

WILEY FINANCE

With CD-ROM




Principles of quantitative development

MANOJ THULASIDAS

Principles of Quantitative Development

Manoj Thulasidas



 **WILEY**

A John Wiley and Sons, Ltd., Publication

This edition first published 2010

© 2010 Manoj Thulasidas

Registered office

John Wiley & Sons, Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com.

The right of the author to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. 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 or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Thulasidas, Manoj.

Principles of quantitative development / Manoj Thulasidas.
p. cm.

ISBN 978-0-470-74570-0 (hardback), ISBN 978-0-470-66737-8 (ebk),

ISBN 978-0-470-97151-2 (ebk), ISBN 978-0-470-97152-9 (ebk)

1. Speculation. 2. Risk management. I. Title.

HG6015.T48 2010

332.1068'1-dc22

2010005779

A catalogue record for this book is available from the British Library.

ISBN 978-0-470-74570-0

Set in 10/12pt Times by Aptara Inc., New Delhi, India

Printed in Great Britain by TJ International Ltd, Padstow, Cornwall, UK

Preface

Times have changed. Not long ago, all that you needed for a successful career in the financial industry were good people skills, and modest arithmetic and accounting abilities. The derivatives market requiring advanced quantitative (or computing) skills used to be miniscule. Now, with the explosive growth of derivatives, those simple times are gone for good. Derivatives and structures have taken the centre stage in recent years.

The world of trading and structuring today relies heavily on mathematical models and sophisticated computational techniques. In order to stay competitive, financial institutions need quantitative talent now. They cannot expect their bankers, specialized professionals themselves, to be adept in other professional domains like computer science and applied mathematics as well. Banks, therefore, import talent from other unorthodox domains. They find a helping hand among physicists and mathematicians who can model and price complex derivatives. They also employ a small army of computer scientists and programmers to deploy the pricing models in systems and platforms that help their traders make money.

This development has created a knowledge gap between the conventional banking staff and the newcomers. Bankers and regulators do not fully understand what the quantitative professionals are up to. The professionals themselves have very little appreciation for how the business side of the financial institution works. Bridging this gap is one of my objectives in embarking on this work.

Coming from a quantitative professional's perspective, I may have tilted this book a bit in our favour. That bias notwithstanding, I feel that it is important for 'quants' (the mathematicians who develop pricing models) and 'quantitative developers' (computer scientists who deploy the models) to understand and appreciate the business aspects of trading so that our efforts may be as fruitful as possible. Once, one of our top traders pointed out that even the most intelligent pricing model is useless unless it can be deployed effectively and in a timely manner. Effective deployment crucially depends on our appreciation of the 'Big Picture'.

Most quantitative professionals, especially at junior levels, are indifferent to the Big Picture. They think of it as a distraction from their real work of taming stochastic differential equations and writing programs. Through this book, I hope to change their mindset to some degree.

However, the scope of the book is bigger than the world of quantitative professionals. There is a flip side to our reluctance fully to appreciate the rest of the value chain. It is the inability of other business groups to realize what is possible and desirable in the quant world or in a trading system. Every business group in a bank is so specialized and requires skills and focus so specific to their duties that they all end up being distinct silos of knowledge. The professionals in other business units, especially in the middle office, will benefit from this book, in being able to formulate their requirements within the realms of feasibility. More than that, they will also gain an insight into how the quantitative work is done and how a mathematical model finally becomes profit-generating systems.

Another group of professionals who are supposed to know everything this book has to say are business analysts whose job it is to interface business needs to systems development. They usually work for technology consultancy firms and talk to their banking clients in order to come up with solutions specifically targeted to the needs of trading, financial computing and risk management. However, their existence is not confined to the technology side of the value chain. Some business associates and analysts work in financial institutions, translating internal business requirements to quantitative development objectives and rolling out new products and processes. These professionals will find this book invaluable as well.

When we change our narrow, albeit efficient, focus on the work at hand to an understanding of our role and value in the organization, we will be able to leverage off each other's work and enhance the overall productivity by leaps and bounds. This big picture view is especially relevant in the current climate (2008–2009) of financial turmoil and in our efforts to avoid or mitigate similar crises in the future. Of the many causes behind the turmoil that we see clearly with the benefit of hindsight, one that this book addresses is the general lack of a high-level understanding of the whole trading process. For instance, most quants, who are often blamed for the hole in which we find ourselves, do not fully appreciate the processes and methodology in managing the credit and market risks, or the interplay between them. Most professionals who manage credit and market risks do not quite understand the risk profiles straddling these two classes of risks that they manage independently of each other. Yet, this deficiency in understanding is neither deliberate nor premeditated. It is a natural consequence of the silo-like concentration and focus, which this book aims to change. In setting its sight at this holistic picture, this book indeed starts with an ambitious goal, but a timely one.

Manoj Thulasidas
Singapore

Contents

List of Figures, Tables and Big Pictures	xi
Preface	xv
1 Introduction	1
1.1 What is a trading platform?	1
1.1.1 Model archival	1
1.1.2 Incremental deployability	2
1.1.3 Live data feeds	2
1.1.4 Trade persistence	2
1.1.5 Regular processing	3
1.2 Quants and quantitative developers	3
1.3 Need for speed	4
1.4 Implementation options	5
1.4.1 Outsource to vendor	5
1.4.2 Use vendor API	5
1.4.3 Develop in-house	6
1.4.4 Replace vended systems	7
1.5 Current trends	7
1.6 Technical and business aspects of platform design	8
1.7 Importance of processes	10
1.8 Objectives and organization	12
Quiz	13
2 Overview of Banking	15
2.1 The offices	16
2.2 Front office	16
2.2.1 Economists	18

2.2.2	Structurers	18
2.2.3	Sales	18
2.2.4	Trading desks	19
2.2.5	Desk quants	19
2.2.6	Platform or quantitative developers	19
2.2.7	Desk risk management	20
2.3	Middle office	20
2.3.1	Product control	20
2.3.2	Treasury control unit	22
2.3.3	Market risk management	23
2.3.4	Credit risk management	23
2.3.5	Operational risk management	24
2.3.6	Rates management	26
2.3.7	Static data management	26
2.3.8	Compliance and reporting	27
2.3.9	Market risk management analytics	27
2.3.10	Asset and liability management	27
2.4	Back office	29
2.5	Supporting units	29
2.6	Summary	30
	Quiz	31
3	Trade Life Cycle	33
3.1	Pre-trade activities	34
3.2	Inception	35
3.3	Validation	38
3.4	Regular processing	40
3.5	Life-cycle events	43
3.6	Termination and settlement	45
3.7	Post-trade Activities	47
3.8	Summary	47
	Quiz	48
4	Trade Perspectives	49
4.1	Trade-centric view	49
4.2	Model-centric view	51
4.3	Product-centric view	53
4.4	Asset-class view	54
4.5	Queues and status flags	55
4.6	Aggregate views	57
4.7	Bottom-line view	60
4.8	Other perspectives	61

4.9	Summary	62
	Quiz	63
5	Programming Languages – Basics and Choices	65
5.1	Language choice	66
5.2	Basics of computing	67
5.2.1	Development cycle	67
5.2.2	Dependencies and Makefiles	69
5.2.3	Lexical structure of a programming language	70
5.3	Object-oriented languages	75
5.3.1	Basics of object-oriented methodology	75
5.3.2	Advantages and disadvantages of the object-oriented approach	78
5.4	Functional programming	79
5.5	Summary	84
	Quiz	85
6	Trading Platform Design	87
6.1	General design considerations	87
6.1.1	Maintainability and documentation	88
6.1.2	Scalability and extensibility	89
6.1.3	Security and access control	90
6.2	Architecture components	91
6.2.1	Trade representation	92
6.2.2	Market representation	94
6.2.3	Static data management	97
6.2.4	Quant library	97
6.2.5	Trade transformations and operations	98
6.2.6	Settlement triggers or pathways	100
6.2.7	Batch processing and grid computing	102
6.2.8	Credit replication	106
6.2.9	Security and audit model	107
6.2.10	Documentation and support model	110
6.3	Example architecture	111
6.4	Advanced architecture	113
6.5	Summary	117
	Quiz	117
7	Computing Patterns for Trading	119
7.1	Facade pattern	121
7.2	Visitor pattern	123
7.3	Singleton pattern	129

7.4	Factory patterns	131
7.5	Component object model	133
7.6	Summary	134
	Quiz	135
8	Flexible Derivatives Pricing Tool	137
8.1	Design scope	137
8.2	Design goals and features	138
8.2.1	Pricing tool features	138
8.2.2	Data types support	140
8.3	User interface	141
8.3.1	Main control interface	143
8.3.2	Create/edit product	145
8.3.3	Create/edit model	146
8.3.4	Generating a function template	147
8.3.5	Complex parameter visualization	150
8.3.6	Pricing interface	154
8.3.7	Visualization interface	156
8.3.8	Finite difference engine	159
8.4	Summary	161
	Quiz	163
9	Pricing Tool to Trading Platform	165
9.1	Pricing tool: internals	165
9.1.1	Common quant classes	167
9.1.2	Main interfaces	170
9.1.3	Pluggable pricing models	176
9.1.4	Reusable components	177
9.1.5	Source code documentation	179
9.1.6	External packages	179
9.2	Future enhancements	180
9.2.1	Trade database	180
9.2.2	Market data	181
9.2.3	Life-cycle management	190
9.2.4	Security and access control	190
9.2.5	Batch process	191
9.3	Summary	191
	Quiz	191
10	Summing Up	193
10.1	Epilogue	193
10.2	Further reading	195

10.2.1 Quantitative finance	195
10.2.2 Computing	199
10.2.3 Economics	201
Appendix	203
A.1 CD Contents	203
A.2 Historical perspective	204
Glossary of Terms	213
Index	227

List of Figures, Tables and Big Pictures

Figures

2.1	High-level overview of the banking structure	16
2.2	The structure of a typical front office	17
2.3	The structure of a typical middle office	21
3.1	Pre-trade activities of the trade life cycle	35
3.2	Trade inception activities	36
3.3	Trade validation process	38
3.4	Regular processes during the life of a trade	40
3.5	Trade termination activities	45
4.1	The model-centric view of trading activities	52
4.2	The product-centric view as in the trading platform	53
4.3	A sample scenario analysis of one trade	57
4.4	Middle-office process queues	58
5.1	Flowchart of a simple portfolio processing	74
5.2	Binomial tree pricing model	81
6.1	Settlement workflow through a vended trading system	101
6.2	Grid-enabled trading platform	104
6.3	Information flow in credit replication	107
6.4	An example architecture of a trading platform	112
6.5	A high-level schema of a professional architecture of a trading platform	114
7.1	Façade pattern in action	122
7.2	The UML diagram for a façade pattern	123
7.3	Visitor pattern in a trading platform	124
7.4	The factory method pattern	131
7.5	The abstract factory pattern	132

8.1	User interaction with the pricing tool	142
8.2	Main control interface of the pricing tool	143
8.3	Context menus in the pricing tool	144
8.4	Product definition in the pricing tool	145
8.5	Model definition in the pricing tool	146
8.6	Code generation in the pricing tool	148
8.7	Help on the code generation in the interface	150
8.8	Viewing Vector, Curve and Matrix	153
8.9	Plotting complex data types	154
8.10	Pricing interface of the attached program	155
8.11	Plotting interface (two-dimensional plots)	157
8.12	Plotting interface (three-dimensional plots)	158
8.13	Plotting interface (parametric plots)	160
8.14	Finite difference Greeks computation interface of the pricing tool	162
9.1	Use case depiction of the pricing tool	166
9.2	UML class hierarchy diagram of the pricing tool	168
9.3	UML diagrams of the product parameter and model output classes	169
9.4	Mapping of use cases to user interface modules in the pricing tool	172
9.5	Main user interface of the pricing tool and its active regions	174
9.6	Product editing interface showing its dynamic nature	175

Tables

4.1	Example of a trade scenario risk analysis	56
6.1	Components that make up the market data set	96
8.1	Example of Vector and DateVec data types	151
8.2	Flexible date input formats	152

Big Pictures

1.1	Conflicts by design	10
2.1	Operational risk in in-house platforms	25
2.2	Risk–reward seesaw	28
3.1	Collateral versus credit lines	36
3.2	Workflow engine	39
3.3	To hedge or not to hedge	42
3.4	Exchange-traded versus over-the-counter	44
4.1	Back-to-back trades	50
4.2	Pluses and minuses of VaR	58
4.3	Big Picture and narrow perspectives	62
5.1	Zeros and ones	66
5.2	Magic of object-oriented languages	78
5.3	Back to basics	83
6.1	Straight-through processing	101

6.2	Throwing the baby out with the bath water	105
6.3	Access control	108
6.4	Untenable supports	110
7.1	Paradigms in computing	119
7.2	Paradigms versus patterns	120
9.1	Market data – quantization errors	181
10.1	Effect of modelling on models	195
10.2	Group dynamics	199
10.3	Philosophy of money	202
A.1	Software nightmares	205

Introduction

Purpose-built trading platforms have become indispensable tools for exotic trading in most modern banks. One reason for their success and prevalence is that exotics and structured products business is highly lucrative to a financial institution. One prerequisite for success in exotic business is the ability to price the products. For pricing model development and hedging or risk analysis, the financial institution can resort to its mathematicians and the quantitative models they develop.

Another vitally important prerequisite is the capability to assess and manage the associated risks. The risks arising from market movements are related to pricing and are managed through sensitivity analyses. However, in order to manage credit and operational risks, especially when the business volume grows, banks need robust trading platforms supporting and augmenting the mathematical effort.

A carefully designed and deployed trading platform can easily recuperate its development costs within a few short months of going live, and turn into an important profit centre for banks. Once they started appreciating the necessity and profitability of an in-house trading system, banks and other financial institutions began to invest resources into developing in-house trading systems.

Resource allocation attracts professionals with highly specialized skills to move into the field. However, their skills, at times, work against them. They think of quantitative development as a purely technical endeavour. Although, at its heart, trading platform development is a technical challenge, it is the myriad of business processes around it that will eventually make or break its viability. The proverbial devil is really in the details.

1.1 WHAT IS A TRADING PLATFORM?

The goal of this book is to help design a robust, durable and reliable trading platform. The first question then is: What exactly is a trading platform (or system)? For our purpose in this book, I define it as a program or a collection of programs that meets the following broad requirements.

1.1.1 Model archival

One of the problems that the quantitative analysis effort in any modern bank faces is the archival and reuse of their pricing models. A trading platform acts as a repository for the quantitative intelligence generated within the bank. A model

developed for a particular product is all too often so specialized that it is difficult to deploy it to tackle another product. This lack of reusability results in duplicated effort. A well-designed trading platform can help by imposing the right quantum of structure and standards to encourage generic programming and to enforce reusability. Furthermore, it will hold all the pricing models with standardized interfaces in one place where they can be found, examined and reused.

1.1.2 Incremental deployability

A trading platform should provide means by which the new quantitative models developed or implemented in the bank can be easily integrated and deployed into the revenue generating work streams of the bank. In other words, even after the trading platform is commissioned, as new and innovative products are developed by the mathematicians and structurers, we should be in a position to augment our system to make use of the innovation. This requirement of incremental deployability drives the whole architecture and design of the trading platform, as you will see in later chapters. It also imposes the necessary rigidity in the form in which quantitative analysts deliver their pricing models, facilitating the first requirement.

1.1.3 Live data feeds

A trading platform should talk to the external world to obtain market data (either as live feeds or as snapshots at regular intervals) and archive snapshots of market conditions for bulk processing. Market data feeds reach the trading platform from several data providers who use dissimilar, and often incompatible, technologies and interfaces. For this reason, the market data handling may become another auxiliary project in its own right, providing a uniform interface to the trading platform regardless of the origin of the data provided. This market data project often has its own database back-end for data persistence.

1.1.4 Trade persistence

A trading platform should be able to save trade data into a database and manage all the associated inception and life-cycle events. The database layer of a trading platform has many stringent requirements on security, performance and record-keeping integrity. It should be capable of handling the inception events such as cancellation, cancel and amend, reissue, novation (changing the ownership) etc., consistent with the policies of the financial institution. A cancellation, for instance, is never a database operation of deleting a record, for a complete deletion would erase the necessary audit trails. In addition, the database layer may reside in geographically dispersed locations, and the trading platform may be called upon to provide data replication services that can be challenging when the trade volumes grow.

1.1.5 Regular processing

A trading platform should facilitate and mediate all aspects of regular downstream processing, such as risk management, trade transformation, settlements, etc. This requirement, although summarized in one bullet point, is very vast in its scope, as you will appreciate by the time you finish this book.

It is instructive to benchmark a vended system (such as the trading system currently in use in your bank) against these requirements. All vended systems do an admirable job in meeting the last two requirements. However, they fail miserably in the first two. In fact, the pricing models implemented in vended solutions are the principal intellectual properties that the vendors jealously guard. They have no incentive in facilitating easy deployment of in-house custom models using their trading systems.

The pricing tool that comes with this book is near the other extreme: it implements the first two requirements wonderfully. Using the pricing tool, you can define and deploy new models without even recompiling the core program. However, it has no market data feeds, not even a clear notion of a market, and it does not provide a database layer. The only data persistence it offers comes in the form of XML files that store pricing scenarios.

Most in-house trading platforms find their place in between these two extremes. They do not try to handle every aspect of trade life-cycle management, but they are far more than the prototyping or proof-of-concept pricing programs that the mathematical brains of the bank generate. Depending on the implementation strategy of a particular bank, the in-house trading platform may end up focusing on different aspects of the requirements listed above. In all cases, however, an in-house trading platform attempts to bridge the gap between quantitative pricing models and a deployed and supported system out of which the bank can generate profit. Its emergence has engendered the relatively new domain of quantitative development. Trading platforms are the domain of the quantitative developers (computer scientists who deploy the models) in the bank, who are distinct from the so-called ‘quants’ (mathematicians who develop pricing models).

1.2 QUANTS AND QUANTITATIVE DEVELOPERS

The term ‘quant’ is short for quantitative analyst. They are mathematicians who develop pricing models and keep abreast with the cutting edge research on stochastic calculus and quantitative finance. They transform the academic knowledge they assimilate into programs that can generate revenue and profit. Quants usually have an advanced degree in mathematics, physics or other quantitative fields. Ideally, quants have a sound knowledge of markets, business and products, as well as computer science.

The input to quants' work is research and the academia. Their outputs are pricing models, often delivered as programs or functions in a library, aptly called the quant library. They may also deliver standalone pricing programs, usually as spreadsheets and add-ins.

Quantitative developers are computing professionals who make the output from the quants widely usable, often through the in-house trading platform. Their functional duties fall in between the quant output and the programs used at trading desks. Ideally, quantitative developers would have the same skill set as the quants (mathematical aptitude, product/business knowledge and computer science), but with less emphasis on mathematics and much more focus on computing and software engineering. How separated quants and quantitative developers are in terms of job functions and organizational hierarchies depends on the human resourcing strategies of the bank.

In our discussions in this book, we will assume a clean separation between the duties of quants and quantitative developers. The developers start their work from the quant library, the fruits of the quant's labour. The end point of their work is an in-house trading platform and other maintainable quantitative tools.

1.3 NEED FOR SPEED

In today's financial markets, opportunities are transient. The skyrocketing commodity prices of early 2008, for instance, generated a strong demand for cost-effective hedging. Banks that could roll out structured products to meet customized hedging requirements reaped handsome benefits from this demand. Hedging essentially caps the customer's upside exposure while subjecting them to unlimited downside risks. Because of the dizzying fall in commodity prices (especially in energy) that soon followed during the second half of 2008, the banks that could provide the hedging structures again made even more profit. Rolling out such hedging solutions on short notice to benefit from the market fluctuations of this kind necessarily calls for the agility and flexibility that only an in-house trading system can provide.

Due to such enticing profit potential, an increasing amount of assets under management gets earmarked for exotics trading. In fact, this influx of institutional investments was at least partly to blame for the wild fluctuations in commodity prices. However, the influx also underscores the importance of exotic and structured products. Coupled with this emphasis on exotics trading is the increasing sophistication of the clients who demand customized structures reflecting their risk appetite and market views.

The net result of the changing financial market attitude is the need for more mathematical modelling and speedier deployment than ever before. The need for speed, in fact, goes beyond exotics business. In spot FX trading, the time scale of profit making opportunities is measured in seconds or milliseconds. The so-called