


A Practical Guide
to Algorithmic Strategies and
Trading Systems

HIGH-FREQUENCY TRADING



Irene Aldridge

High-Frequency Trading

*A Practical Guide to Algorithmic
Strategies and Trading Systems*

IRENE ALDRIDGE



WILEY

John Wiley & Sons, Inc.

Copyright © 2010 by Irene Aldridge. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
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 also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Aldridge, Irene, 1975–

High-frequency trading : a practical guide to algorithmic strategies and trading system / Irene Aldridge.

p. cm. – (Wiley trading series)

Includes bibliographical references and index.

ISBN 978-0-470-56376-2 (cloth)

1. Investment analysis. 2. Portfolio management. 3. Securities. 4. Electronic trading of securities. I. Title.

HG4529.A43 2010

332.64–dc22

2009029276

Printed in the United States of America

10 9 8 7 6 5 4 3

Acknowledgments

This book was made possible by a terrific team at John Wiley & Sons: Deb Englander, Laura Walsh, Bill Falloon, Tiffany Charbonier, Cristin Riffle-Lash, and Michael Lisk. I am also immensely grateful to all reviewers for their comments, and to my immediate family for their encouragement, edits, and good cheer.

Contents

Acknowledgments	xi
CHAPTER 1 Introduction	1
<hr/>	
CHAPTER 2 Evolution of High-Frequency Trading	7
<hr/>	
Financial Markets and Technological Innovation	7
Evolution of Trading Methodology	13
CHAPTER 3 Overview of the Business of High-Frequency Trading	21
<hr/>	
Comparison with Traditional Approaches to Trading	22
Market Participants	24
Operating Model	26
Economics	32
Capitalizing a High-Frequency Trading Business	34
Conclusion	35
CHAPTER 4 Financial Markets Suitable for High-Frequency Trading	37
<hr/>	
Financial Markets and Their Suitability for High-Frequency Trading	38
Conclusion	47

CHAPTER 5 Evaluating Performance of High-Frequency Strategies	49
<hr/>	
Basic Return Characteristics	49
Comparative Ratios	51
Performance Attribution	57
Other Considerations in Strategy Evaluation	58
Conclusion	60
CHAPTER 6 Orders, Traders, and Their Applicability to High-Frequency Trading	61
<hr/>	
Order Types	61
Order Distributions	70
Conclusion	73
CHAPTER 7 Market Inefficiency and Profit Opportunities at Different Frequencies	75
<hr/>	
Predictability of Price Moves at High Frequencies	78
Conclusion	89
CHAPTER 8 Searching for High-Frequency Trading Opportunities	91
<hr/>	
Statistical Properties of Returns	91
Linear Econometric Models	97
Volatility Modeling	102
Nonlinear Models	108
Conclusion	114
CHAPTER 9 Working with Tick Data	115
<hr/>	
Properties of Tick Data	116
Quantity and Quality of Tick Data	117
Bid-Ask Spreads	118

<i>Contents</i>	vii
Bid-Ask Bounce	120
Modeling Arrivals of Tick Data	121
Applying Traditional Econometric Techniques to Tick Data	123
Conclusion	125
CHAPTER 10 Trading on Market Microstructure: Inventory Models	127
<hr/>	
Overview of Inventory Trading Strategies	129
Orders, Traders, and Liquidity	130
Profitable Market Making	134
Directional Liquidity Provision	139
Conclusion	143
CHAPTER 11 Trading on Market Microstructure: Information Models	145
<hr/>	
Measures of Asymmetric Information	146
Information-Based Trading Models	149
Conclusion	164
CHAPTER 12 Event Arbitrage	165
<hr/>	
Developing Event Arbitrage Trading Strategies	165
What Constitutes an Event?	167
Forecasting Methodologies	168
Tradable News	173
Application of Event Arbitrage	175
Conclusion	184
CHAPTER 13 Statistical Arbitrage in High-Frequency Settings	185
<hr/>	
Mathematical Foundations	186
Practical Applications of Statistical Arbitrage	188
Conclusion	199

CHAPTER 14 Creating and Managing Portfolios of High-Frequency Strategies	201
<hr/>	
Analytical Foundations of Portfolio Optimization	202
Effective Portfolio Management Practices	211
Conclusion	217
CHAPTER 15 Back-Testing Trading Models	219
<hr/>	
Evaluating Point Forecasts	220
Evaluating Directional Forecasts	222
Conclusion	231
CHAPTER 16 Implementing High-Frequency Trading Systems	233
<hr/>	
Model Development Life Cycle	234
System Implementation	236
Testing Trading Systems	246
Conclusion	249
CHAPTER 17 Risk Management	251
<hr/>	
Determining Risk Management Goals	252
Measuring Risk	253
Managing Risk	266
Conclusion	271
CHAPTER 18 Executing and Monitoring High-Frequency Trading	273
<hr/>	
Executing High-Frequency Trading Systems	274
Monitoring High-Frequency Execution	280
Conclusion	281

CHAPTER 19 Post-Trade Profitability Analysis	283
<hr/>	
Post-Trade Cost Analysis	284
Post-Trade Performance Analysis	295
Conclusion	301
References	303
About the Web Site	323
About the Author	325
Index	327

Introduction

High-frequency trading has been taking Wall Street by storm, and for a good reason: its immense profitability. According to *Alpha* magazine, the highest earning investment manager of 2008 was Jim Simons of Renaissance Technologies Corp., a long-standing proponent of high-frequency strategies. Dr. Simons reportedly earned \$2.5 billion in 2008 alone. While no institution was thoroughly tracking performance of high-frequency funds when this book was written, colloquial evidence suggests that the majority of high-frequency managers delivered positive returns in 2008, whereas 70 percent of low-frequency practitioners lost money, according to the *New York Times*. The profitability of high-frequency enterprises is further corroborated by the exponential growth of the industry. According to a February 2009 report from Aite Group, high-frequency trading now accounts for over 60 percent of trading volume coming through the financial exchanges. High-frequency trading professionals are increasingly in demand and reap top-dollar compensation. Even in the worst months of the 2008 crisis, 50 percent of all open positions in finance involved expertise in high-frequency trading (Aldridge, 2008). Despite the demand for information on this topic, little has been published to help investors understand and implement high-frequency trading systems.

So what is high-frequency trading, and what is its allure? The main innovation that separates high-frequency from low-frequency trading is a high turnover of capital in rapid computer-driven responses to changing market conditions. High-frequency trading strategies are characterized by a higher number of trades and a lower average gain per trade. Many traditional money managers hold their trading positions for weeks or even

months, generating a few percentage points in return per trade. By comparison, high-frequency money managers execute multiple trades each day, gaining a fraction of a percent return per trade, with few, if any, positions carried overnight. The absence of overnight positions is important to investors and portfolio managers for three reasons:

1. The continuing globalization of capital markets extends most of the trading activity to 24-hour cycles, and with the current volatility in the markets, overnight positions can become particularly risky. High-frequency strategies do away with overnight risk.
2. High-frequency strategies allow for full transparency of account holdings and eliminate the need for capital lock-ups.
3. Overnight positions taken out on margin have to be paid for at the interest rate referred to as an overnight carry rate. The overnight carry rate is typically slightly above LIBOR. With volatility in LIBOR and hyperinflation around the corner, however, overnight positions can become increasingly expensive and therefore unprofitable for many money managers. High-frequency strategies avoid the overnight carry, creating considerable savings for investors in tight lending conditions and in high-interest environments.

High-frequency trading has additional advantages. High-frequency strategies have little or no correlation with traditional long-term buy and hold strategies, making high-frequency strategies valuable diversification tools for long-term portfolios. High-frequency strategies also require shorter evaluation periods because of their statistical properties, which are discussed in depth further along in this book. If an average monthly strategy requires six months to two years of observation to establish the strategy's credibility, the performance of many high-frequency strategies can be statistically ascertained within a month.

In addition to the investment benefits already listed, high-frequency trading provides operational savings and numerous benefits to society. From the operational perspective, the automated nature of high-frequency trading delivers savings through reduced staff headcount as well as a lower incidence of errors due to human hesitation and emotion.

Among the top societal benefits of high-frequency strategies are the following:

- Increased market efficiency
- Added liquidity
- Innovation in computer technology
- Stabilization of market systems

High-frequency strategies identify and trade away temporary market inefficiencies and impound information into prices more quickly. Many high-frequency strategies provide significant liquidity to the markets, making the markets work more smoothly and with fewer frictional costs for all investors. High-frequency traders encourage innovation in computer technology and facilitate new solutions to relieve Internet communication bottlenecks. They also stimulate the invention of new processors that speed up computation and digital communication. Finally, high-frequency trading stabilizes market systems by flushing out toxic mispricing.

A fit analogy was developed by Richard Olsen, CEO of Oanda, Inc. At a March 2009 FXWeek conference, Dr. Olsen suggested that if financial markets can be compared to a human body, then high-frequency trading is analogous to human blood that circulates throughout the body several times a day flushing out toxins, healing wounds, and regulating temperature. Low-frequency investment decisions, on the other hand, can be thought of as actions that destabilize the circulatory system by reacting too slowly. Even a simple decision to take a walk in the park exposes the body to infection and other dangers, such as slips and falls. It is high-frequency trading that provides quick reactions, such as a person rebalancing his footing, that can stabilize markets' reactions to shocks.

Many successful high-frequency strategies run on foreign exchange, equities, futures, and derivatives. By its nature, high-frequency trading can be applied to any sufficiently liquid financial instrument. (A "liquid instrument" can be a financial security that has enough buyers and sellers to trade at any time of the trading day.)

High-frequency trading strategies can be executed around the clock. Electronic foreign exchange markets are open 24 hours, 5 days a week. U.S. equities can now be traded "outside regular trading hours," from 4 A.M. EST to midnight EST every business day. Twenty-four-hour trading is also being developed for selected futures and options.

Many high-frequency firms are based in New York, Connecticut, London, Singapore, and Chicago. Many Chicago firms use their proximity to the Chicago Mercantile Exchange to develop fast trading strategies for futures, options, and commodities. New York and Connecticut firms tend to be generalist, with a preference toward U.S. equities. European time zones give Londoners an advantage in trading currencies, and Singapore firms tend to specialize in Asian markets. While high-frequency strategies can be run from any corner of the world at any time of day, natural affiliations and talent clusters emerge at places most conducive to specific types of financial securities.

The largest high-frequency names worldwide include Millennium, DE Shaw, Worldquant, and Renaissance Technologies. Most of the high-frequency firms are hedge funds or other proprietary investment vehicles

TABLE 1.1 Classification of High-Frequency Strategies

Strategy	Description	Typical Holding Period
Automated liquidity provision	Quantitative algorithms for optimal pricing and execution of market-making positions	<1 minute
Market microstructure trading	Identifying trading party order flow through reverse engineering of observed quotes	<10 minutes
Event trading	Short-term trading on macro events	<1 hour
Deviations arbitrage	Statistical arbitrage of deviations from equilibrium: triangle trades, basis trades, and the like	<1 day

that fly under the radar of many market participants. Proprietary trading desks of major banks, too, dabble in high-frequency products, but often get spun out into hedge fund structures once they are successful.

Currently, four classes of trading strategies are most popular in the high-frequency category: automated liquidity provision, market microstructure trading, event trading, and deviations arbitrage. Table 1.1 summarizes key properties of each type.

Developing high-frequency trading presents a set of challenges previously unknown to most money managers. The first is dealing with large volumes of intra-day data. Unlike the daily data used in many traditional investment analyses, intra-day data is much more voluminous and can be irregularly spaced, requiring new tools and methodologies. As always, most prudent money managers require any trading system to have at least two years worth of back testing before they put money behind it. Working with two or more years of intra-day data can already be a great challenge for many. Credible systems usually require four or more years of data to allow for full examination of potential pitfalls.

The second challenge is the precision of signals. Since gains may quickly turn to losses if signals are misaligned, a signal must be precise enough to trigger trades in a fraction of a second.

Speed of execution is the third challenge. Traditional phone-in orders are not sustainable within the high-frequency framework. The only reliable way to achieve the required speed and precision is computer automation of order generation and execution. Programming high-frequency computer systems requires advanced skills in software development. Run-time mistakes can be very costly; therefore, human supervision of trading in production remains essential to ensure that the system is running within

prespecified risk boundaries. Such discretion is embedded in human supervision. However, the intervention of the trader is limited to one decision only: whether the system is performing within prespecified bounds, and if it is not, whether it is the right time to pull the plug.

From the operational perspective, the high speed and low transparency of computer-driven decisions requires a particular comfort level with computer-driven execution. This comfort level may be further tested by threats from Internet viruses and other computer security challenges that could leave a system paralyzed.

Finally, just staying in the high-frequency game requires ongoing maintenance and upgrades to keep up with the “arms race” of information technology (IT) expenditures by banks and other financial institutions that are allotted for developing the fastest computer hardware and execution engines in the world.

Overall, high-frequency trading is a difficult but profitable endeavor that can generate stable profits under various market conditions. Solid footing in both theory and practice of finance and computer science are the normal prerequisites for successful implementation of high-frequency environments. Although past performance is never a guarantee of future returns, solid investment management metrics delivered on auditable returns net of transaction costs are likely to give investors a good indication of a high-frequency manager’s abilities.

This book offers the first applied “how to do it” manual for building high-frequency systems, covering the topic in sufficient depth to thoroughly pinpoint the issues at hand, yet leaving mathematical complexities to their original publications, referenced throughout the book.

The following professions will find the book useful:

- Senior management in investment and broker-dealer functions seeking to familiarize themselves with the business of high-frequency trading
- Institutional investors, such as pension funds and funds of funds, desiring to better understand high-frequency operations, returns, and risk
- Quantitative analysts looking for a synthesized guide to contemporary academic literature and its applications to high-frequency trading
- IT staff tasked with supporting a high-frequency operation
- Academics and business students interested in high-frequency trading
- Individual investors looking for a new way to trade
- Aspiring high-frequency traders, risk managers, and government regulators

The book has five parts. The first part describes the history and business environment of high-frequency trading systems. The second part reviews the statistical and econometric foundations of the common types of

high-frequency strategies. The third part addresses the details of modeling high-frequency trading strategies. The fourth part describes the steps required to build a quality high-frequency trading system. The fifth and last part addresses the issues of running, monitoring, and benchmarking high-frequency trading systems.

The book includes numerous quantitative trading strategies with references to the studies that first documented the ideas. The trading strategies discussed illustrate practical considerations behind high-frequency trading. Chapter 10 considers strategies of the highest frequency, with position-holding periods of one minute or less. Chapter 11 looks into a class of high-frequency strategies known as the market microstructure models, with typical holding periods seldom exceeding 10 minutes. Chapter 12 details strategies capturing abnormal returns around ad hoc events such as announcements of economic figures. Such strategies, known as “event arbitrage” strategies, work best with positions held from 30 minutes to 1 hour. Chapter 13 addresses a gamut of other strategies collectively known as “statistical arbitrage” with positions often held up to one trading day. Chapter 14 discusses the latest scientific thought in creating multistrategy portfolios.

The strategies presented are based on published academic research and can be readily implemented by trading professionals. It is worth keeping in mind, however, that strategies made public soon become obsolete, as many people rush in to trade upon them, erasing the margin potential in the process. As a consequence, the best-performing strategies are the ones that are kept in the strictest of confidence and seldom find their way into the press, this book being no exception. The main purpose of this book is to illustrate how established academic research can be applied to capture market inefficiencies with the goal of stimulating readers’ own innovations in the development of new, profitable trading strategies.

Evolution of High-Frequency Trading

Advances in computer technology have supercharged the transmission and execution of orders and have compressed the holding periods required for investments. Once applied to quantitative simulations of market behavior conditioned on large sets of historical data, a new investment discipline, called “high-frequency trading,” was born.

This chapter examines the historical evolution of trading to explain how technological breakthroughs impacted financial markets and facilitated the emergence of high-frequency trading.

FINANCIAL MARKETS AND TECHNOLOGICAL INNOVATION

Among the many developments affecting the operations of financial markets, technological innovation leaves the most persistent mark. While the introduction of new market securities, such as EUR/USD in 1999, created large-scale one-time disruptions in market routines, technological changes have a subtle and continuous impact on the markets. Over the years, technology has improved the way news is disseminated, the quality of financial analysis, and the speed of communication among market participants. While these changes have made the markets more transparent and reduced the number of traditional market inefficiencies, technology has also made available an entirely new set of arbitrage opportunities.

Many years ago, securities markets were run in an entirely manual fashion. To request a quote on a financial security, a client would contact

his sales representative in person or via messengers and later via telegraph and telephone when telephony became available. The salesperson would then walk over or shout to the trading representative a request for prices on securities of interest to the client. The trader would report back the market prices obtained from other brokers and exchanges. The process would repeat itself when the client placed an order.

The process was slow, error-prone, and expensive, with the costs being passed on to the client. Most errors arose from two sources:

1. Markets could move significantly between the time the market price was set on an exchange and the time the client received the quote.
2. Errors were introduced in multiple levels of human communication, as people misheard the market data being transmitted.

The communication chain was as costly as it was unreliable, as all the links in the human chain were compensated for their efforts and market participants absorbed the costs of errors.

It was not until the 1980s that the first electronic dealing systems appeared and were immediately heralded as revolutionary. The systems aggregated market data across multiple dealers and exchanges, distributed information simultaneously to a multitude of market participants, allowed parties with preapproved credits to trade with each other at the best available prices displayed on the systems, and created reliable information and transaction logs. According to Leinweber (2007), designated order turnaround (DOT), introduced by the New York Stock Exchange (NYSE), was the first electronic execution system. DOT was accessible only to NYSE floor specialists, making it useful only for facilitation of the NYSE's internal operations. Nasdaq's computer-assisted execution system, available to broker-dealers, was rolled out in 1983, with the small-order execution system following in 1984.

While computer-based execution has been available on selected exchanges and networks since the mid-1980s, systematic trading did not gain traction until the 1990s. According to Goodhart and O'Hara (1997), the main reasons for the delay in adopting systematic trading were the high costs of computing as well as the low throughput of electronic orders on many exchanges. NASDAQ, for example, introduced its electronic execution capability in 1985, but made it available only for smaller orders of up to 1,000 shares at a time. Exchanges such as the American Stock Exchange (AMEX) and the NYSE developed hybrid electronic/floor markets that did not fully utilize electronic trading capabilities.

Once new technologies are accepted by financial institutions, their applications tend to further increase demand for automated trading. To wit, rapid increases in the proportion of systematic funds among all hedge