism, option traders were aware that directional risk could be mitigated by combining their options with a position in the underlying. So let's assume we hold the delta-hedged option position,

$$C - \Delta S_t \tag{1.1}$$

where C is the value of the option, S_t is the underlying price at time, t , and Δ is the number of shares we are short. Over the next time step the underlying changes to S_{t+1} . The change in the value of the portfolio is given by the change in the option and stock positions together with any financing charges we incur by borrowing money to pay for the position.

$$C(S_{t+1}) - C(S_t) - \Delta(S_{t+1} - S_t) - r(C - \Delta S_t) \tag{1.2}$$

To see why the last term is positive we need to consider our cash flows. We bought the option, so we need to finance that cost, but we shorted stock so we receive money for this. Over a single time step we gain $r\Delta S_t$ from this.

Note also that we assume that the time step is small enough that we can take delta to be unchanged.

The change in the option value due to the underlying price change can be approximated by a second-order Taylor expansion. Also we know that when “other things are held constant,” the option will decrease due to the passing of time by an amount denoted by θ .

At this point in our argument we have assumed that we need to consider second derivatives with respect to price but only first derivatives with respect to time. Why is either of these choices valid? Ignoring higher derivatives with respect to price really cannot be justified at this point. We have only done it because we are trying to recover the BSM equation. In a more formal derivation this would be related to the assumption of a normal distribution of underlying returns. This is a major simplification that I am not ignoring. I’m postponing the discussion until later. The assumption that we need fewer derivatives with respect to time is easier to justify. Underlying price changes are stochastic and so they are a source of *risk*. Time change is predictable and the effect of time on options is merely a *cost*.

So we get

$$\Delta(S_{t+1} - S_t) + \frac{1}{2}(S_{t+1} - S_t)^2 \frac{\partial^2 C}{\partial S^2} + \theta - \Delta(S_{t+1} - S_t) - r(C - \Delta S_t) \quad (1.3)$$

or

$$\frac{1}{2}(S_{t+1} - S_t)^2 \Gamma + \theta - r(C - \Delta S_t) \quad (1.4)$$

where Γ is the second derivative of the option price with respect to the underlying. Equation 1.4 gives the change in value of the portfolio, or the profit the trader makes when the stock price changes by a small amount. It has three separate components.

1. The first term gives the effect of gamma. Since gamma is positive, the option holder makes money. The return is proportional to half the square of the underlying price change.
2. The second term gives the effect of theta. The option holder loses money due to the passing of time.
3. The third term gives the effect of financing. Holding a hedged long option portfolio is equivalent to lending money.

Further, we see in the next chapter that on average

$$(S_{t+1} - S_t)^2 \cong \sigma^2 S^2$$

where σ is the standard deviation of the underlying’s returns, generally known as *volatility*. So we can rewrite Equation 1.4 as

$$\frac{1}{2}\sigma^2 S^2 \Gamma + \theta - r(C - \Delta S_t) \quad (1.5)$$

If we accept that this position should not earn any abnormal profits because it is riskless and financed with borrowed money, the equation can be set equal to zero. Therefore, the equation for the fair value of the option is

$$\frac{1}{2}\sigma^2 S^2 \Gamma + \theta - r(C - \Delta S_t) = 0 \quad (1.6)$$

Before continuing, we need to make explicit some of the assumptions that this informal derivation has hidden.

- To write down Equation 1.1, we needed to assume the existence of a tradable underlying asset. In fact we assumed that it could be shorted and the underlying could be traded in any size necessary without incurring transaction costs.
- Equation 1.2 has assumed that the proceeds from the short sale can be reinvested at the same interest rate at which we have borrowed to finance the purchase of the call. We have also taken this rate to be constant.
- Equation 1.3 has assumed that the underlying changes are continuous and smooth. And as we mentioned earlier, we have considered second order derivatives with respect to price but only first order with respect to time. This is a very limiting assumption and will be returned to in some depth.

But something that we haven't made any assumptions about at all is whether the underlying has any drift. This is remarkable. We may naively assume that an instrument whose value increases as the underlying asset rises would be dependent on its drift. However, the effect of drift can be negated by combining the option with the share in the correct proportion. As the drift can be hedged away, the holder of the option is not compensated for it. Later in the chapter on hedging we see that in the real world, where the assumptions about continuity fail, directional dependence reemerges.

However, note that although the price change does not appear in Equation 1.6, the square of the price change does. So the magnitude of the price changes is central to whether the trader makes a profit with a delta-hedged position. This is true whether returns are normally distributed or not. This result holds as long as the variance of returns is finite. In fact if we had included higher order price terms in the Taylor expansion, we would see that the option's price change also depended on higher order price differences.

With appropriate final conditions, Equation 1.6 holds for a variety of instruments: European and American options, calls and puts, and many exotics. It can be solved with any of the usual methods for solving partial differential equations. The closed forms for these solutions (when they

- **[HTML5: The Missing Manual \(Missing Manuals\) online](#)**
- [read online The Stash Plan: Your 21-Day Guide to Shed Weight, Feel Great, and Take Charge of Your Health](#)
- [download online Into the Silence: The Great War, Mallory, and the Conquest of Everest](#)
- [download Diffusion MRI: Theory, Methods, and Applications](#)

- <http://anvilpr.com/library/Red-Queen--Red-Queen--Book-1-.pdf>
- <http://ramazotti.ru/library/The-Stash-Plan--Your-21-Day-Guide-to-Shed-Weight--Feel-Great--and-Take-Charge-of-Your-Health.pdf>
- <http://damianfoster.com/books/Shards-and-Ashes.pdf>
- <http://junkrobots.com/ebooks/Man-of-Smoke.pdf>