Artificial Intelligence in Finance

Dr. Yves J. Hilpisch

ODSC East, Boston, 30. April 2019

THE PYTHON QUANTS

finaince

The Group

THE PYTHON QUANTS

http://tpq.io


16 week program
150+ hours of instruction
PROGRAM DIRECTOR
Dr. Yves J. Hilpisch is founder and managing partner of The Python Quants (http://tpq.io), a group focusing on the use of open source technologies for financial data science, algorithmic trading and computational finance. He is the author of the books
He is the author of the books:
Python for Finance (O'Reilly)
Derivatives Analytics with Python (Wiley)
Listed Volatility and Variance Derivatives (Wiley)
He has written the financial analytics library DX Analytics (http://dx-analytics.com) and organizes conferences and Meetup events about Python for finance and algorithmic trading in Frankfurt, London and New York. He has given keynote speeches at technology conferences in the United States, Europe and Asia.
The Python Quants GmbH
recognized by Banking CIO Outlook magazine as
TOP 10 BankingAnalytics SOLUTION PROVIDERS - 2017
An annual listing of 10 companies that are at the forefront of providing banking analytics solutions and impacting the marketplace
The Python Quants GmbH
66333 Voelklingen
Germany
T/F +49 3212 112 91 94
http://training.tpq.io
training@tpq.io
April 2017
UNIVERSITY CERTIFICATE IN PYTHON FOR ALGORITHMIC TRADING
THE PYTHON QUANTS
5,000+ lines of code
1,200 pages PDF
http://certificate.tpq.io

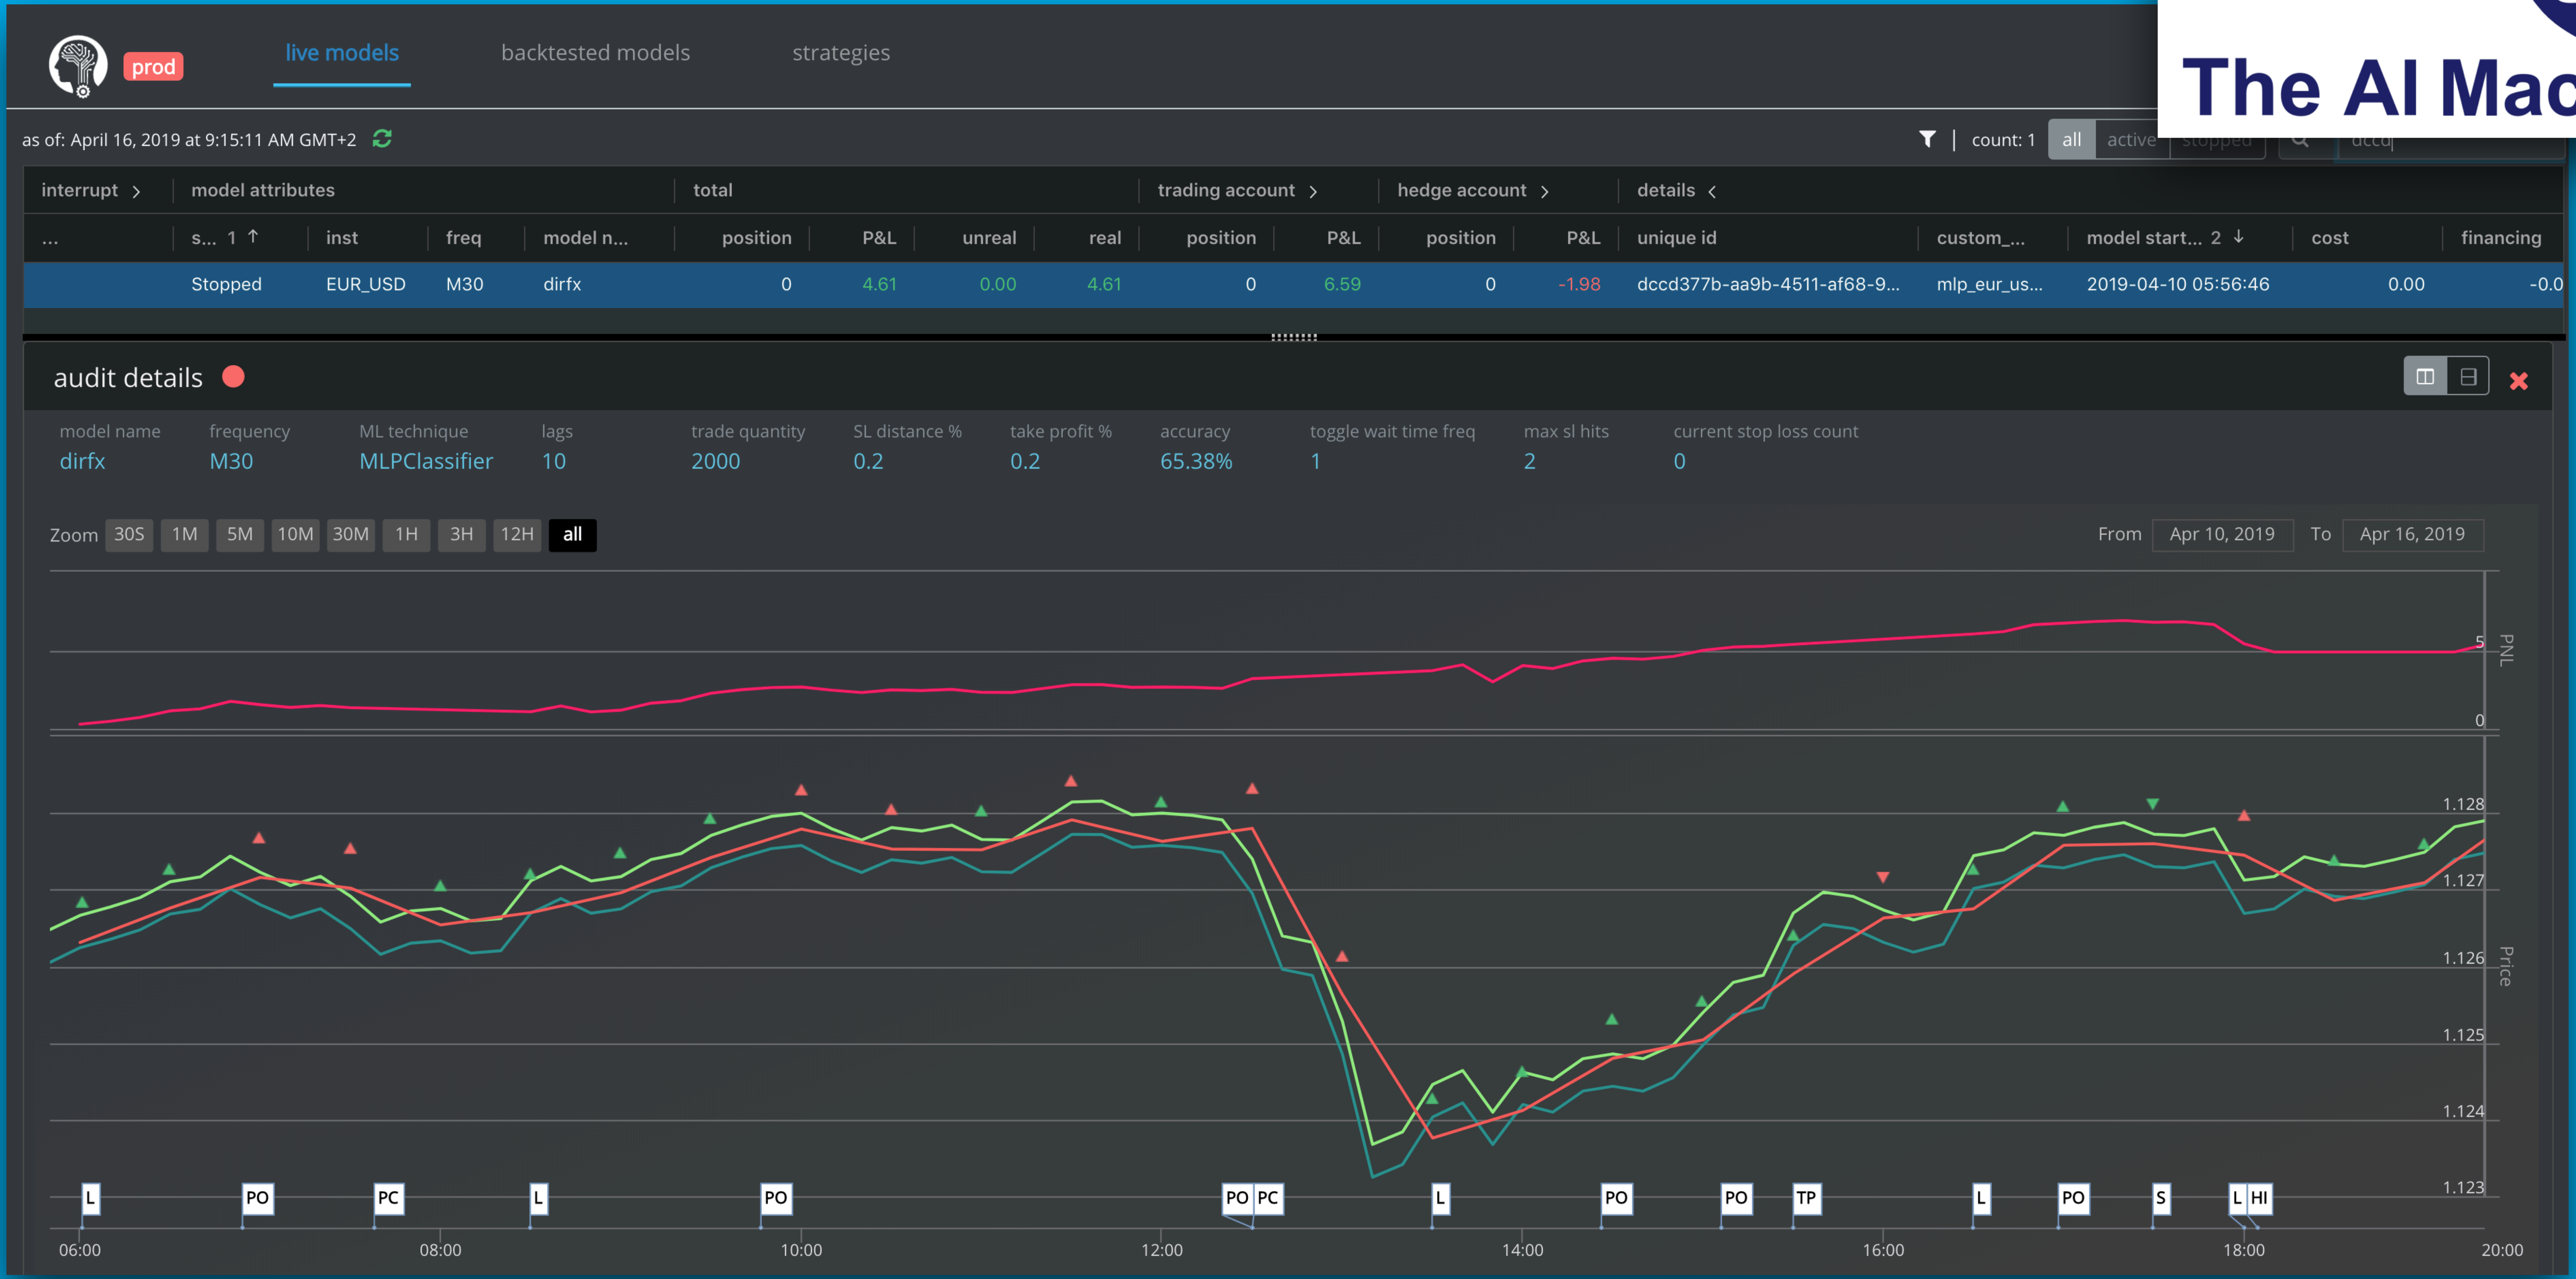

The AI Machine

http://aimachine.io

THE PYTHON QUANTS

recognized by Capital Markets CIO Outlook *magazine as*

TOP 10
ALGO TRADING
SOLUTION PROVIDERS - 2019

An annual listing of 10 companies that are at the forefront of providing Algo Trading solutions

http://certificate.tpq.io/tpq_top_algo_2019.pdf

The Python Quants

First University Certificate in Python for Algorithmic Trading

Python programming has become a key skill in the financial industry. In areas such as financial data science, computational finance or algorithmic trading, Python has established itself as the primary technological platform. At the same time, the level of Python sophistication the industry is expecting from its employees and applicants is increasing steadily. The Python Quants Group is one of the leading providers of Python for Finance training programs.

Among others, The Python Quants have tailored a comprehensive online training program leading to the first University Certificate in Python for Algorithmic Trading. Be it an ambitious student with intrigue for algorithmic trading, or a major financial institution, The Python Quants, through this systematic training program, is equipping delegates with requisite skills and tools to formulate, backtest and deploy algorithmic trading strategies based on Python.

The topics covered in the training programs offered by The Python Quants are generally not found in the typical curriculum of financial engineering or quantitative finance Master programs. Dr. Yves Hilpisch, the firm's founder and managing partner, explains, "There are courses out there that show students how to apply machine learning for the formulation and backtesting of algorithmic trading strategies. However, none of them explains the difficulties or the skills required in deploying such algorithmic trading strategies in the real world. Besides providing an introductory course that teaches Python and financial concepts from scratch, we train our delegates and clients on how best to deploy algorithmic trading strategies in automated fashion in the cloud, with, among others, real-time risk management and monitoring," explains Hilpisch, an author of three books on the topic, with "Python for Finance" (2nd ed., O'Reilly) being the standard reference in the field.

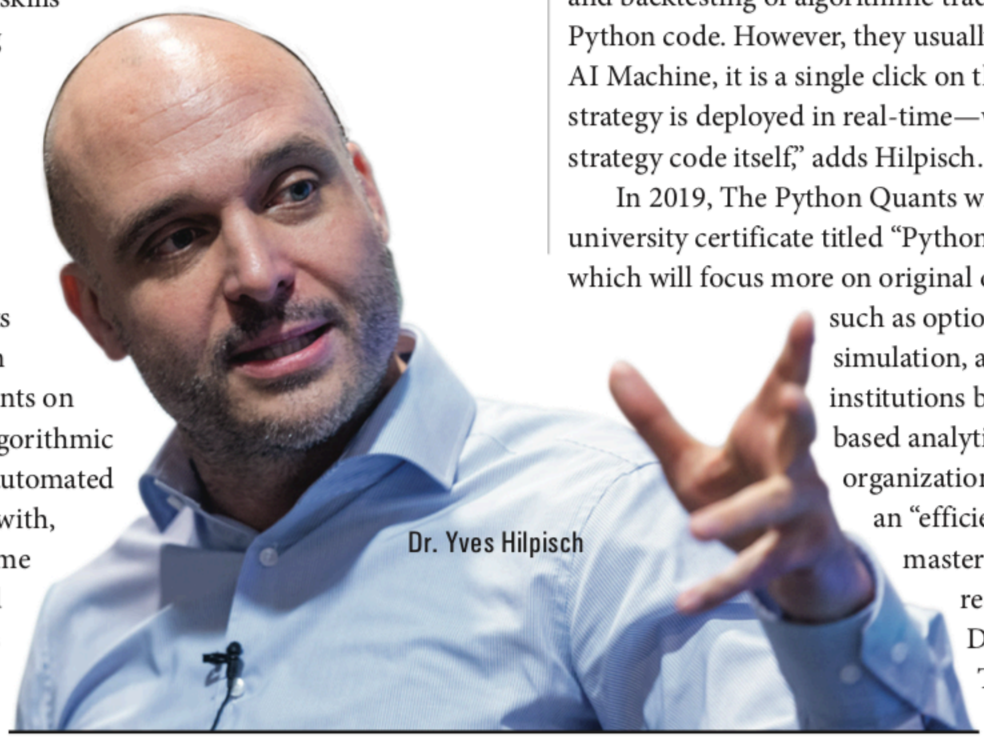

Dr. Yves Hilpisch

The organization's "Python for Algorithmic Trading University Certificate" consists of 200 hours of instruction, 1,200 pages of documentation and 1,000s of lines of Python code. In addition to offering both online and offline Python training, Hilpisch and his team also organize bespoke training events for financial institutions, hedge funds, banks, and asset management companies. "Most of the training is online since we have students and delegates from about 65 different countries in general. Most recently, we noticed that it's not just financial firms and students who want to deepen their algorithmic trading knowledge, but even professors of finance who want to get more involved in this popular topic," says Hilpisch.

While the Quant Platform is the most popular choice, especially for users in the financial sector who don't have access to a full-fledged, interactive, financial analytics environment, the team at The Python Quants is currently developing The AI Machine—a new platform which leverages artificial intelligence to formulate and deploy algorithmic trading strategies in a standardized manner. Hilpisch explains that it's relatively easy to write Python code for an algorithmic trading strategy, but the same can't be said about the deployment of such a strategy. "There are a few platforms out there that allow the formulation and backtesting of algorithmic trading strategies by the use of Python code. However, they usually stop exactly there. With The AI Machine, it is a single click on the 'GO LIVE' button and the strategy is deployed in real-time—without any changes to the strategy code itself," adds Hilpisch.

In 2019, The Python Quants will be introducing a new university certificate titled "Python for Computational Finance," which will focus more on original quantitative finance topics, such as option pricing, Monte Carlo simulation, and hedging. As financial institutions begin to perceive Python-based analytics as a prerequisite skill, the organization will continue to provide an "efficient and structured way of mastering all the tools and skills required in Python for Financial Data Science, Algorithmic Trading, and Computational Finance." CM

About Myself

http://hilpisch.com

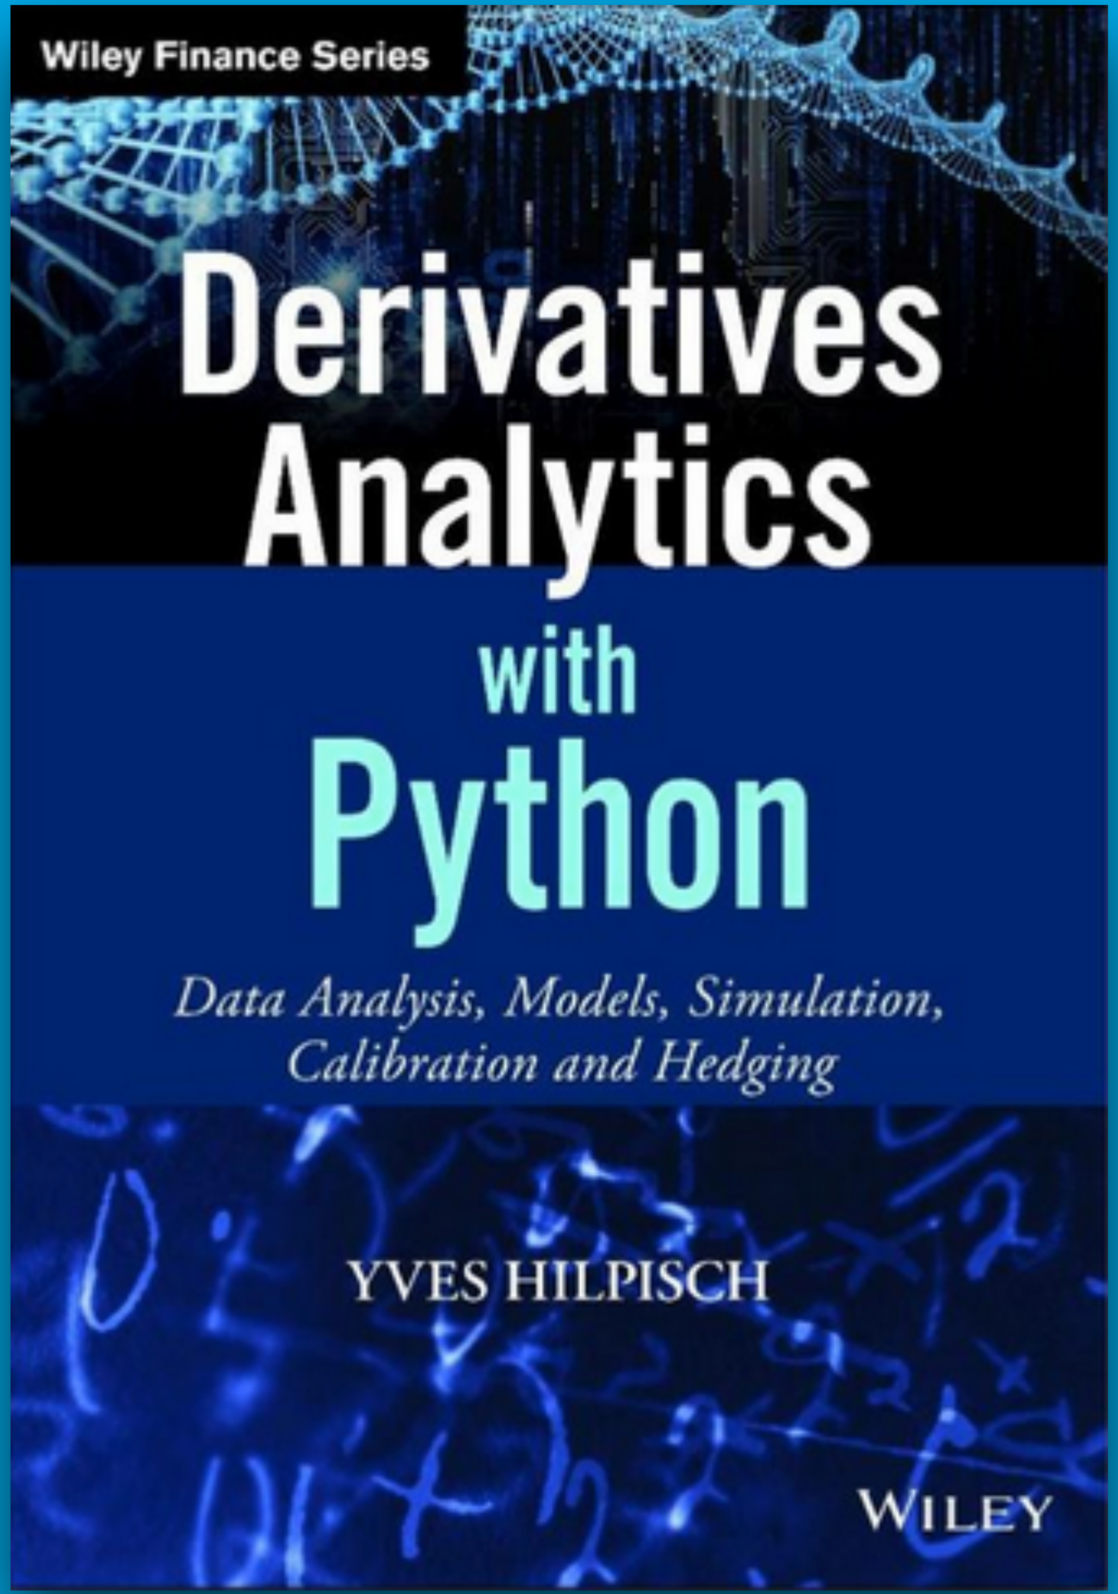

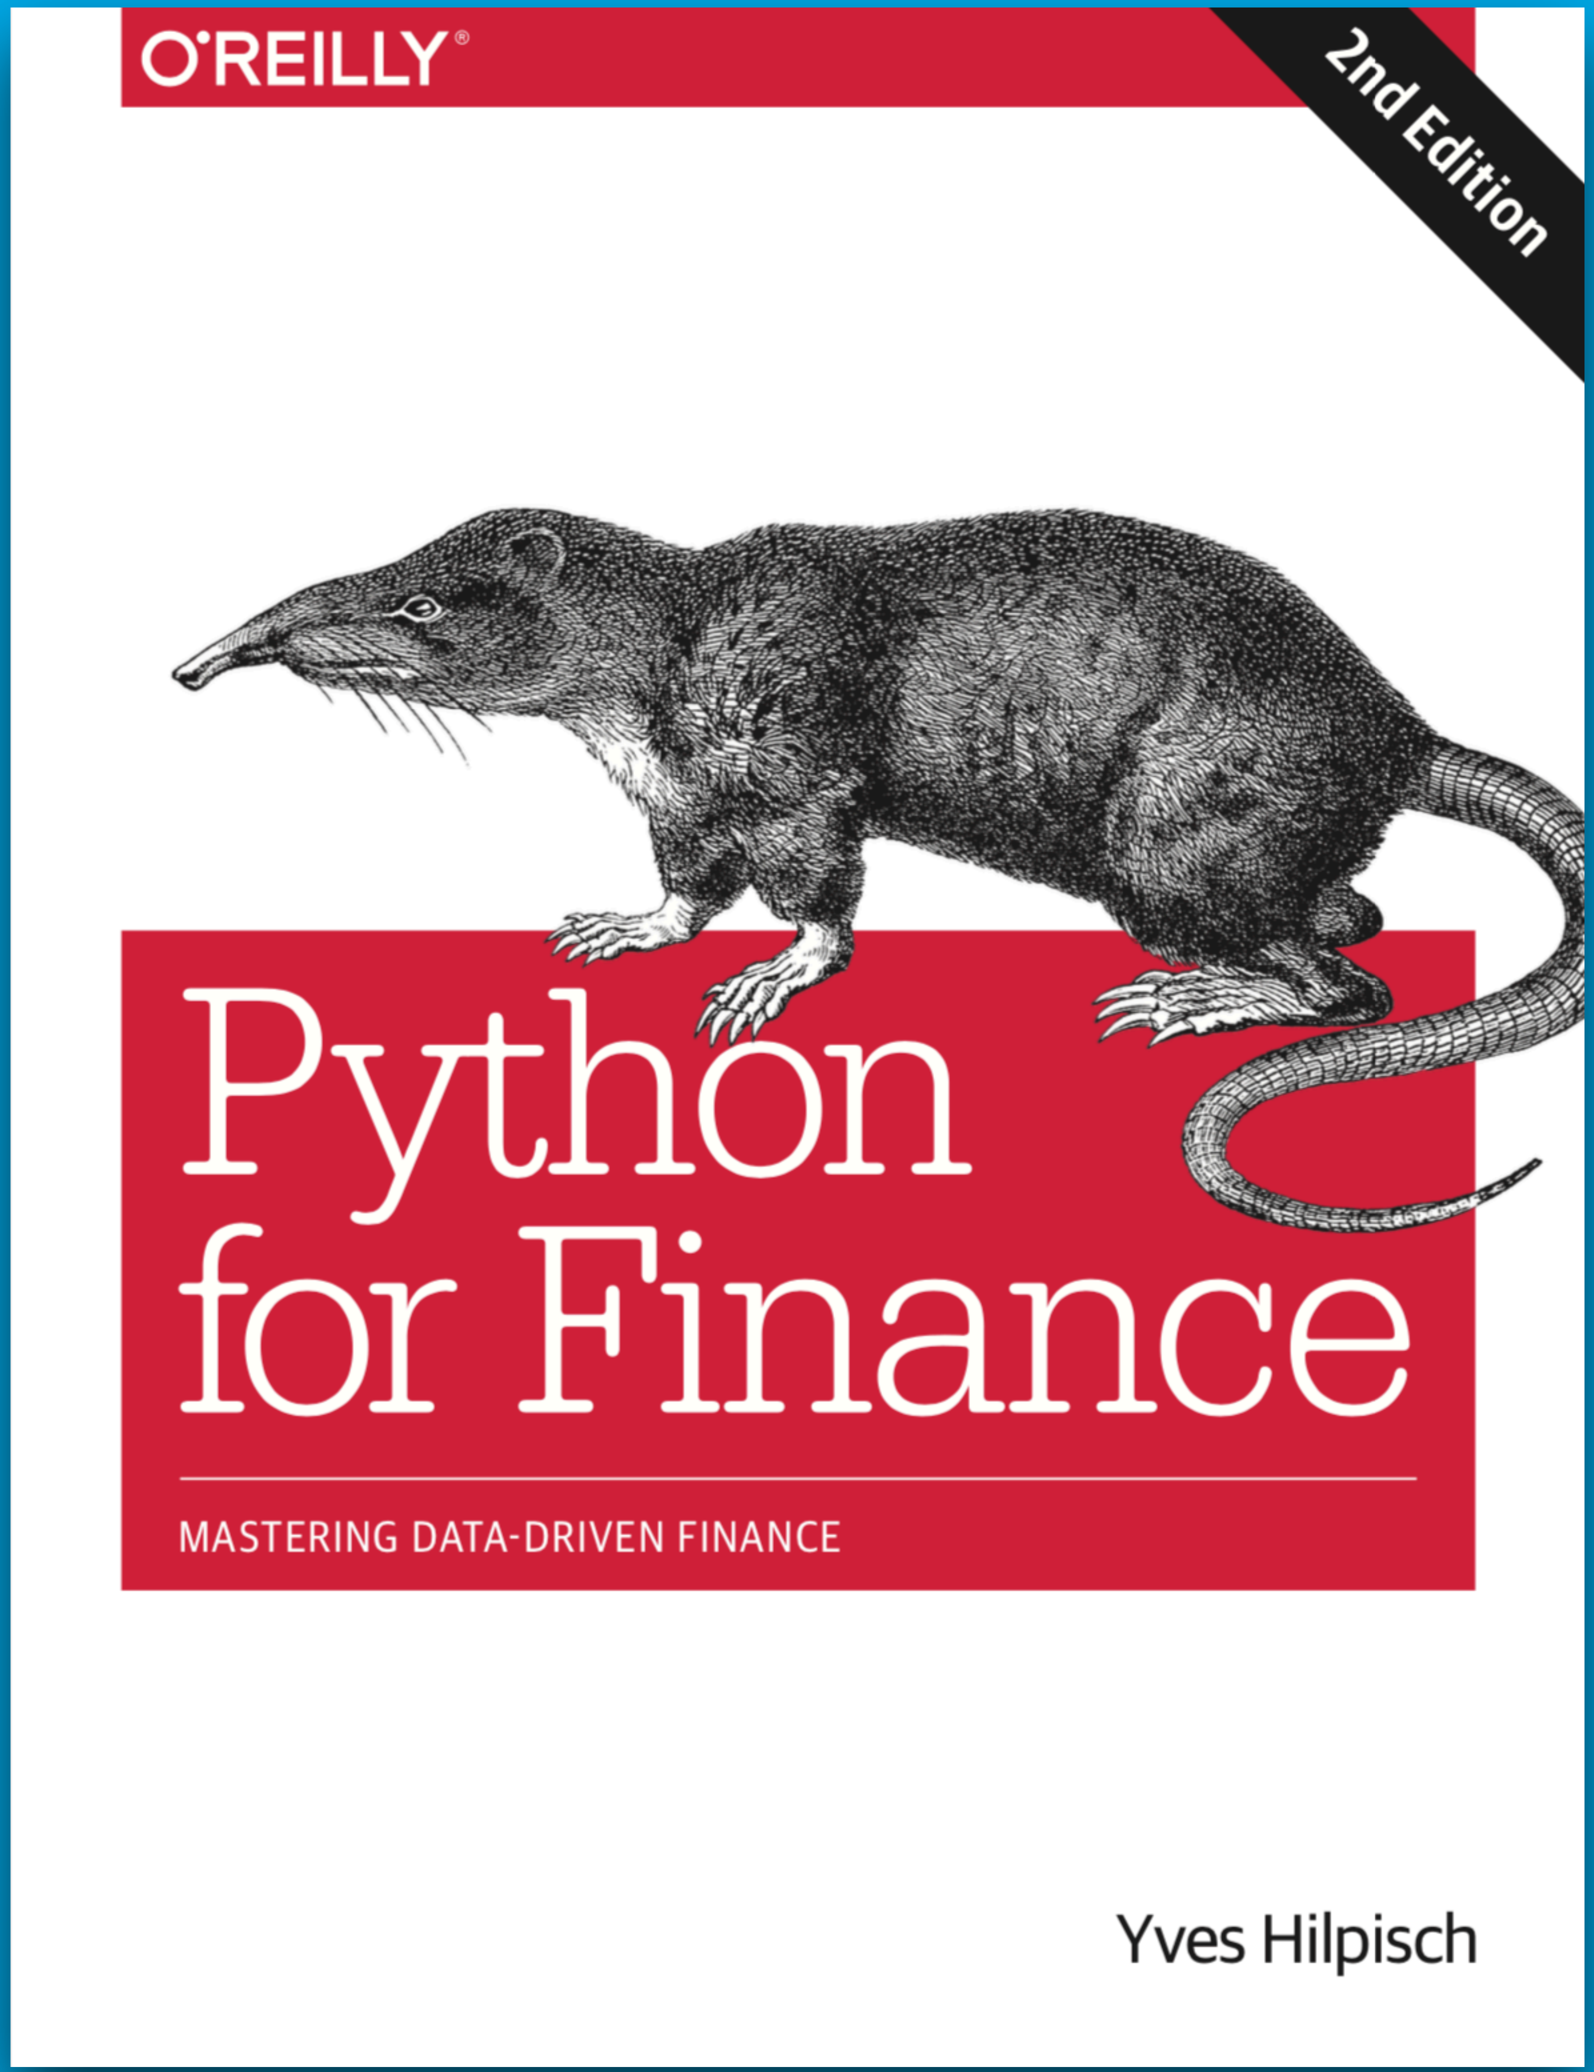

NEW book project:
Artificial Intelligence in Finance — A Python-based Guide

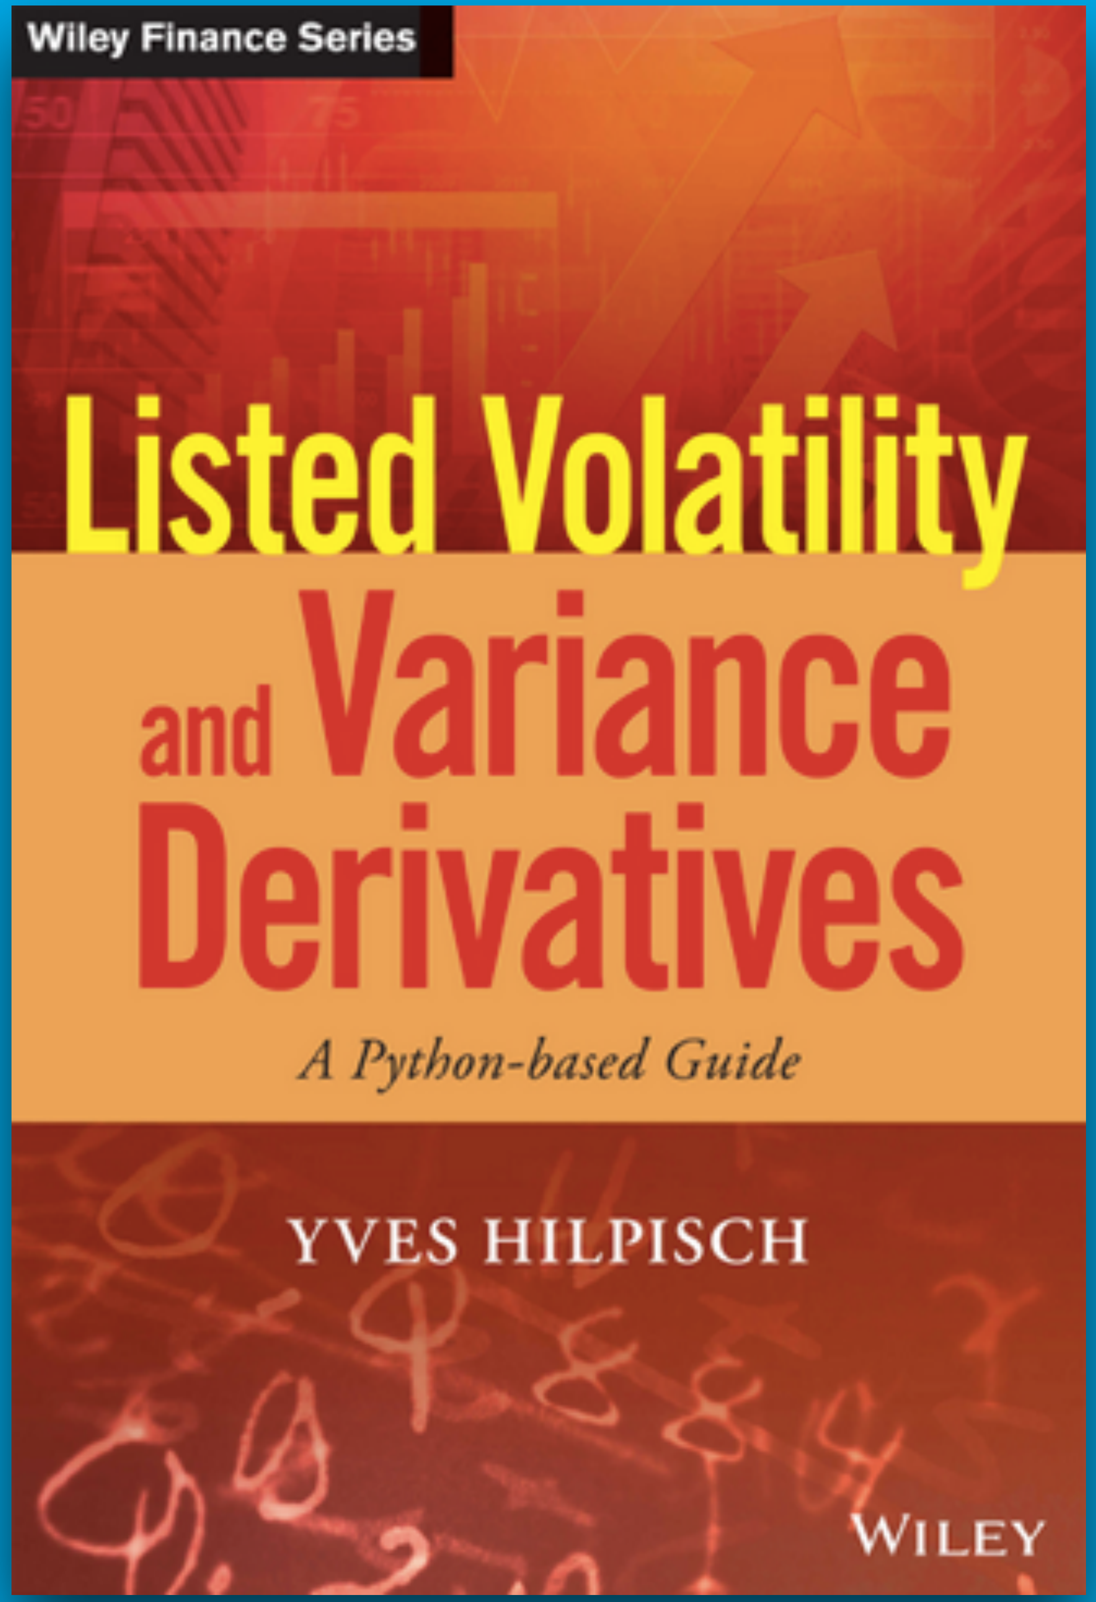

http://books.tpq.io

Resources (Gist):

http://bit.ly/odsc_east

Overview

1. The Beauty Myth
2. Data-Driven Finance
3. Statistical Learning
4. OLS Regression
5. Efficient Markets
6. AI-First Finance
7. Algorithms
8. Deep Learning
9. Market Prediction
10. Conclusions

The Beauty Myth

Particle physics

Fundamental physics is frustrating physicists

The Economist

No GUTs, no glory

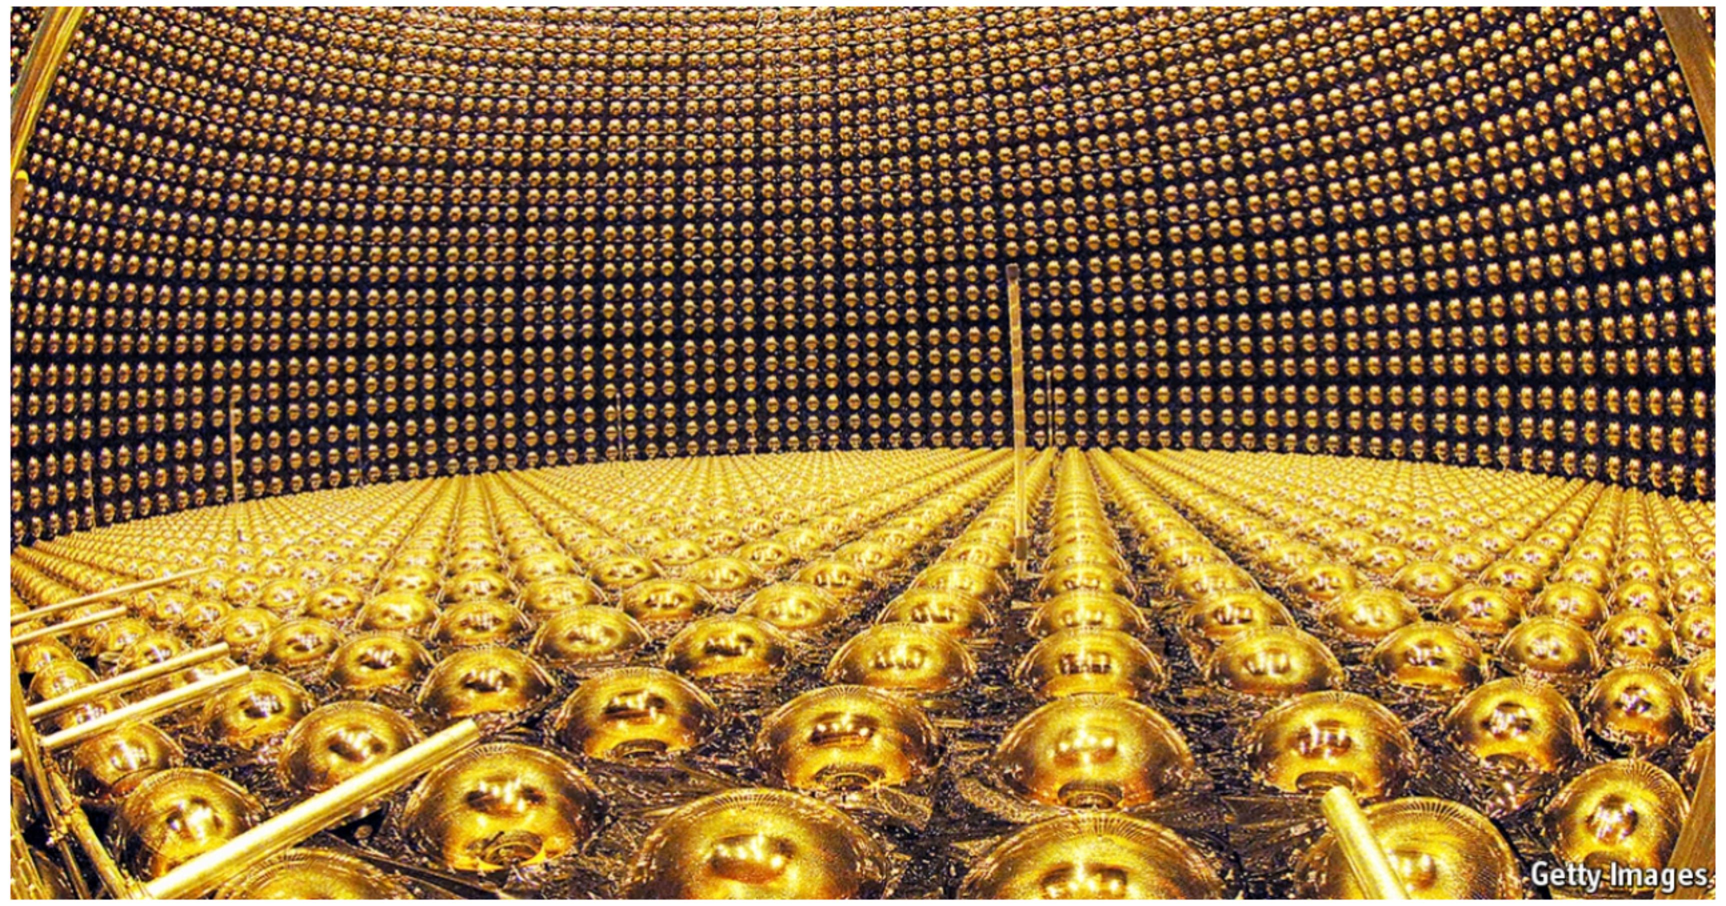

Print edition | Science and technology ›

Jan 13th 2018

DEEP in a disused zinc mine in Japan, 50,000 tonnes of purified water held in a vast cylindrical stainless-steel tank are quietly killing theories long cherished by physicists. Since 1996, the photomultiplier-tube detectors (pictured above) at Super-Kamiokande, an experiment under way a

The beauty myth

One such is Sabine Hossenfelder of the Frankfurt Institute for Advanced Studies, in Germany. She argues that the appeal of GUTs, supersymmetry and the like rests on their ability to explain "numerological coincidences" that do not need to be explained. Perhaps, to take one example, the universe simply started out with more matter than antimatter in it, rather than this being a consequence of its subsequent evolution. As she points out, no theory precludes this possibility—it is just that it is not very elegant. Similarly, she says, "It's not like anybody actually needs supersymmetry to explain anything. It's an idea widely praised for its aesthetic appeal. Well, that's nice, but it's not science."

Particle physics

Fundamental physics is frustrating physicists

GUTs are among several long-established theories that remain stubbornly unsupported by the big, costly experiments testing them. …
Despite the dearth of data, the answers that all these theories offer to some of the most vexing questions in physics are so elegant that they populate postgraduate textbooks. As Peter Woit of Columbia University observes, “Over time, these ideas became institutionalised. People stopped thinking of them as speculative.” That is understandable, for they appear to have great explanatory power.

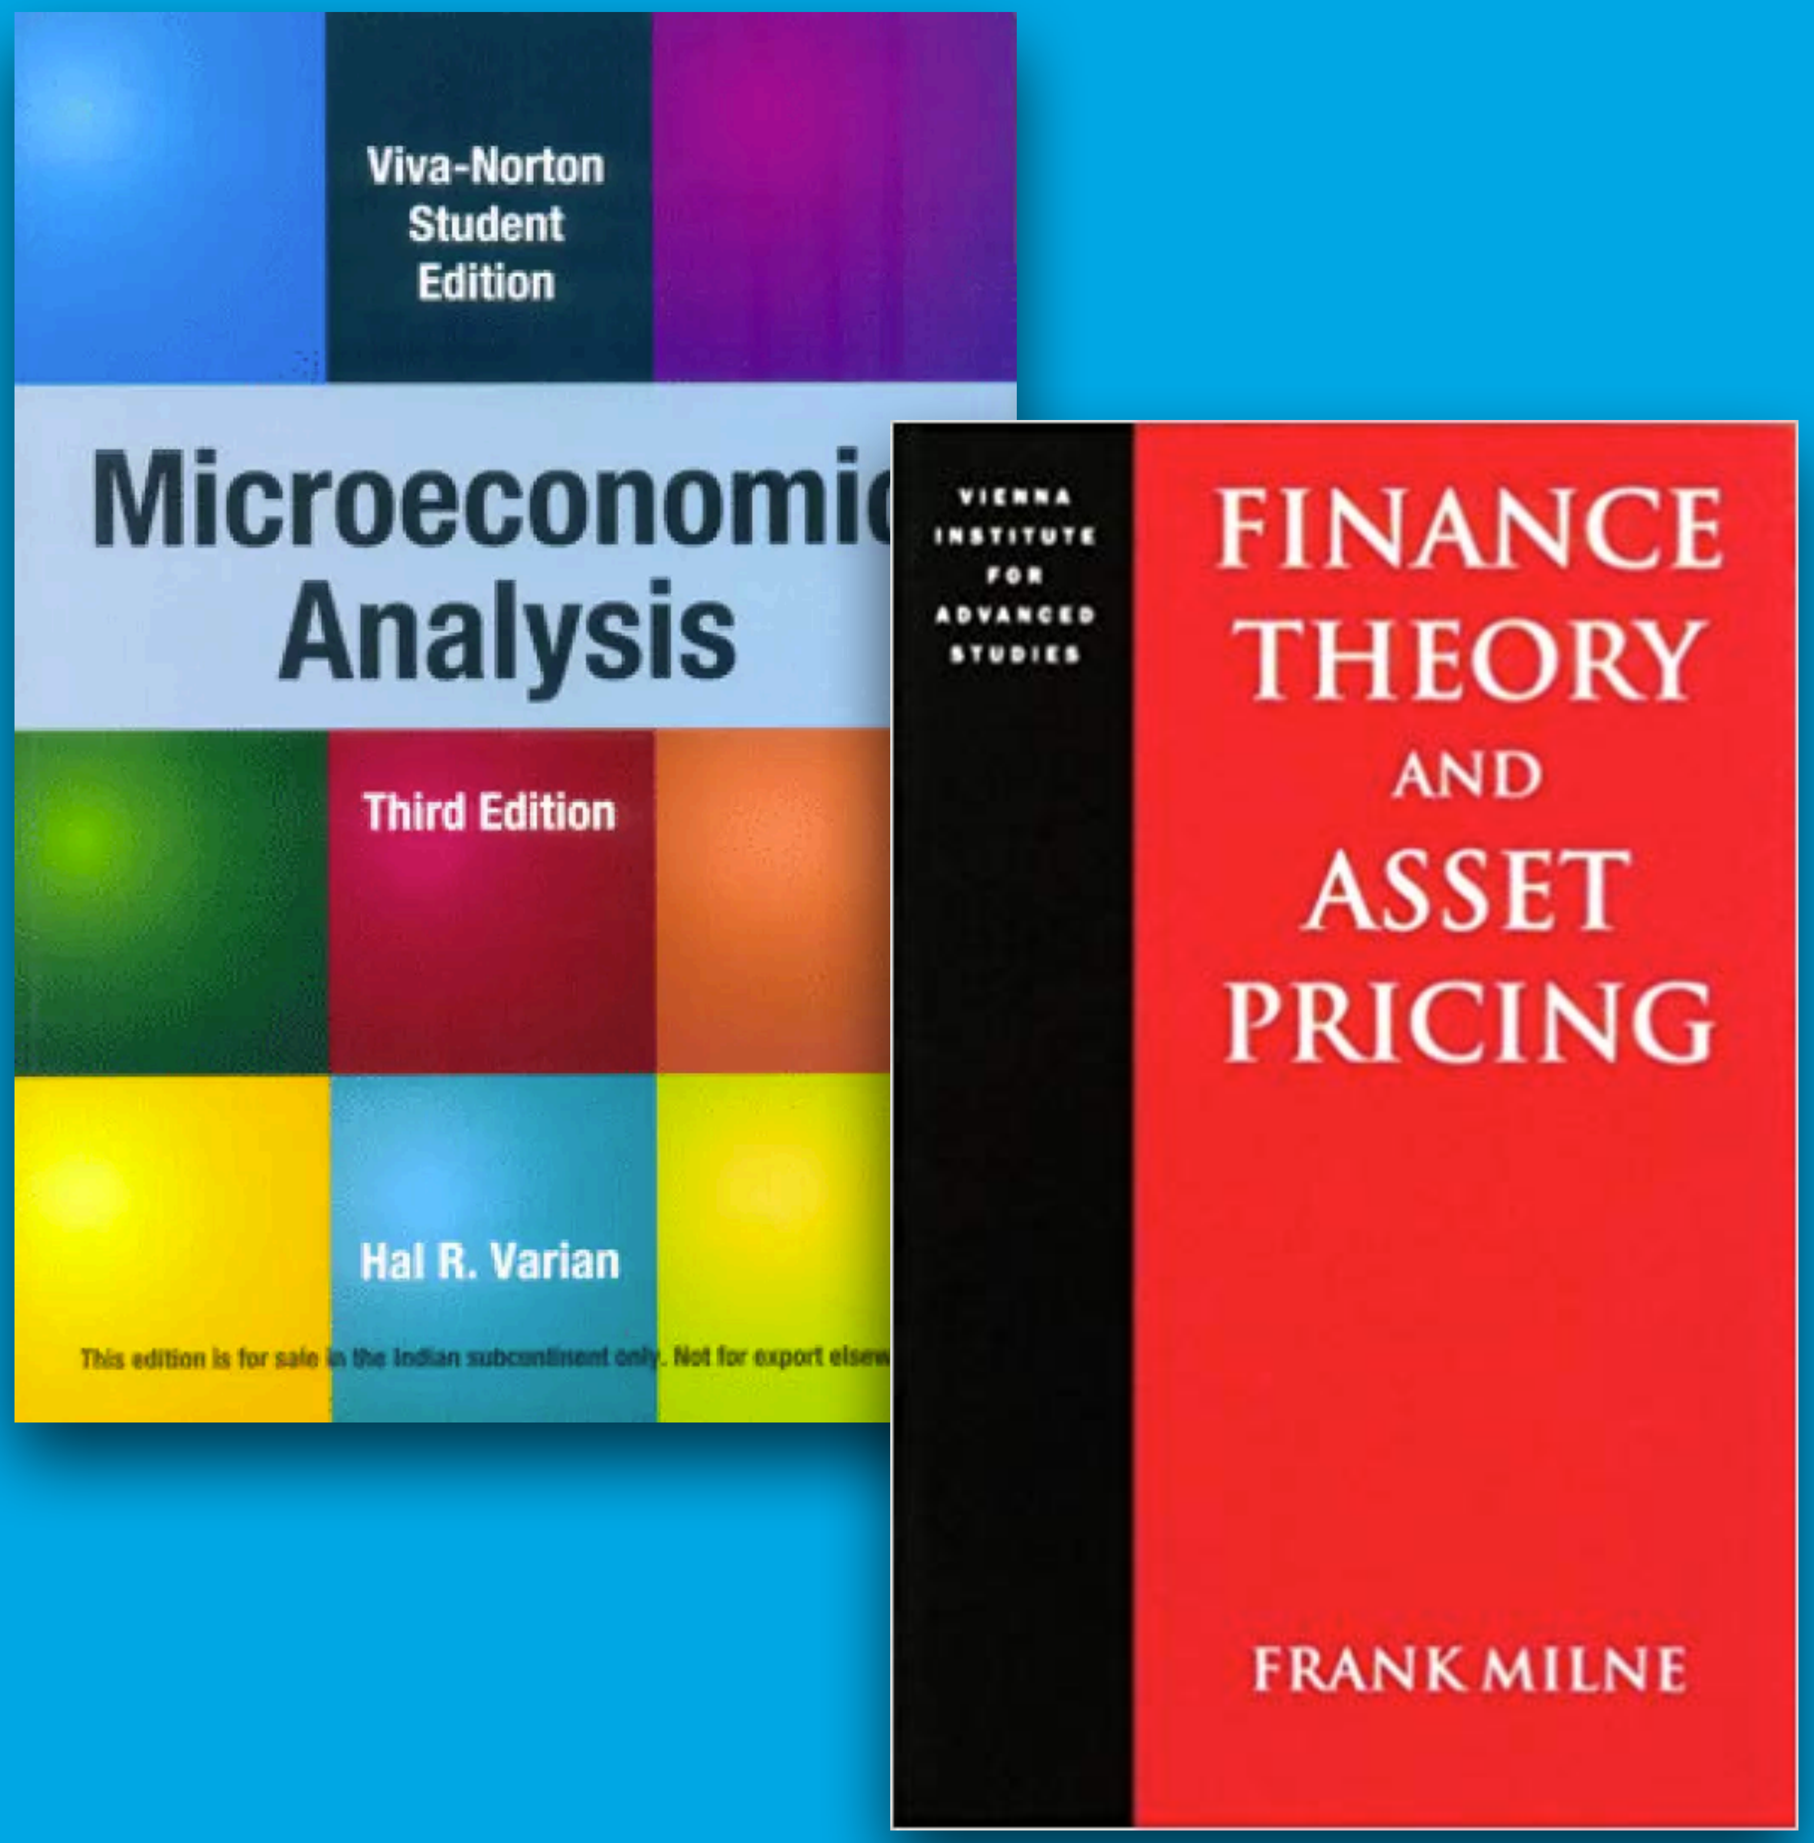

Cornerstones of Economics

A. Arbitrage Pricing
B. Expected Utility
C. Equilibrium
D. Normal Distributions
E. Linear Relationships
F. Efficient Markets

Theory

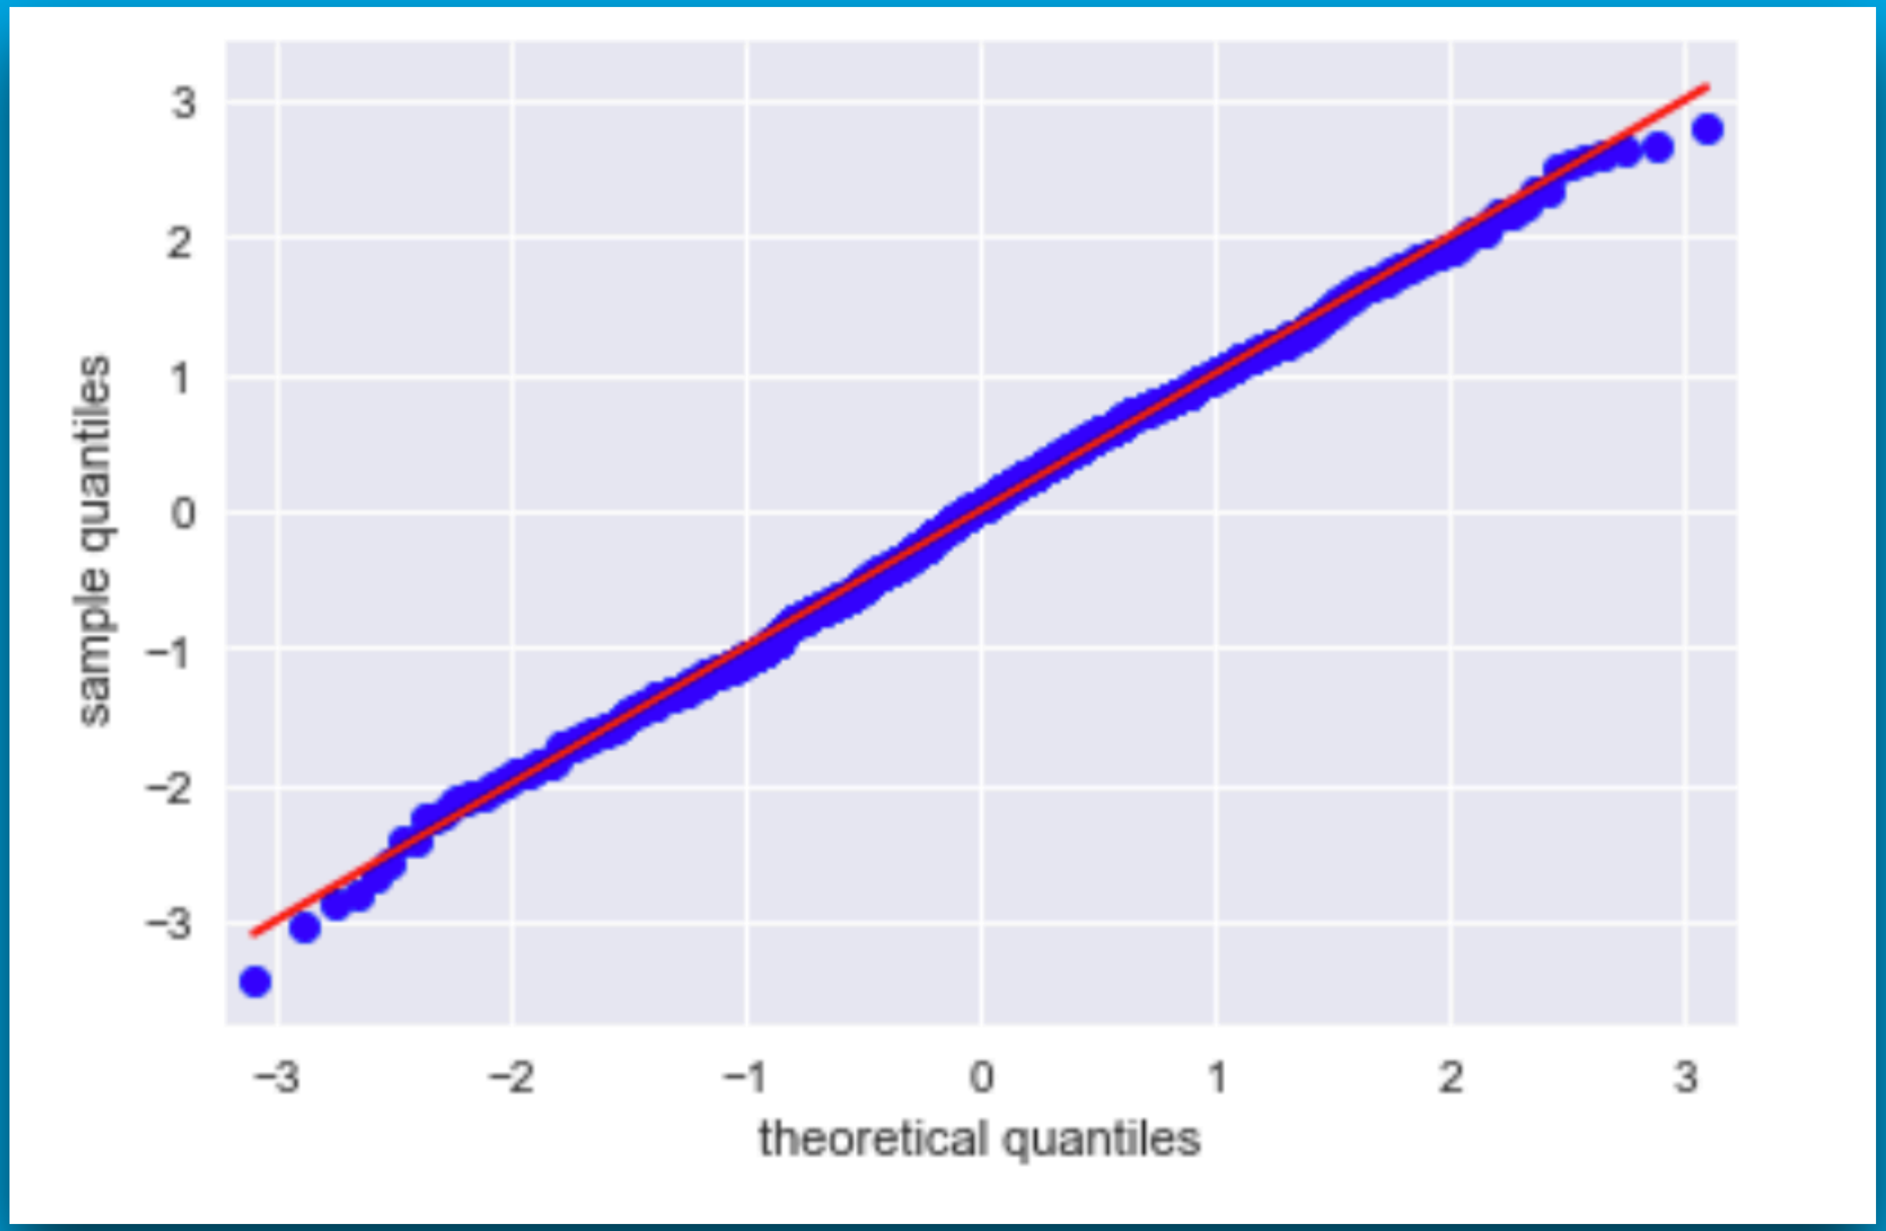

Reality

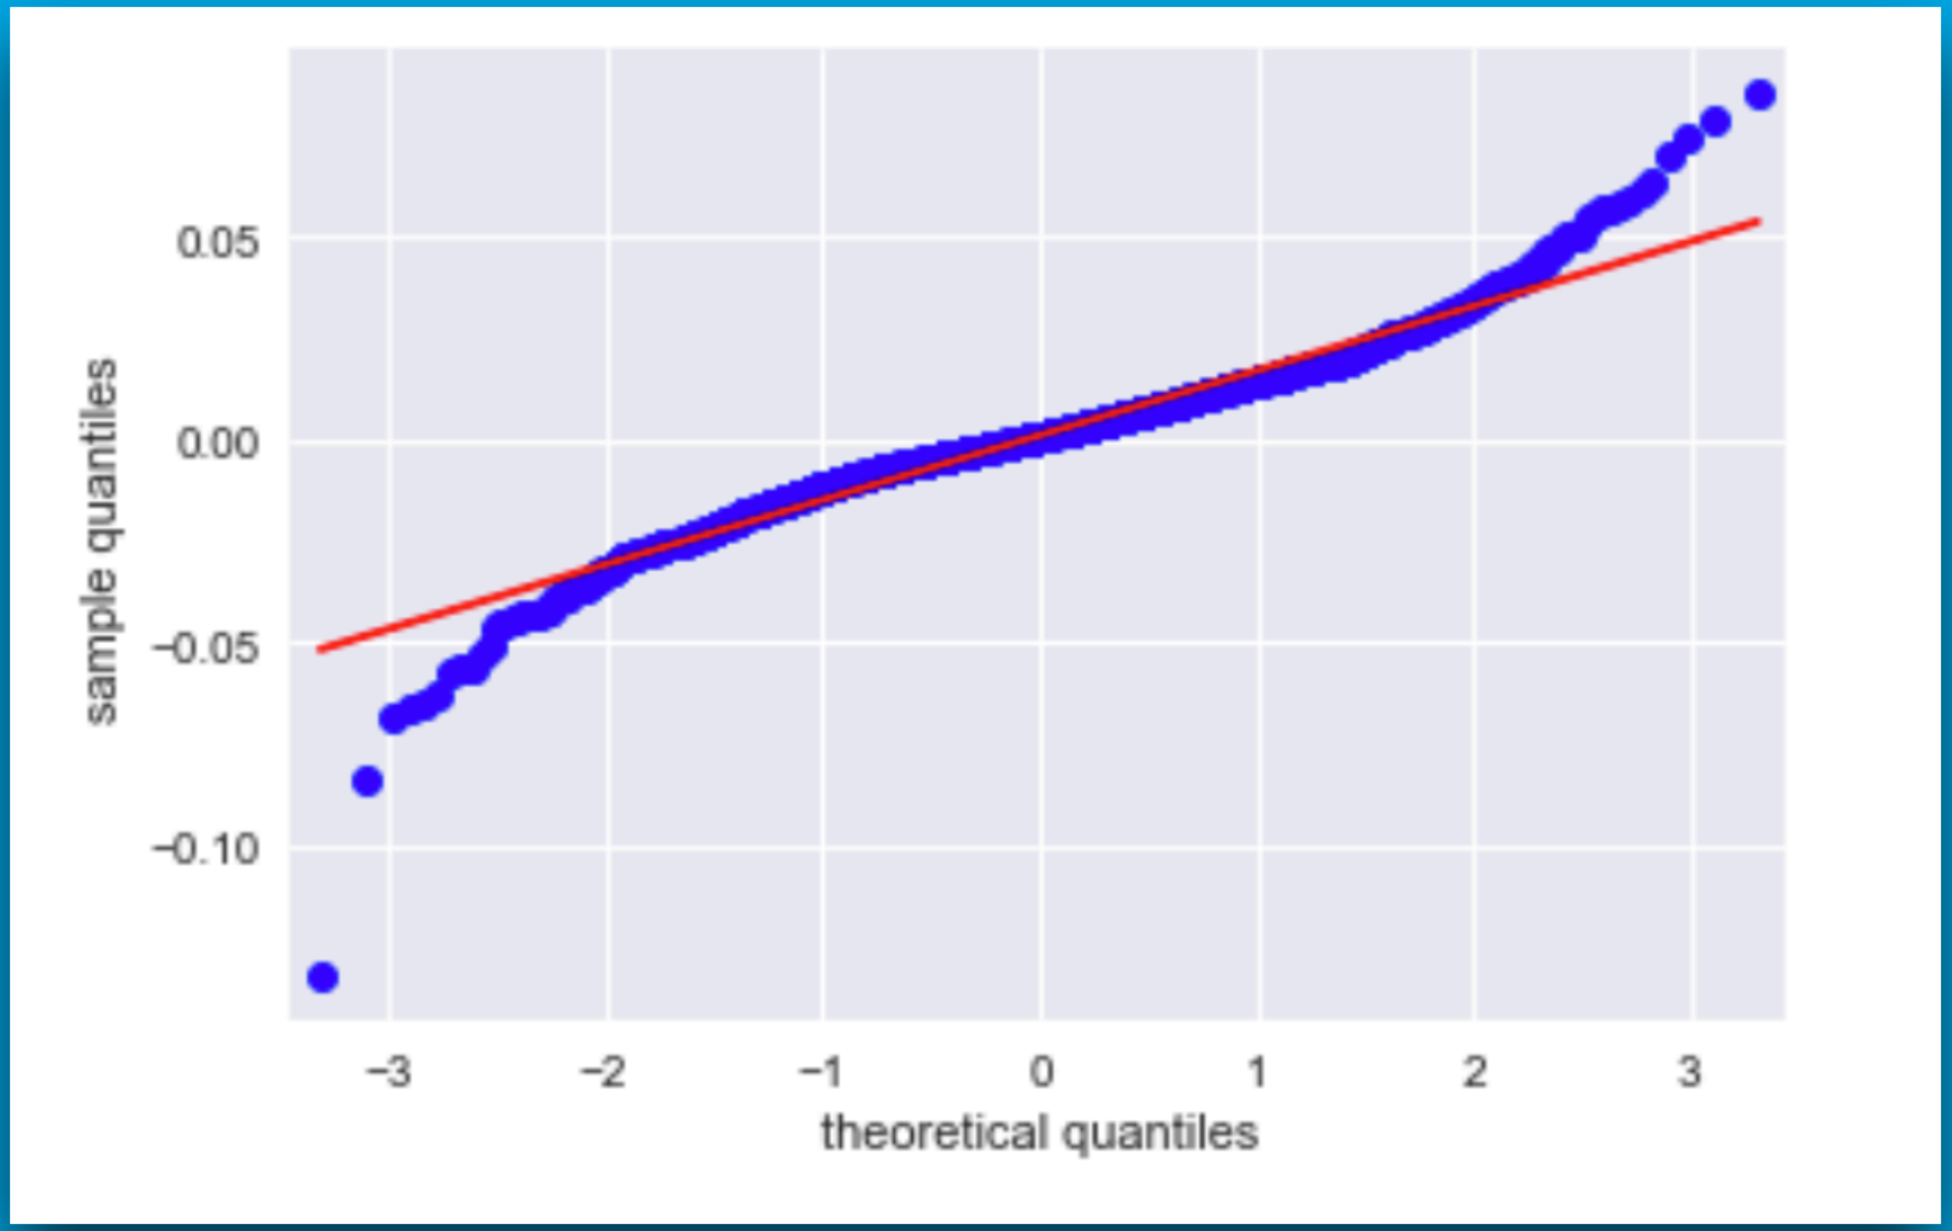

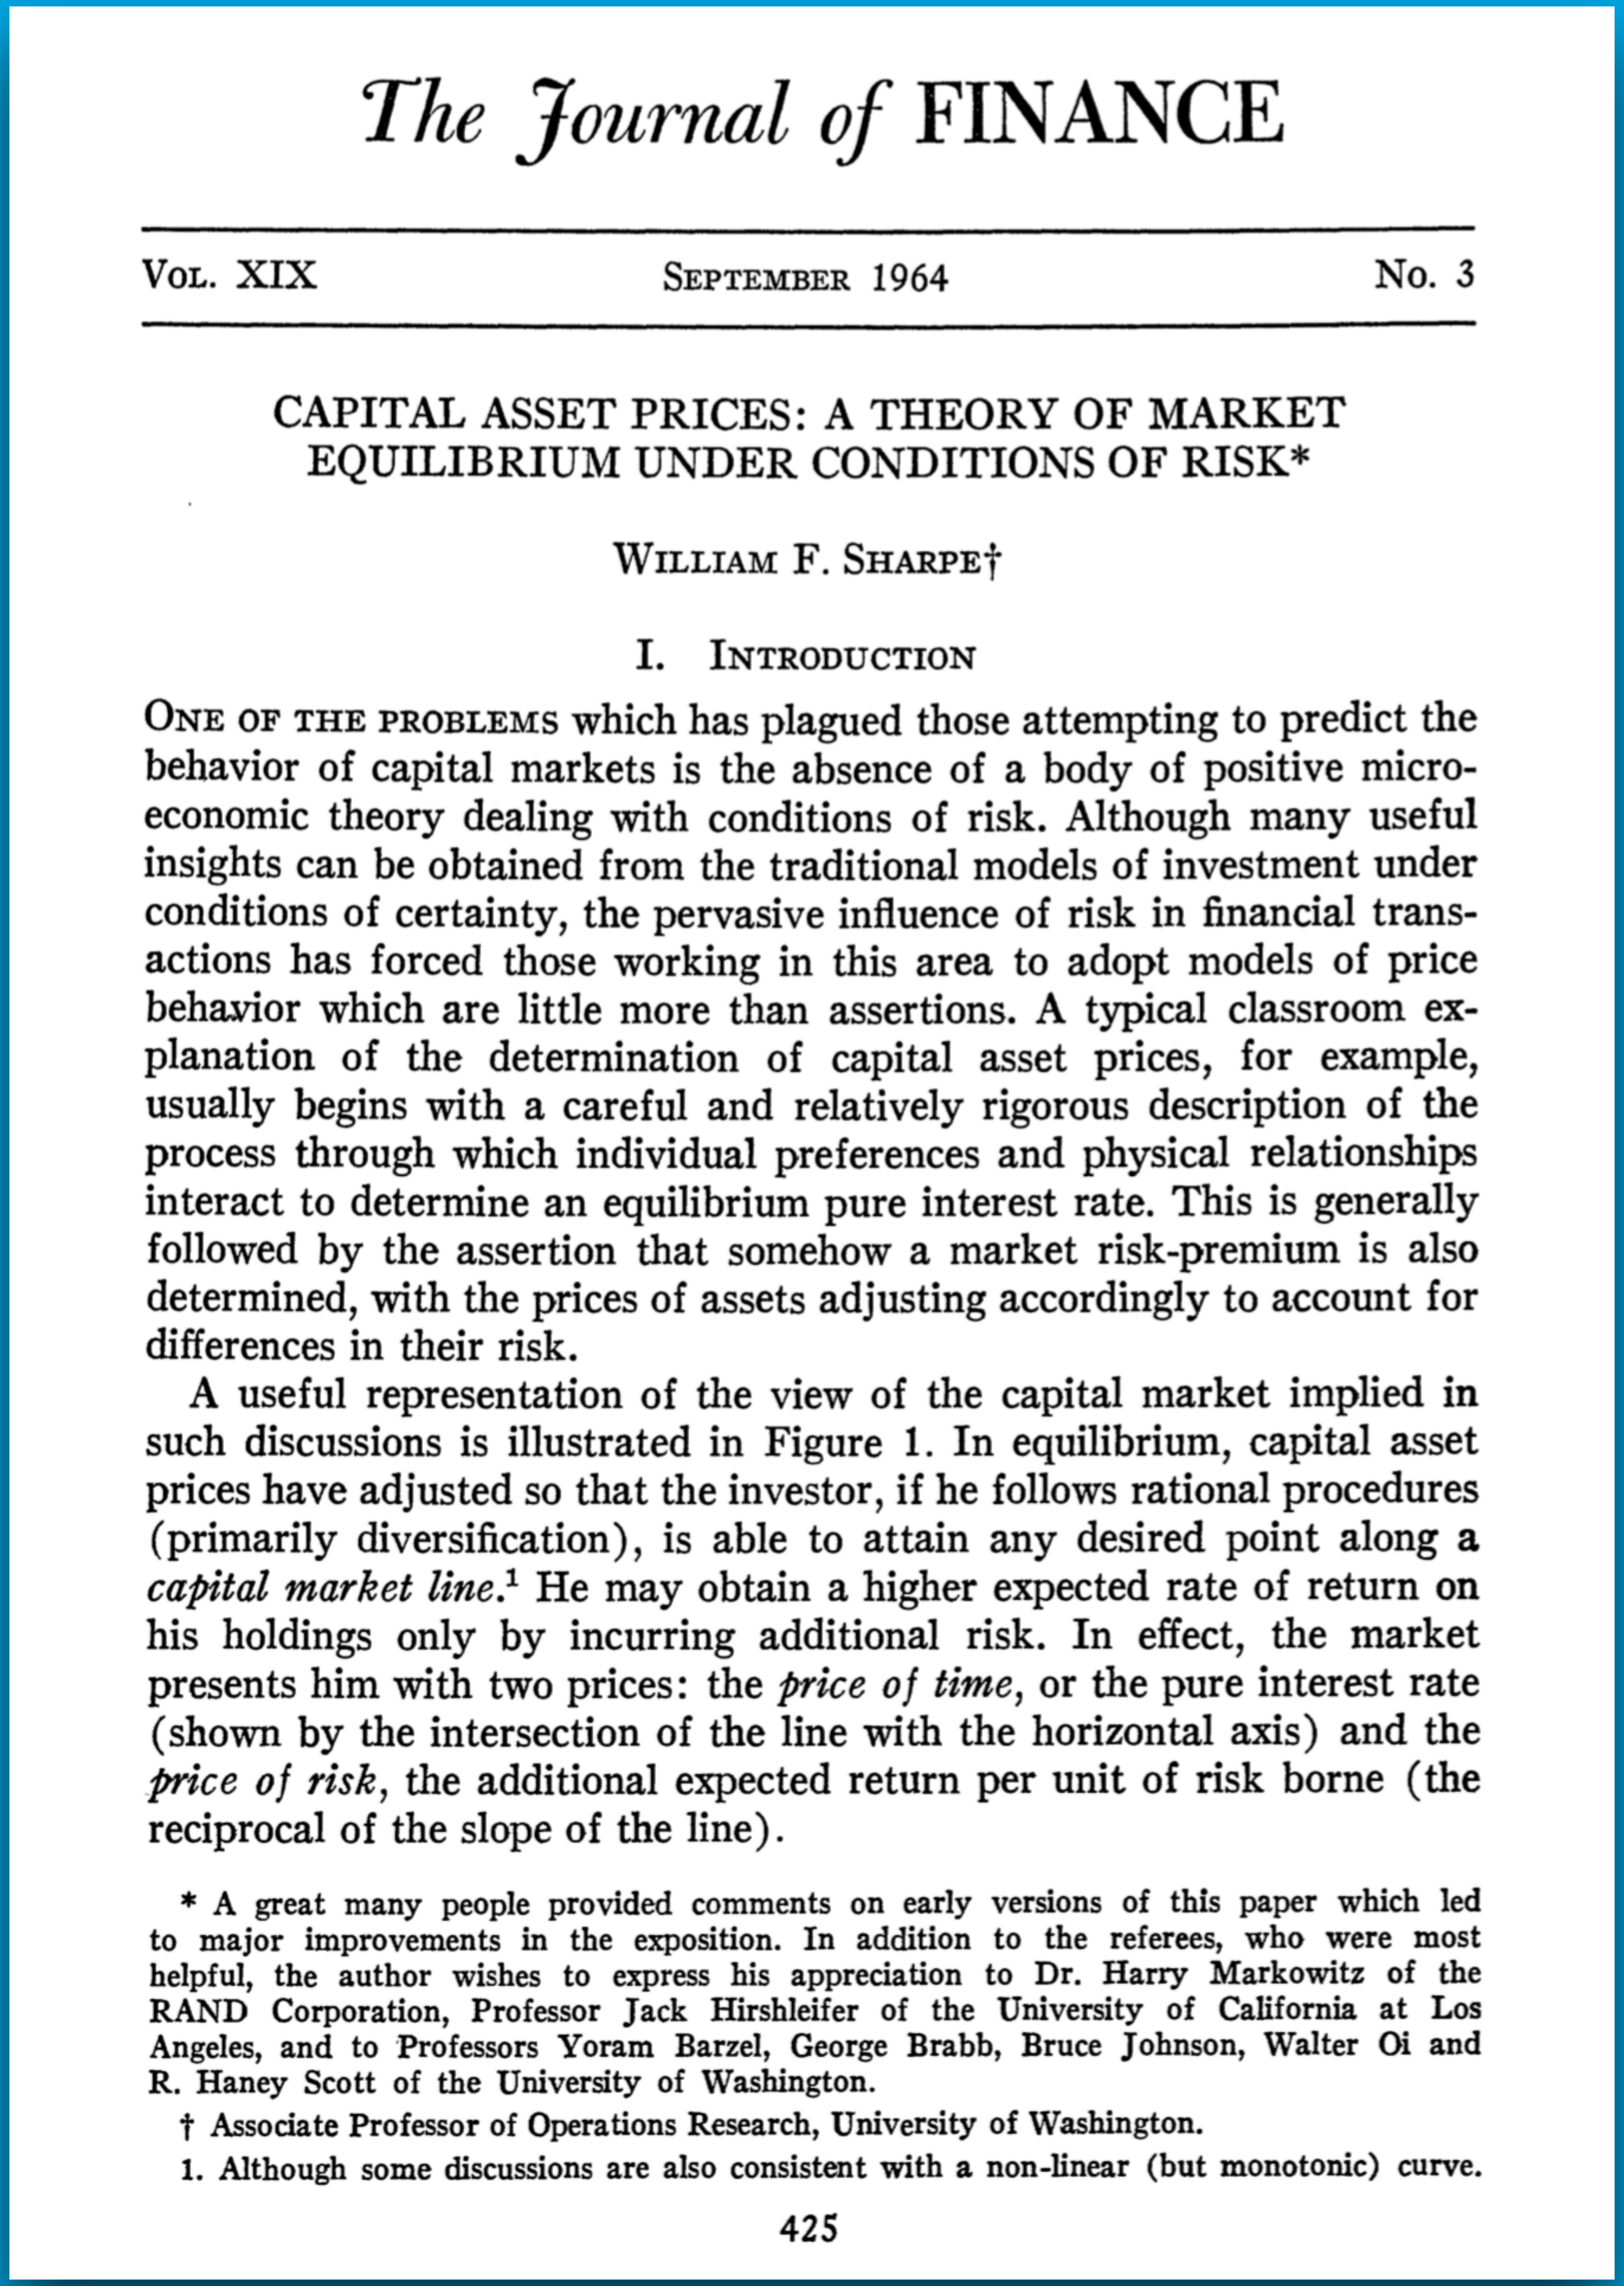

The Journal of FINANCE

Vol. XIX — September 1964 — No. 3

CAPITAL ASSET PRICES: A THEORY OF MARKET EQUILIBRIUM UNDER CONDITIONS OF RISK*

William F. Sharpe†

I. Introduction

One of the problems which has plagued those attempting to predict the behavior of capital markets is the absence of a body of positive microeconomic theory dealing with conditions of risk. Although many useful insights can be obtained from the traditional models of investment under conditions of certainty, the pervasive influence of risk in financial transactions has forced those working in this area to adopt models of price behavior which are little more than assertions. A typical classroom explanation of the determination of capital asset prices, for example, usually begins with a careful and relatively rigorous description of the process through which individual preferences and physical relationships interact to determine an equilibrium pure interest rate. This is generally followed by the assertion that somehow a market risk-premium is also determined, with the prices of assets adjusting accordingly to account for differences in their risk.

A useful representation of the view of the capital market implied in such discussions is illustrated in Figure 1. In equilibrium, capital asset prices have adjusted so that the investor, if he follows rational procedures (primarily diversification), is able to attain any desired point along a *capital market line*.[1] He may obtain a higher expected rate of return on his holdings only by incurring additional risk. In effect, the market presents him with two prices: the *price of time*, or the pure interest rate (shown by the intersection of the line with the horizontal axis) and the *price of risk*, the additional expected return per unit of risk borne (the reciprocal of the slope of the line).

* A great many people provided comments on early versions of this paper which led to major improvements in the exposition. In addition to the referees, who were most helpful, the author wishes to express his appreciation to Dr. Harry Markowitz of the RAND Corporation, Professor Jack Hirshleifer of the University of California at Los Angeles, and to Professors Yoram Barzel, George Brabb, Bruce Johnson, Walter Oi and R. Haney Scott of the University of Washington.

† Associate Professor of Operations Research, University of Washington.

1. Although some discussions are also consistent with a non-linear (but monotonic) curve.

$$\mu_i = r + \beta_i(\mu_M - r)$$

"Market Risk"
"Idiosyncratic Risk"

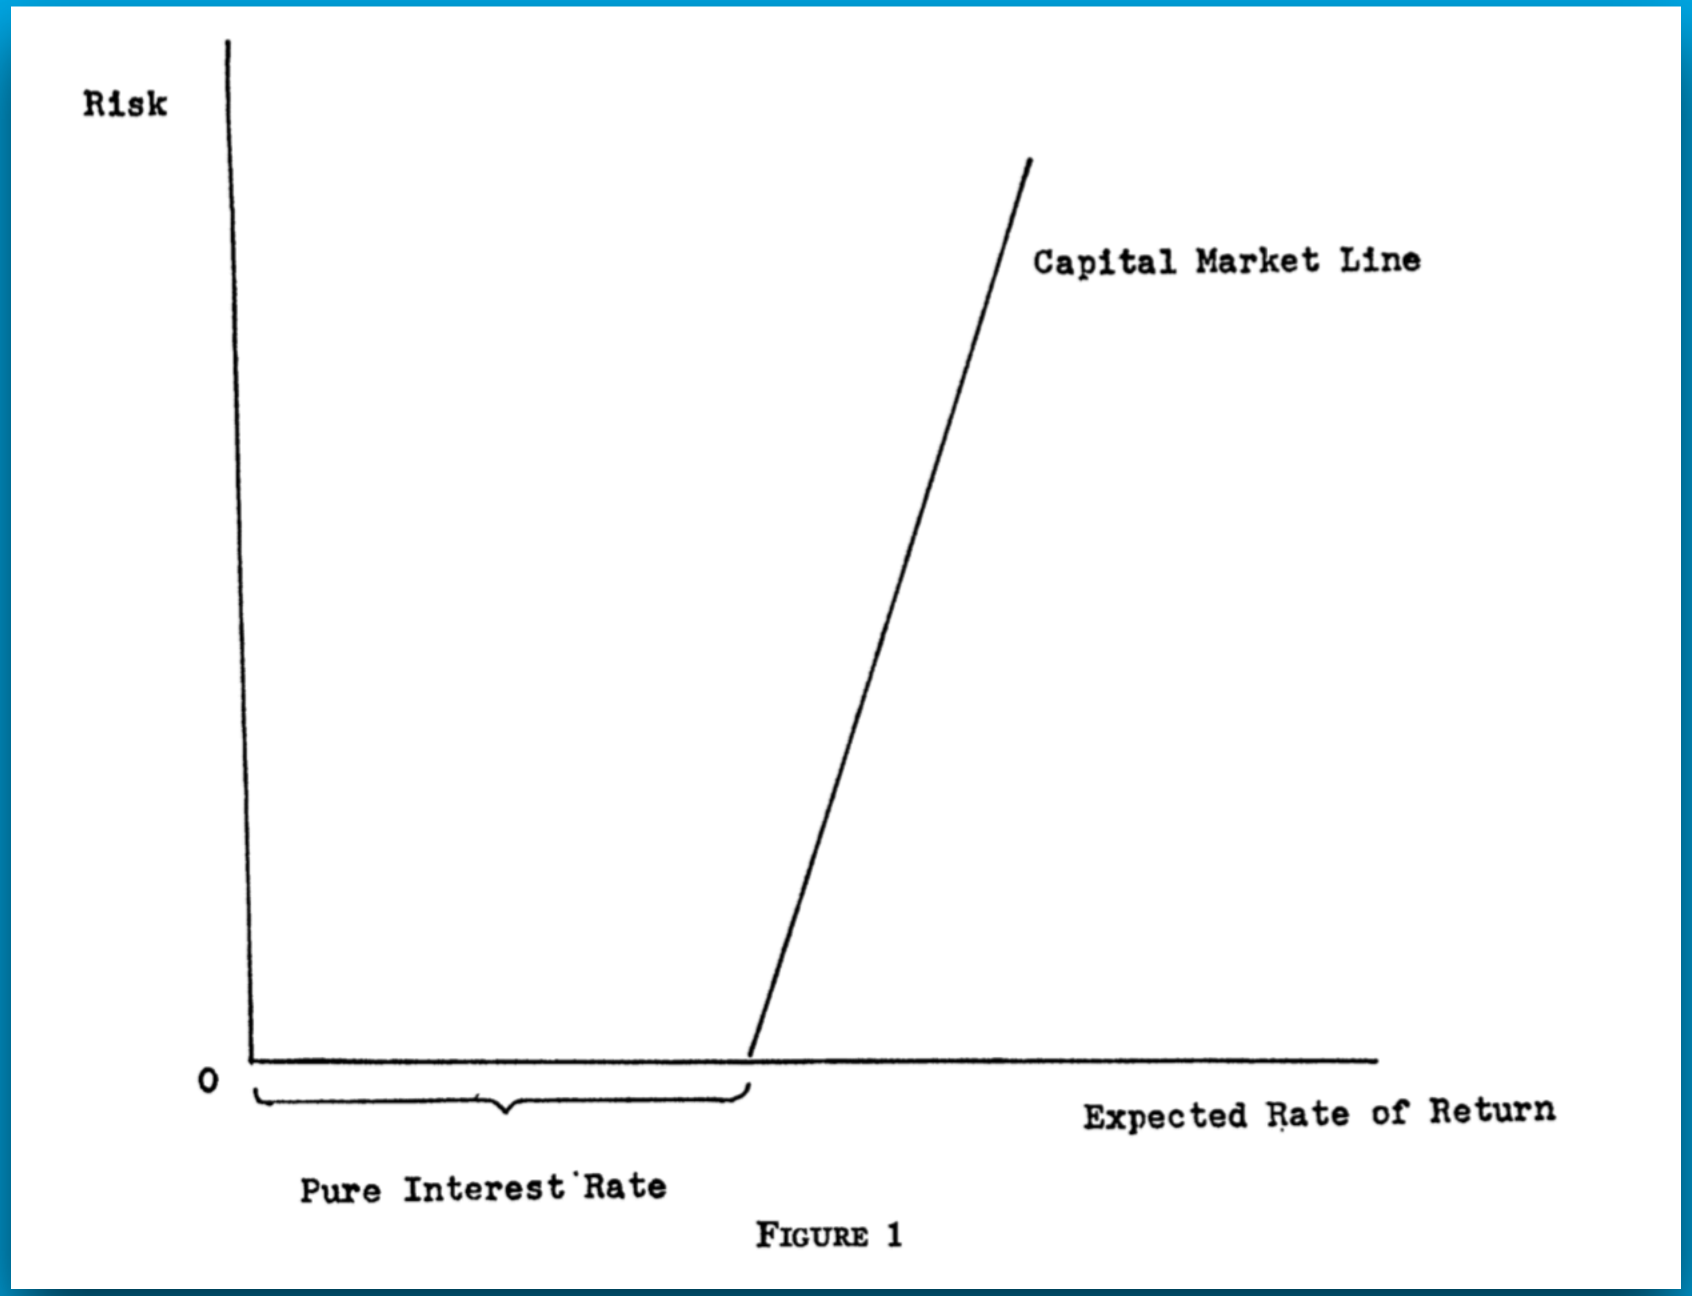

Figure 1

Data-Driven Finance

FINANCIAL TIMES
TUESDAY 3 OCTOBER 2017
WORLD BUSINESS NEWSPAPER
MIDDLE EAST
Mohamed El-Erian
Reasons to worry about the Fed's 'beautiful normalisation' — PAGE 20
Torturing Theresa
Boris Johnson's bid to dictate May's Brexit strategy — JANAN GANESH, PAGE 11
Confined in a circle
The myths that hold back women in Indian society — AMY KAZMIN, PAGE 10
Las Vegas reels from worst US mass shooting
Briefing
Catalan president urges Brussels to mediate in independence clash
Region seeks to avoid 'traumatic split' from Spain
EU says dispute is 'internal matter'
Datawatch
Knowledge. Integrity. Discretion.
A Nikkei Company
Financial Times
Secure https://www.ft.com
FTSE 100 -0.08%
S&P 500 +0.37%
Euro/Dollar -0.17%
Pound/Dollar -0.20%
Brent Crude Oil +0.33%
Visit Markets Data
FINANCIAL TIMES
myFT
HOME WORLD US COMPANIES MARKETS OPINION WORK & CAREERS LIFE & ARTS
Portfolio Account Settings
fastFT
Fast Europe Open: Italy, eurozone industrial production
Tencent-backed carmaker NIO prices shares at bottom of range
Oil prices climb as hurricane threatens US
Hong Kong bear market worsens as Asia stocks fall
US banks
Jamie Dimon hands more responsibilities to lieutenants
Wall Street's longest-serving chief executive says he is 'more like the coach now'
AN HOUR AGO
Lehman/US bank capital: loss cause
After the crisis, the banks are safer but debt is a danger
Financial crisis 2008: A reporter's memories from the front lines
FT Series Financial crisis: Are we safer now?
JPMorgan: defying attempts to end 'too big to fail'
AN HOUR AGO
Feedback
TODAY IN PERSONAL JOURNAL
A Celebrity Selfie Machine
PLUS Secondhand Stress at Work
THE WALL STREET JOURNAL.
WEDNESDAY, DECEMBER 11, 2013 - VOL. CCLXII NO. 138
WSJ.com
$2.00
What's News
Business & Finance
World-Wide
Deal Brings Stability to U.S. Budget
Congressional Negotiators Avert January Shutdown and Soften Sequester Cuts; Airline Fees to Climb
Thousands Honor Memory of South Africa's Liberator
RIP NELSON MANDELA
Longtime Insider Is GM's First Female CEO
PARTY DISCIPLINE
China Spins New Lesson From Soviet Fall
Ukrainian Forces Confront Protesters
Bank Rule Challenges Wall Street
Here's Your Holiday Bonus, Now Start Running
Workers Win All-They-Can-Grab Sprees From Companies; 'Supermarket Sweep'
Every investor is unique. That's why there's iShares Core.
iShares by BLACKROCK

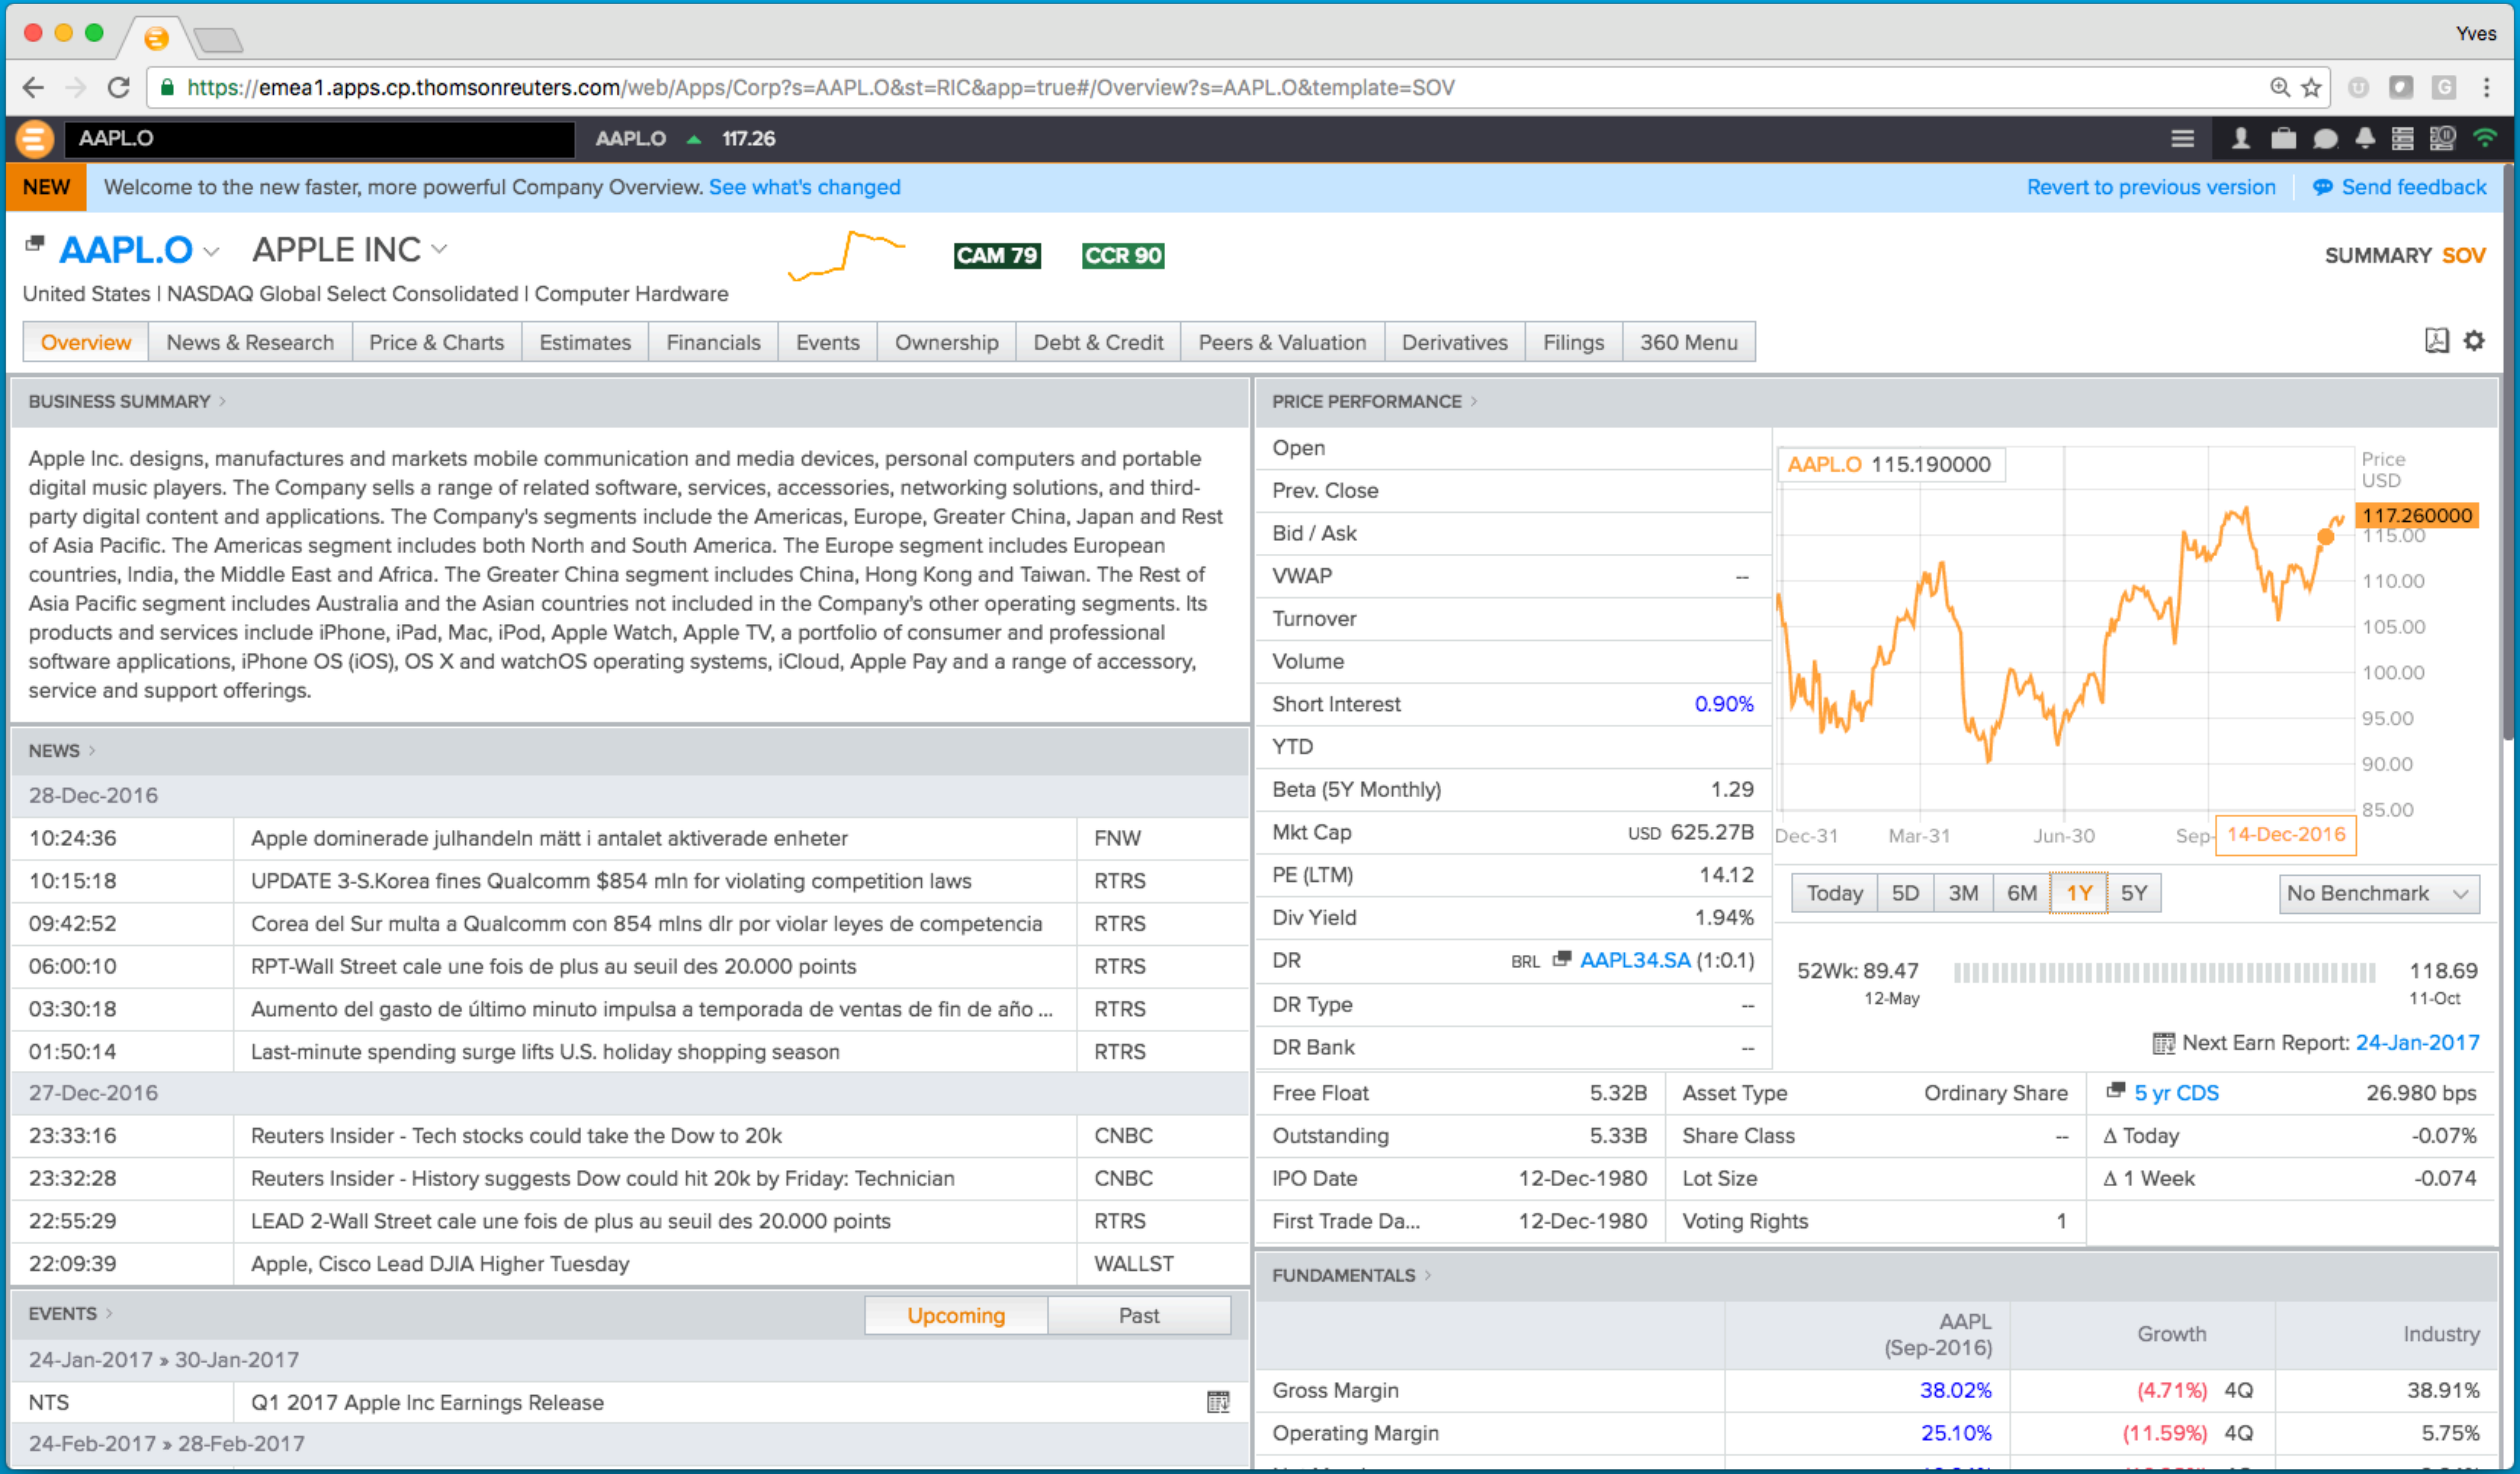
Yves
https://emea1.apps.cp.thomsonreuters.com/web/Apps/Corp?s=AAPL.O&st=RIC&app=true#/Overview?s=AAPL.O&template=SOV
AAPL.O
AAPL.O 117.26
NEW
Welcome to the new faster, more powerful Company Overview. See what's changed
Revert to previous version
Send feedback
AAPL.O
APPLE INC
CAM 79
CCR 90
SUMMARY SOV
United States | NASDAQ Global Select Consolidated | Computer Hardware
Overview
News & Research
Price & Charts
Estimates
Financials
Events
Ownership
Debt & Credit
Peers & Valuation
Derivatives
Filings
360 Menu
BUSINESS SUMMARY
Apple Inc. designs, manufactures and markets mobile communication and media devices, personal computers and portable digital music players. The Company sells a range of related software, services, accessories, networking solutions, and third-party digital content and applications. The Company's segments include the Americas, Europe, Greater China, Japan and Rest of Asia Pacific. The Americas segment includes both North and South America. The Europe segment includes European countries, India, the Middle East and Africa. The Greater China segment includes China, Hong Kong and Taiwan. The Rest of Asia Pacific segment includes Australia and the Asian countries not included in the Company's other operating segments. Its products and services include iPhone, iPad, Mac, iPod, Apple Watch, Apple TV, a portfolio of consumer and professional software applications, iPhone OS (iOS), OS X and watchOS operating systems, iCloud, Apple Pay and a range of accessory, service and support offerings.
NEWS
28-Dec-2016
10:24:36 Apple dominerade julhandeln mätt i antalet aktiverade enheter FNW
10:15:18 UPDATE 3-S.Korea fines Qualcomm $854 mln for violating competition laws RTRS
09:42:52 Corea del Sur multa a Qualcomm con 854 mlns dlr por violar leyes de competencia RTRS
06:00:10 RPT-Wall Street cale une fois de plus au seuil des 20.000 points RTRS
03:30:18 Aumento del gasto de último minuto impulsa a temporada de ventas de fin de año ... RTRS
01:50:14 Last-minute spending surge lifts U.S. holiday shopping season RTRS
27-Dec-2016
23:33:16 Reuters Insider - Tech stocks could take the Dow to 20k CNBC
23:32:28 Reuters Insider - History suggests Dow could hit 20k by Friday: Technician CNBC
22:55:29 LEAD 2-Wall Street cale une fois de plus au seuil des 20.000 points RTRS
22:09:39 Apple, Cisco Lead DJIA Higher Tuesday WALLST
EVENTS
Upcoming
Past
24-Jan-2017 » 30-Jan-2017
NTS Q1 2017 Apple Inc Earnings Release
24-Feb-2017 » 28-Feb-2017
PRICE PERFORMANCE
Open
Prev. Close
Bid / Ask
VWAP --
Turnover
Volume
Short Interest 0.90%
YTD
Beta (5Y Monthly) 1.29
Mkt Cap USD 625.27B
PE (LTM) 14.12
Div Yield 1.94%
DR BRL AAPL34.SA (1:0.1)
DR Type --
DR Bank --
AAPL.O 115.190000
Price USD
117.260000
115.00
110.00
105.00
100.00
95.00
90.00
85.00
Dec-31
Mar-31
Jun-30
Sep-
14-Dec-2016
Today
5D
3M
6M
1Y
5Y
No Benchmark
52Wk: 89.47
12-May
118.69
11-Oct
Next Earn Report: 24-Jan-2017
Free Float 5.32B
Outstanding 5.33B
IPO Date 12-Dec-1980
First Trade Da... 12-Dec-1980
Asset Type Ordinary Share
Share Class --
Lot Size
Voting Rights 1
5 yr CDS 26.980 bps
Δ Today -0.07%
Δ 1 Week -0.074
FUNDAMENTALS
AAPL (Sep-2016)
Growth
Industry
Gross Margin 38.02% (4.71%) 4Q 38.91%
Operating Margin 25.10% (11.59%) 4Q 5.75%

Tick Data

In [23]:
```
tick = ek.get_timeseries(['AAPL.O'],
                         fields='*',
                         start_date='2017-07-11 16:00:0000',
                         end_date='2017-07-11 16:15:0000',
                         interval='tick')
```

In [24]:
```
tick.info()
```

```
<class 'pandas.core.frame.DataFrame'>
DatetimeIndex: 1898 entries, 2017-07-11 16:00:00.686000 to 2017-07-11 16:14:59.708000
Data columns (total 2 columns):
VALUE     1892 non-null float64
VOLUME    1898 non-null float64
dtypes: float64(2)
memory usage: 44.5 KB
```

In [25]:
```
tick.tail()
```

Out[25]:

AAPL.O Date	VALUE	VOLUME
2017-07-11 16:14:59.693	144.9900	100.0
2017-07-11 16:14:59.693	144.9900	100.0
2017-07-11 16:14:59.693	144.9900	100.0
2017-07-11 16:14:59.707	144.9899	400.0
2017-07-11 16:14:59.708	144.9899	1305.0

News

In [29]:
```python
news = ek.get_news_headlines('R:.SPX AND "Trump" AND Language:LEN', count=5)
news
```

Out[29]:

	versionCreated	text	storyId	sourceCode
2017-08-18 16:46:19	2017-08-18 16:46:19	U.S. STOCKS EXTEND GAINS AFTER NEW YORK TIMES ...	urn:newsml:reuters.com:20170818:nL4N1L44L9:1	NS:RTRS
2017-08-18 15:53:08	2017-08-18 15:53:08	CORRECTED-U.S. STOCKS PARE LOSSES, TRADERS CIT...	urn:newsml:reuters.com:20170818:nL4N1L44IK:1	NS:RTRS
2017-08-18 15:16:27	2017-08-18 15:16:27	US STOCKS-Wall St lower on growing concerns ov...	urn:newsml:reuters.com:20170818:nL4N1L44F2:5	NS:RTRS
2017-08-18 11:24:30	2017-08-18 11:24:30	US STOCKS-Futures flat amid growing concerns o...	urn:newsml:reuters.com:20170818:nL4N1L43RR:5	NS:RTRS
2017-08-17 17:09:05	2017-08-17 17:09:05	US STOCKS-Wall St extends losses on Trump poli...	urn:newsml:reuters.com:20170817:nL4N1L34N1:5	NS:RTRS

In [30]:
```python
storyId = news.iloc[4, 2]
storyId
```

Out[30]: 'urn:newsml:reuters.com:20170817:nL4N1L34N1:5'

In [31]:
```python
from IPython.display import display, HTML
```

In [32]:
```python
display(HTML(ek.get_news_story(storyId)))
```

- Gary Cohn resignation rumors knocked down
- Wal-Mart drops after reporting margin fall
- Indexes down: Dow 0.81 pct, S&P 1.03 pct, Nasdaq 1.39 pct

Updates to early afternoon

By Sruthi Shankar and Tanya Agrawal

Aug 17 (Reuters) - U.S stocks hit session lows in early afternoon trading on Thursday as investors worried about President Donald Trump's ability to

	structured data	unstructured data	alternative data
historical data	price data (eod, minute, tick, …) fundamental data	texts news IoT	web texts social media satellite data
streaming data	tick data volume data	news IoT	web texts social media satellite data

EXPERT OPINION

Contact Editor: **Brian Brannon**, bbrannon@computer.org

The Unreasonable Effectiveness of Data

Alon Halevy, Peter Norvig, and Fernando Pereira, *Google*

Eugene Wigner's article "The Unreasonable Effectiveness of Mathematics in the Natural Sciences"[1] examines why so much of physics can be neatly explained with simple mathematical formulas such as $f = ma$ or $e = mc^2$. Meanwhile, sciences that involve human beings rather than elementary particles have proven more resistant to elegant mathematics. Economists suffer from physics envy over their inability to neatly model human behavior. An informal, incomplete grammar of the English language runs over 1,700 pages.[2] Perhaps when it comes to natural language processing and related fields, we're doomed to complex theories that will never have the elegance of physics equations. But if that's so, we should stop acting as if our goal is to author extremely elegant theories, and instead embrace complexity and make use of the best ally we have: the unreasonable effectiveness of data.

One of us, as an undergraduate at Brown University, remembers the excitement of having access to the Brown Corpus, containing one million English words.[3] Since then, our field has seen several notable corpora that are about 100 times larger, and in 2006, Google released a trillion-word corpus with frequency counts for all sequences up to five words long.[4] In some ways this corpus is a step backwards from the Brown Corpus: it's taken from unfiltered Web pages and thus contains incomplete sentences, spelling errors, grammatical errors, and all sorts of other errors. It's not annotated with carefully hand-corrected part-of-speech tags. But the fact that it's a million times larger than the Brown Corpus outweighs these drawbacks. A trillion-word corpus—along with other Web-derived corpora of millions, billions, or trillions of links, videos, images, tables, and user interactions—captures even very rare aspects of human behavior. So, this corpus could serve as the basis of a complete model for certain tasks—if only we knew how to extract the model from the data.

Learning from Text at Web Scale

The biggest successes in natural-language-related machine learning have been statistical speech recognition and statistical machine translation. The reason for these successes is not that these tasks are easier than other tasks; they are in fact much harder than tasks such as document classification that extract just a few bits of information from each document. The reason is that translation is a natural task routinely done every day for a real human need (think of the operations of the European Union or of news agencies). The same is true of speech transcription (think of closed-caption broadcasts). In other words, a large training set of the input-output behavior that we seek to automate is available to us *in the wild*. In contrast, traditional natural language processing problems such as document classification, part-of-speech tagging, named-entity recognition, or parsing are not routine tasks, so they have no large corpus available in the wild. Instead, a corpus for these tasks requires skilled human annotation. Such annotation is not only slow and expensive to acquire but also difficult for experts to agree on, being bedeviled by many of the difficulties we discuss later in relation to the Semantic Web. The first lesson of Web-scale learning is to use available large-scale data rather than hoping for annotated data that isn't available. For instance, we find that useful semantic relationships can be automatically learned from the statistics of search queries and the corresponding results[5] or from the accumulated evidence of Web-based text patterns and formatted tables,[6] in both cases without needing any manually annotated data.

1541-1672/09/$25.00 © 2009 IEEE
Published by the IEEE Computer Society

Eugene Wigner's article "The Unreasonable Effectiveness of Mathematics in the Natural Sciences" examines why so much of physics can be neatly explained with simple mathematical formulas such as f = ma or e = mc². Meanwhile, sciences that involve human beings rather than elementary particles have proven more resistant to elegant mathematics. Economists suffer from physics envy over their inability to neatly [and successfully] model human behavior. An informal, incomplete grammar of the English language runs over 1,700 pages. Perhaps when it comes to natural language processing and related fields, we're doomed to complex theories that will never have the elegance of physics equations. But if that's so, we should stop acting as if our goal is to author extremely elegant theories, and instead embrace complexity and make use of the best ally we have: the unreasonable effectiveness of data.

IP[y]:
IPython

jupyter

NumPy

matplotlib

pandas

$y_it = \beta' x_{it} + \mu_i + \epsilon_{it}$

xlwings

X

Statistical Learning

Mathematics.

$$f(x) = 2 + \frac{1}{2}x$$

$$y_i = f(x_i), i = 1,2,\ldots,n$$

Statistics.

$$(y_i, x_i)_{i=1}^{n}$$

$$\hat{f}(x) = \alpha + \beta x \approx y$$

$$\alpha, \beta = ?, ?$$

Mathematics.

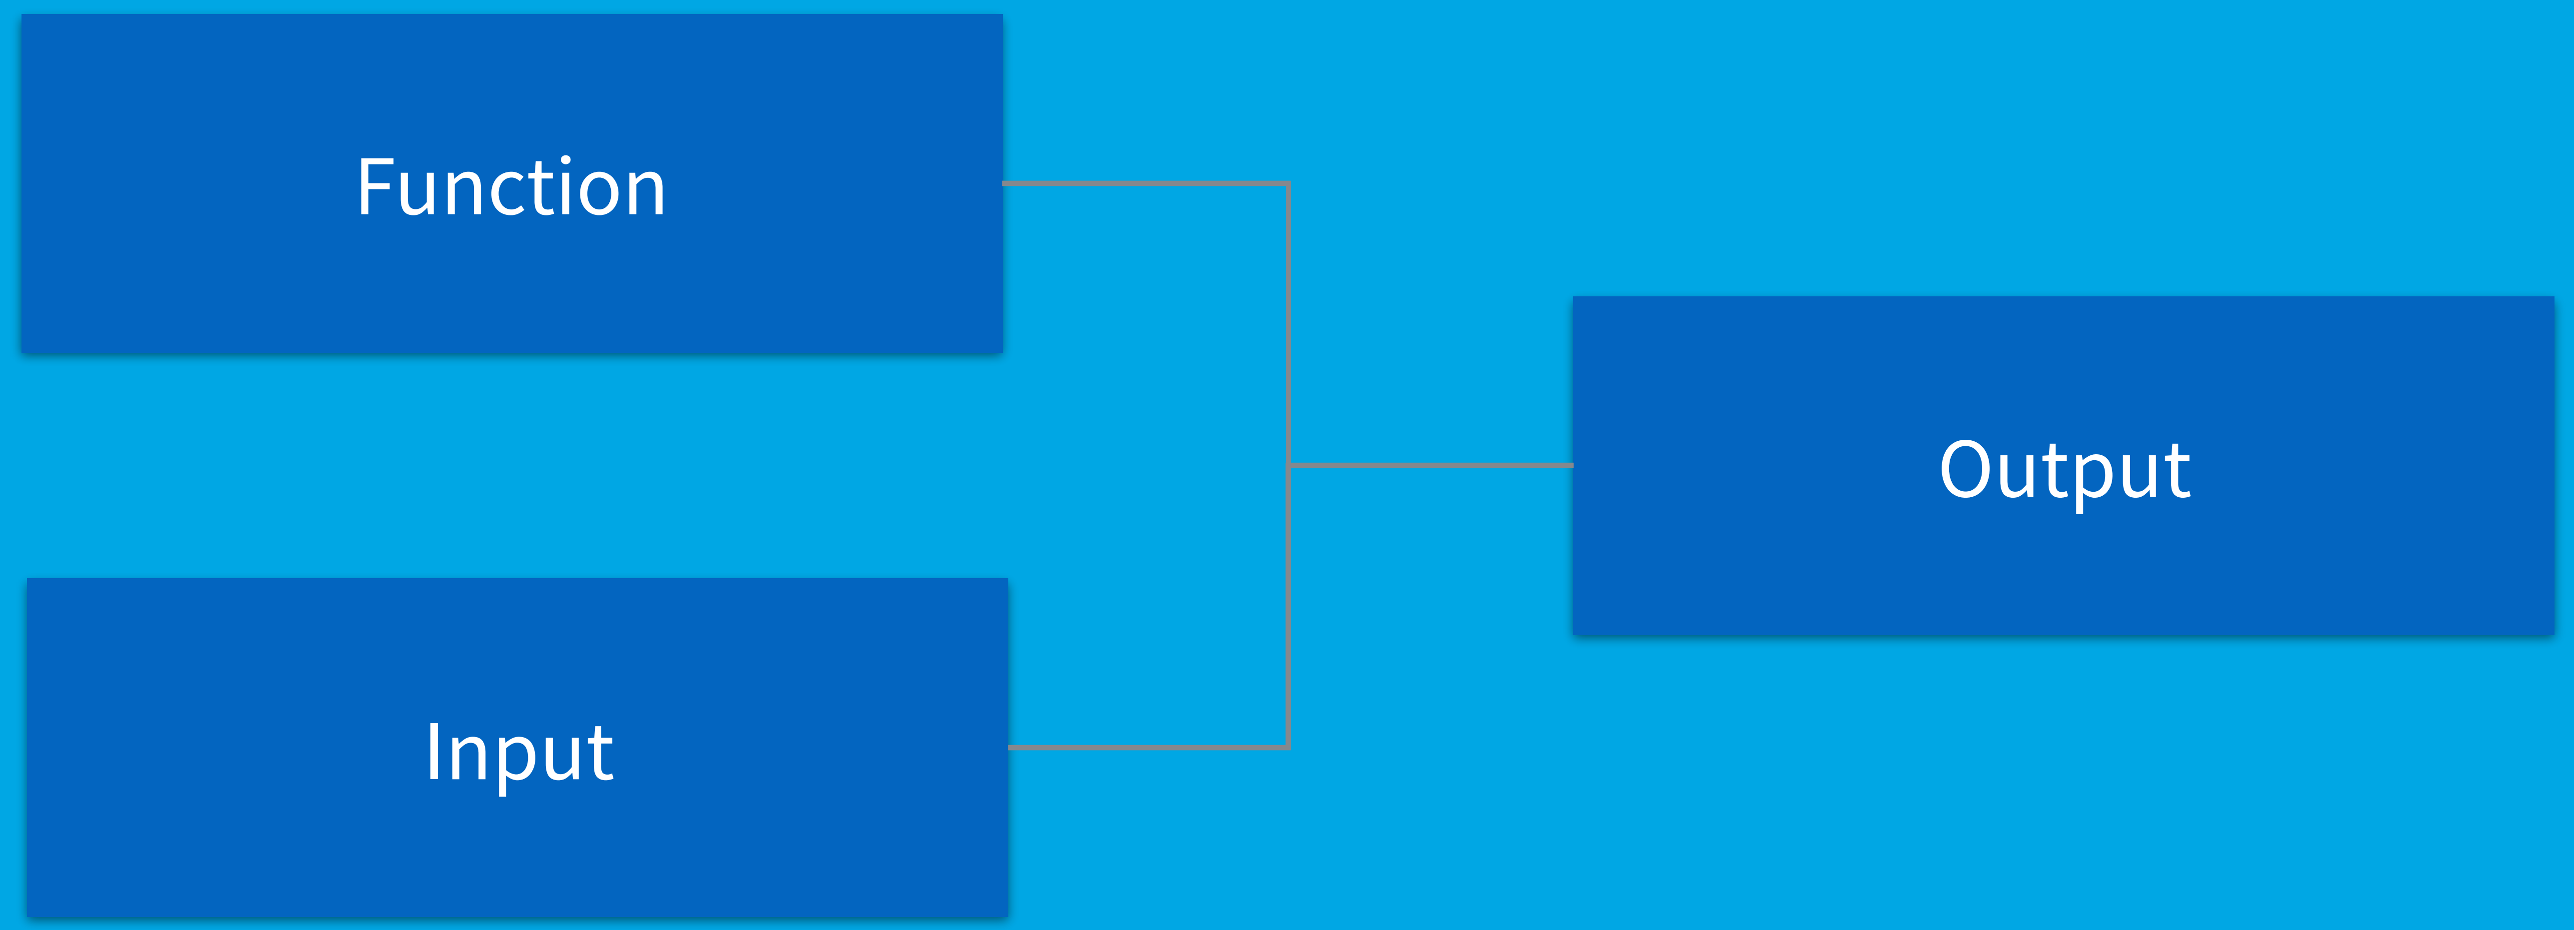

Statistics.

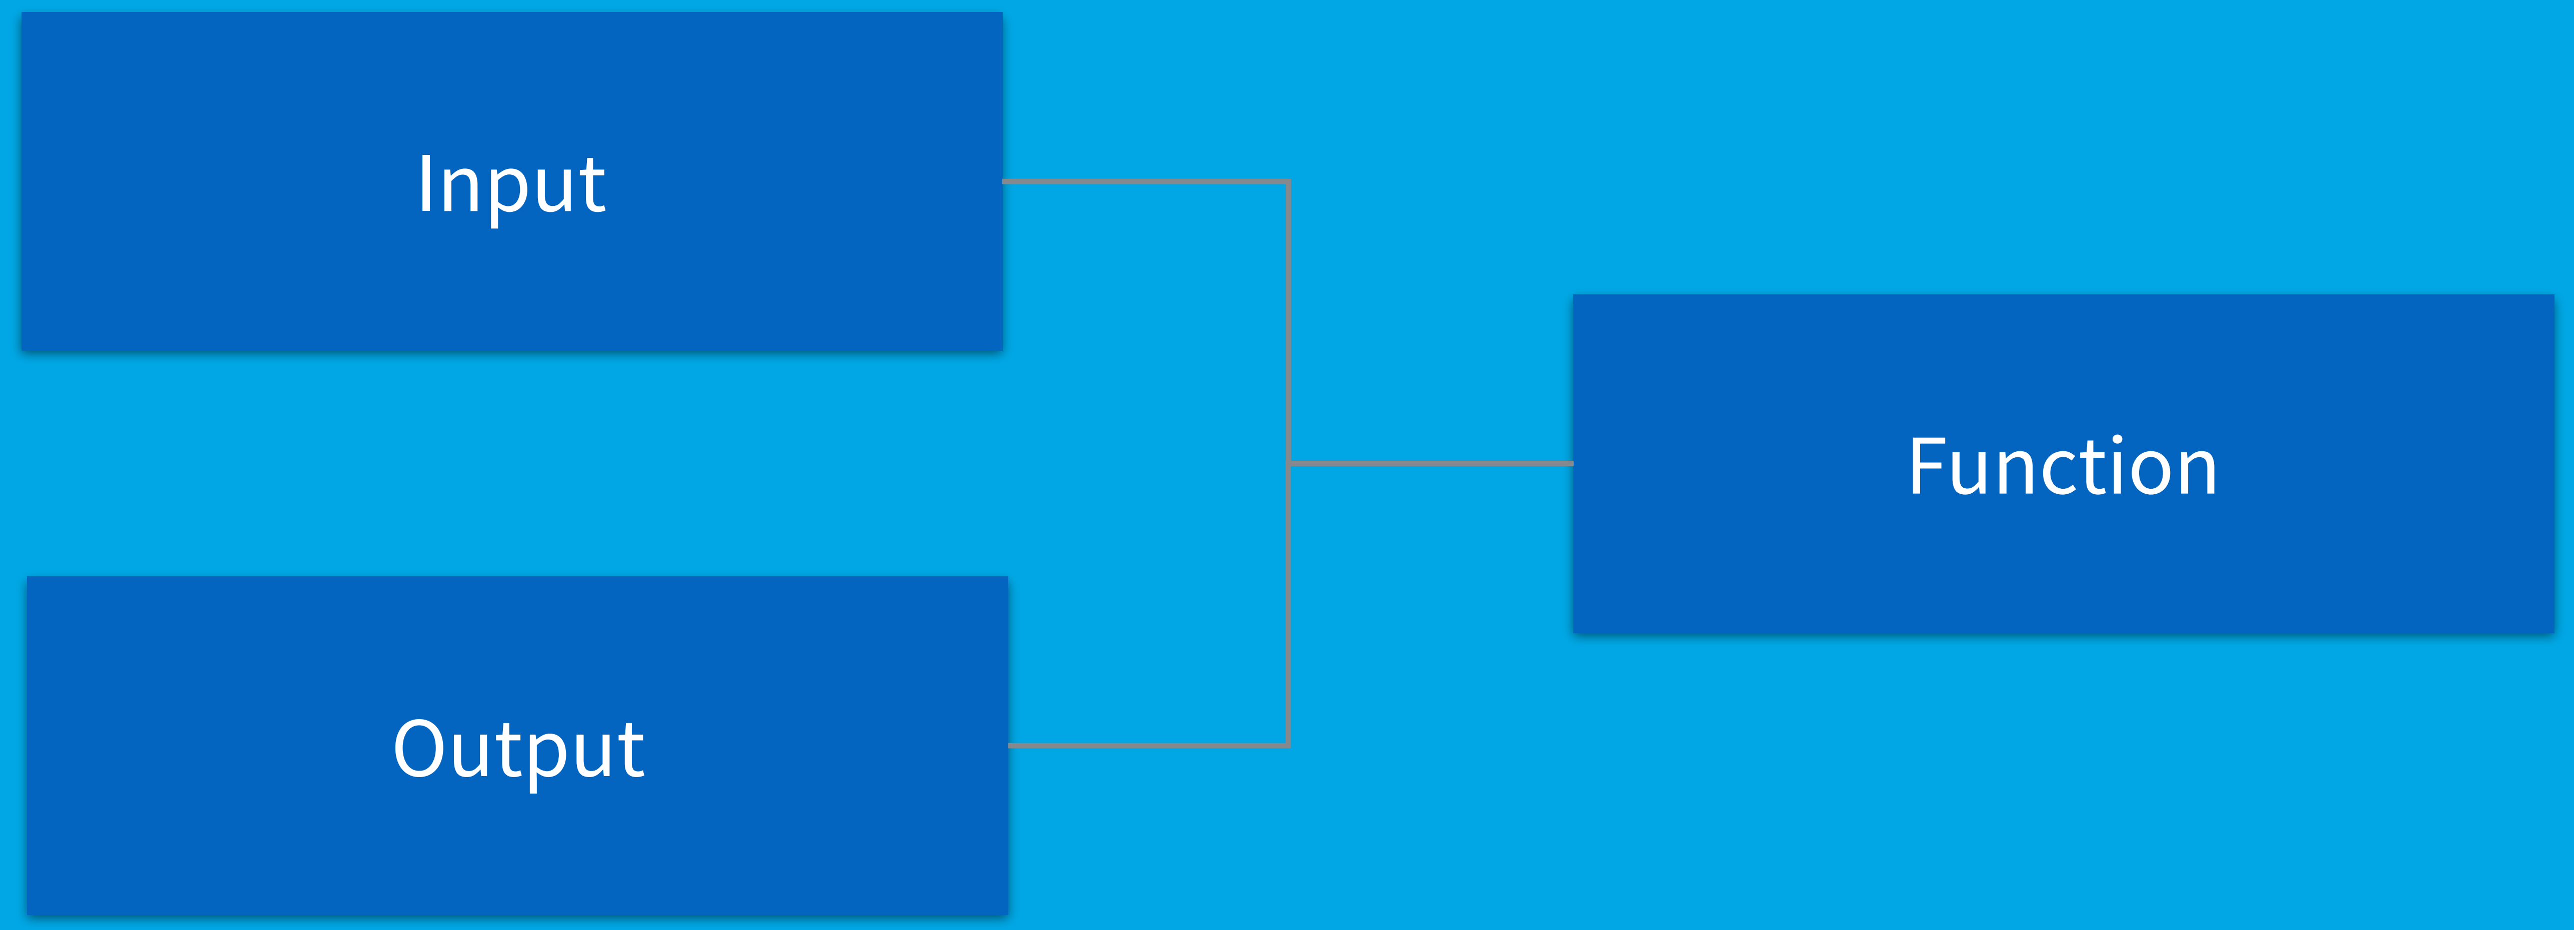

OLS Regression

Why OLS Regression?

1. **centuries old**: least squares approach used since more than 200 years
   (see e.g. this article)
2. **simple math**: easy to understand and transfer to different data sets
3. **lightning fast**: fast to evaluate even on large data sets
4. **scalable**: basically not limit regarding data size
5. **implementation**: efficient implementations (e.g. Python) readily available

Given input data.

$$(y_i, x_i)_{i=1}^{n}$$

Simple linear regression.

$$\hat{y}_i = \alpha + \beta x_i \approx y_i$$

$$y_i = \alpha + \beta x_i + \epsilon_i$$

Minimization problem.

$$\min_{\alpha,\beta} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2$$

Optimal Solution.

$$\beta = \frac{Cov(x, y)}{Var(x)}$$

$$\alpha = \bar{y} - \beta\bar{x}$$

Major assumptions of the linear regression model:

1. **linearity**: the model is linear in its parameters (coefficients and error term)
2. **independence**: independent variables should not be perfectly correlated with each other (no multicollinearity)
3. **zero mean**: the mean of the residuals should be zero
4. **no correlation**: residuals should not be correlated with the independent variables
5. **homoscedasticity**: the standard deviation of the residuals should be constant
6. **no autocorrelation**: the residuals should not be correlated with each other

Efficient Markets

Random Walks in Stock Market Prices

Eugene F. Fama

For many years economists, statisticians, and teachers of finance have been interested in developing and testing models of stock price behavior. One important model that has evolved from this research is the theory of random walks. This theory casts serious doubt on many other methods for describing and predicting stock price behavior—methods that have considerable popularity outside the academic world. For example, we shall see later that if the random walk theory is an accurate description of reality, then the various "technical" or "chartist" procedures for predicting stock prices are completely without value.

In general the theory of random walks raises challenging questions for anyone who has more than a passing interest in understanding the behavior of stock prices. Unfortunately, however, most discussions of the theory have appeared in technical academic journals and in a form which the non-mathematician would usually find incomprehensible. This article describes, briefly and simply, the theory of random walks and some of the important issues it raises concerning the work of market analysts. To preserve brevity some aspects of the theory and its implications are omitted. More complete (and also more technical) discussions of the theory of random walks are available elsewhere; hopefully the introduction provided here will encourage the reader to examine one of the more rigorous and lengthy works listed at the end of this article.

COMMON TECHNIQUES FOR PREDICTING STOCK MARKET PRICES

In order to put the theory of random walks into perspective we first discuss, in brief and general terms, the two approaches to predicting stock prices that are commonly espoused by market professionals. These are (1) "chartist" or "technical" theories and (2) the theory of fundamental or intrinsic value analysis.

The basic assumption of all the chartist or technical theories is that history tends to repeat itself, i.e., past patterns of price behavior in individual securities will tend to recur in the future. Thus the way to predict stock prices (and, of course, increase one's potential gains) is to develop a familiarity with past patterns of price behavior in order to recognize situations of likely recurrence.

Essentially, then, chartist techniques attempt to use knowledge of the past behavior of a price series to predict the probable future behavior of the series. A statistician would characterize such techniques as assuming that successive price changes in individual securities are dependent. That is, the various chartist theories assume that the *sequence* of price changes prior to any given day is important in predicting the price change for that day.[1]

The techniques of the chartist have always been surrounded by a certain degree of mysticism, however, and as a result most market professionals have found them suspect. Thus it is probably safe to say that the pure chartist is relatively rare among stock market analysts. Rather the typical analyst adheres to a technique known as fundamental analysis or the intrinsic value method. The assumption of the fundamental analysis approach is that at any point in time an individual security has an intrinsic value (or in the terms of the economist, an equilibrium price) which depends on the earning potential of the security. The earning potential of the security depends in turn on such fundamental factors as quality of management, outlook for the industry and the economy, etc.

Through a careful study of these fundamental factors the analyst should, in principle, be able to determine whether the actual price of a security is above or below its intrinsic value. If actual prices tend to move toward intrinsic values, then attempting to determine the intrinsic value of a security is equivalent to making a prediction of its future price; and this is the essence of the predictive procedure implicit in fundamental analysis.

THE THEORY OF RANDOM WALKS

Chartist theories and the theory of fundamental analysis are really the province of the market

Reprinted from Financial Analysts Journal *(September/October 1965):55–59.*

Eugene F. Fama (1965):

"For many years, economists, statisticians, and teachers of finance have been interested in developing and testing models of stock price behavior. One important model that has evolved from this research is the theory of random walks. This theory casts serious doubt on many other methods for describing and predicting stock price behavior—methods that have considerable popularity outside the academic world. For example, we shall see later that, if the random-walk theory is an accurate description of reality, then the various "technical" or "chartist" procedures for predicting stock prices are completely without value."—Eugene F. Fama (1965): "Random Walks in Stock Market Prices"

Michael Jensen (1978): "Some Anomalous Evidence Regarding Market Efficiency":

"A market is efficient with respect to an information set S if it is impossible to make economic profits by trading on the basis of information set S."

If a stock price follows a (simple) random walk (no drift & normally distributed returns), then it rises and falls with the same probability of 50% ("toss of a coin").

In such a case, the best predictor of tomorrow's stock price — in a least-squares sense — is today's stock price.

AI-First Finance

scientific method

noun

a method of procedure that has characterized natural science since the 17th century, consisting in systematic observation, measurement, and experiment, and the formulation, testing, and modification of hypotheses.

"criticism is the backbone of **the scientific method**"

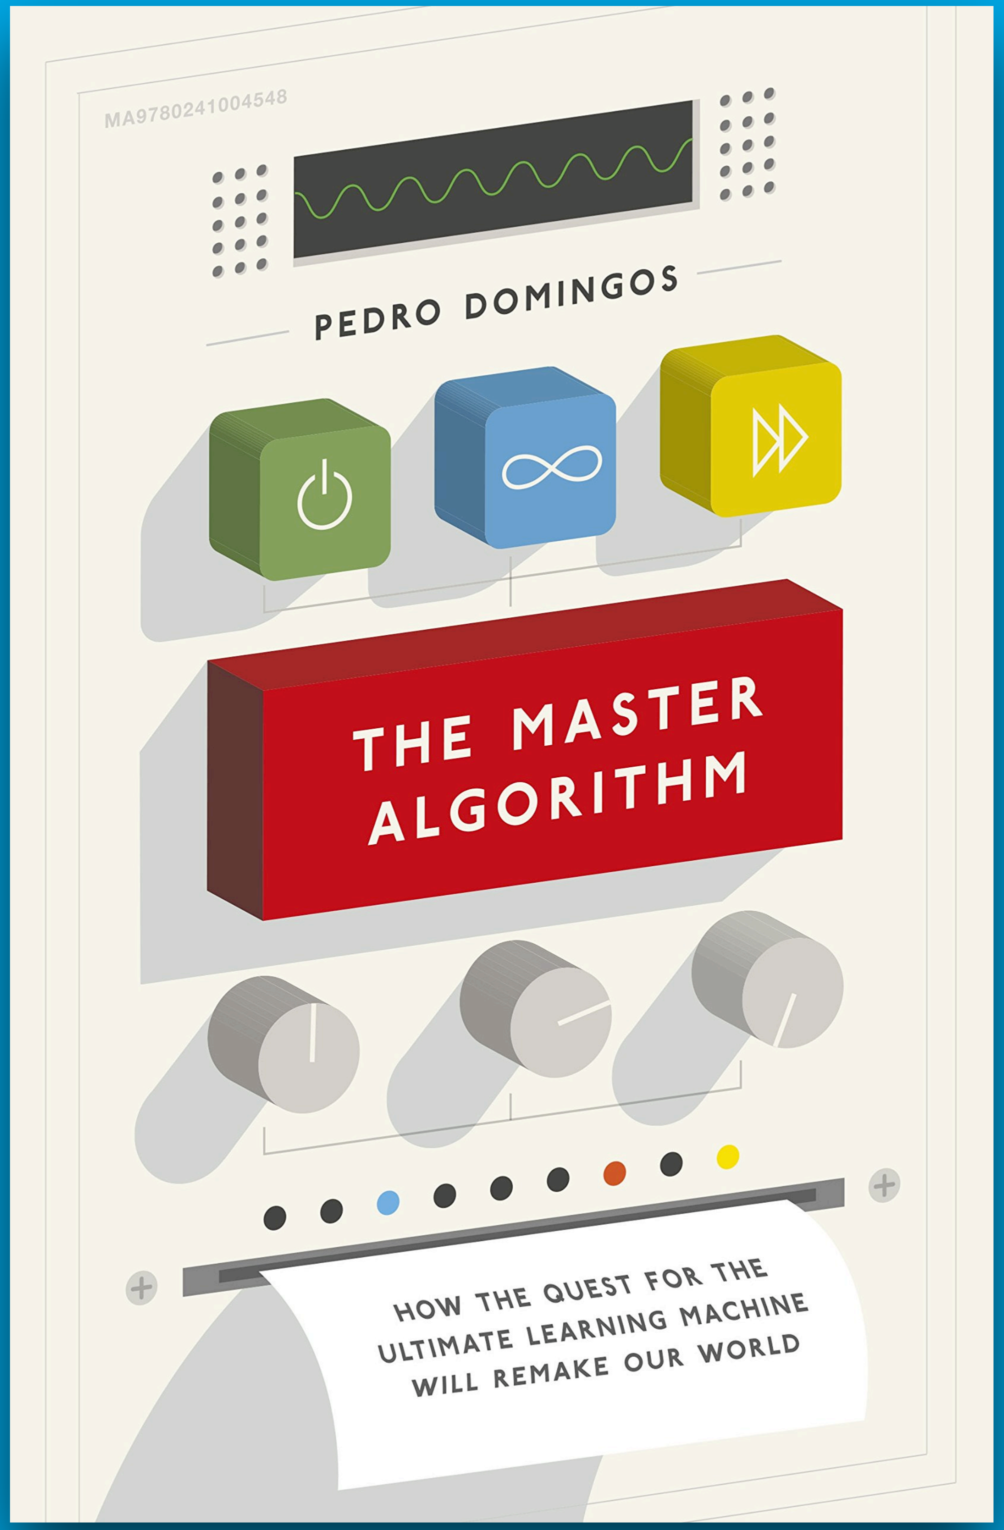

"The grand aim of science is to cover the greatest number of experimental facts by logical deduction from the smallest number of hypotheses or axioms."
— Albert Einstein

"Machine learning is the scientific method on steroids. It follows the same process of generating, testing, and discarding or refining hypotheses. But while a scientist may spend his or her whole life coming up with and testing a few hundred hypotheses, a machine-learning system can do the same in a second. Machine learning automates discovery. It's no surprise, then that it's revolutionizing science as much as it's revolutionizing business."

Programming.

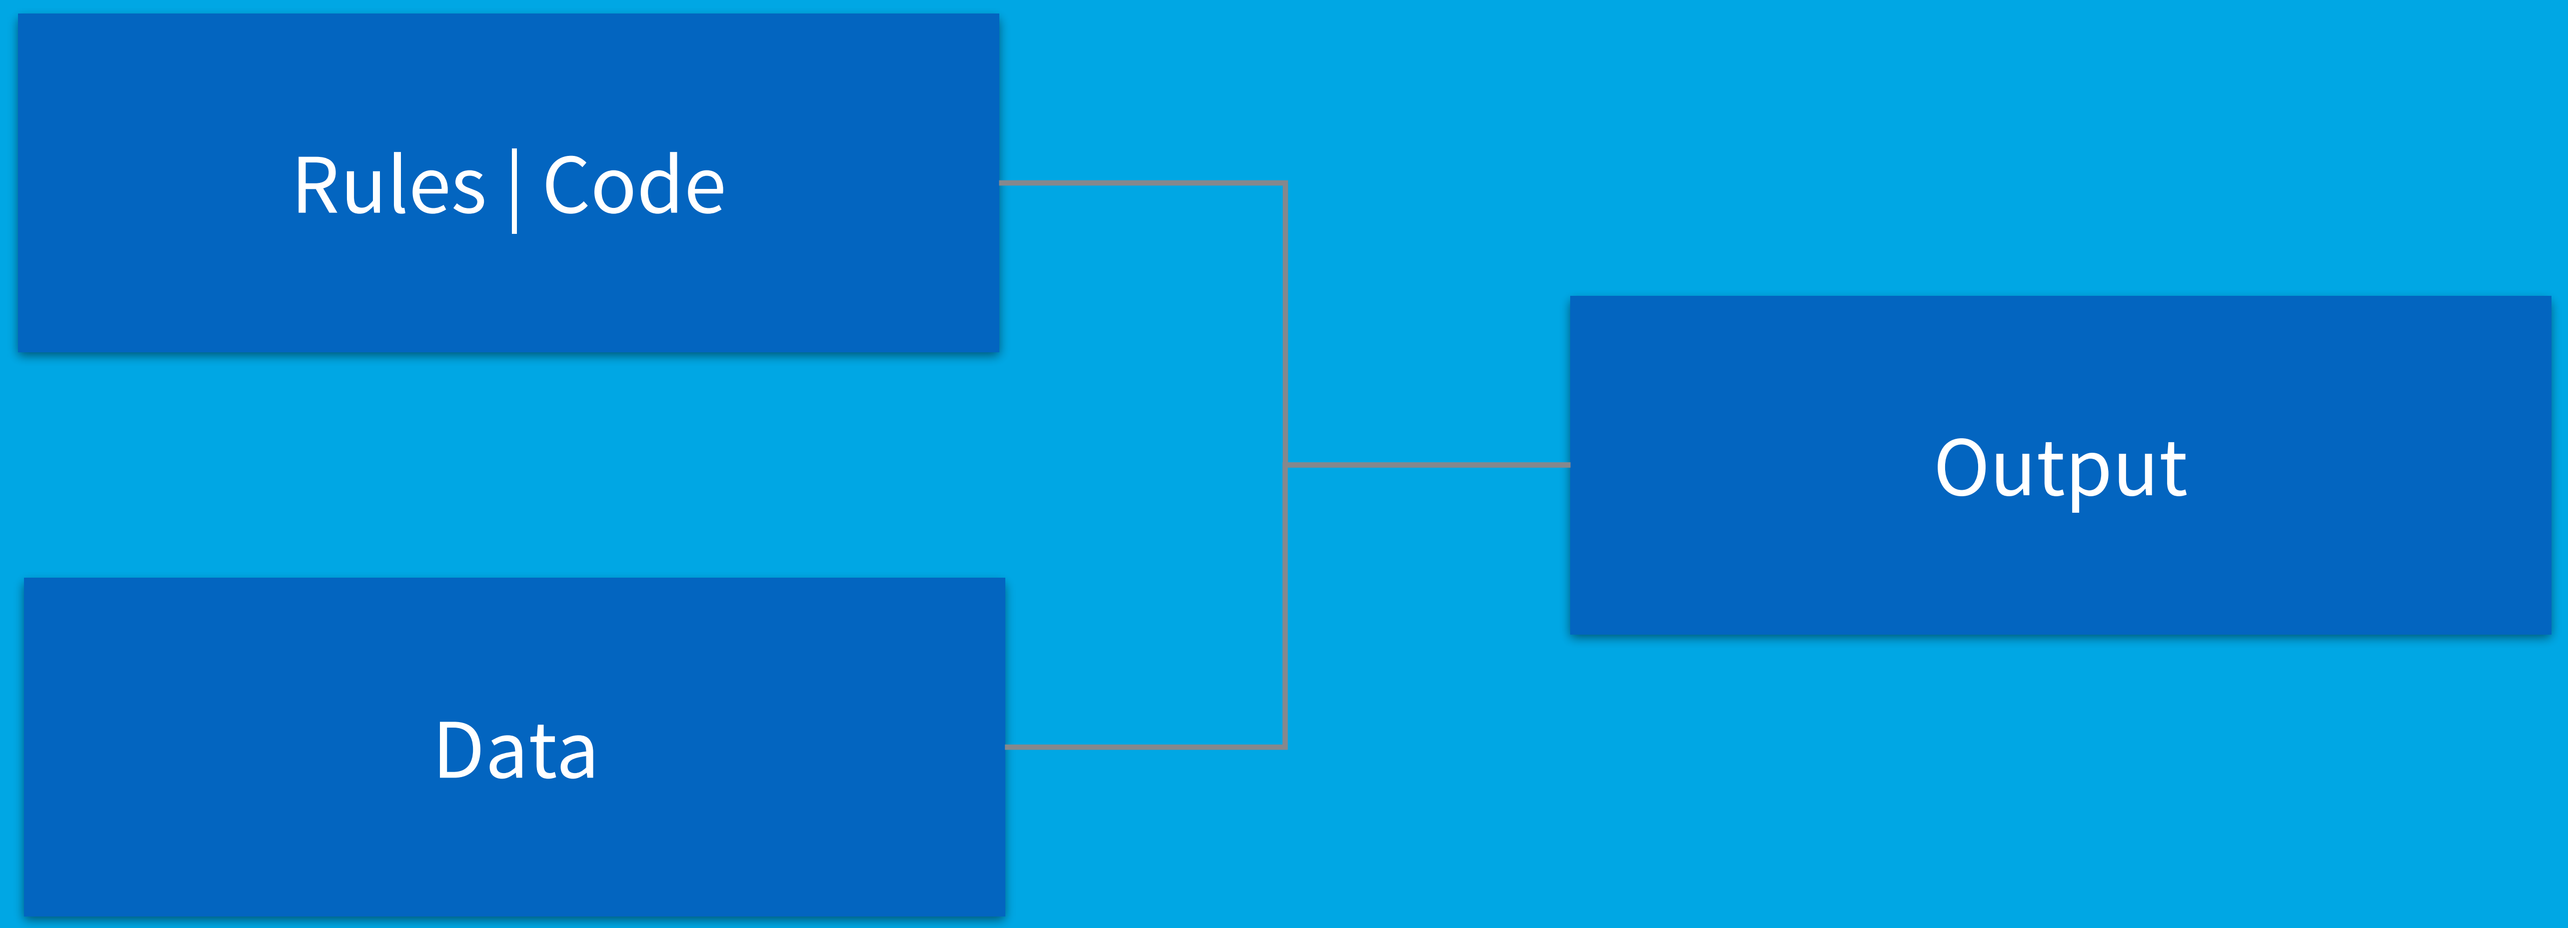

Machine Learning.

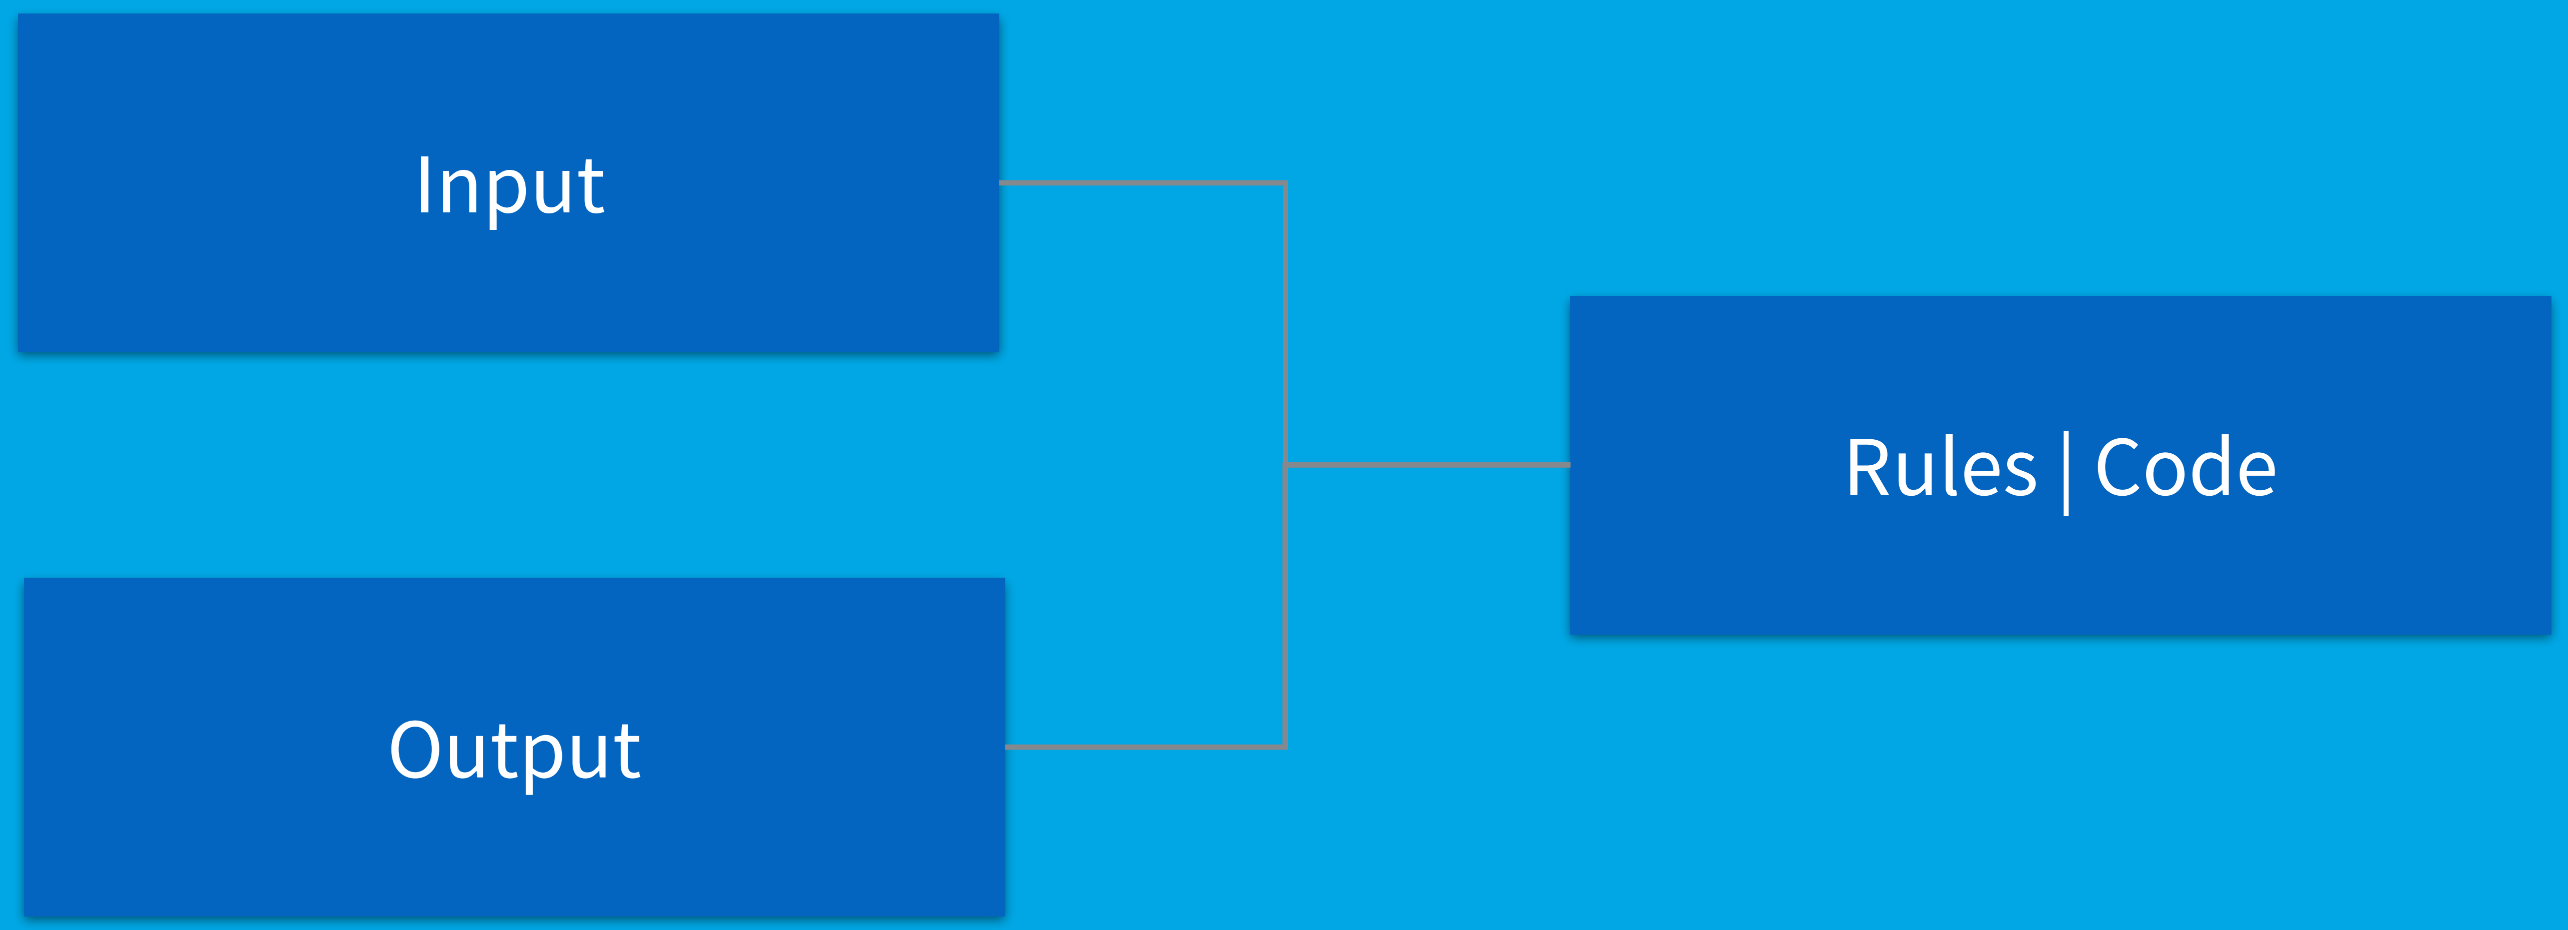

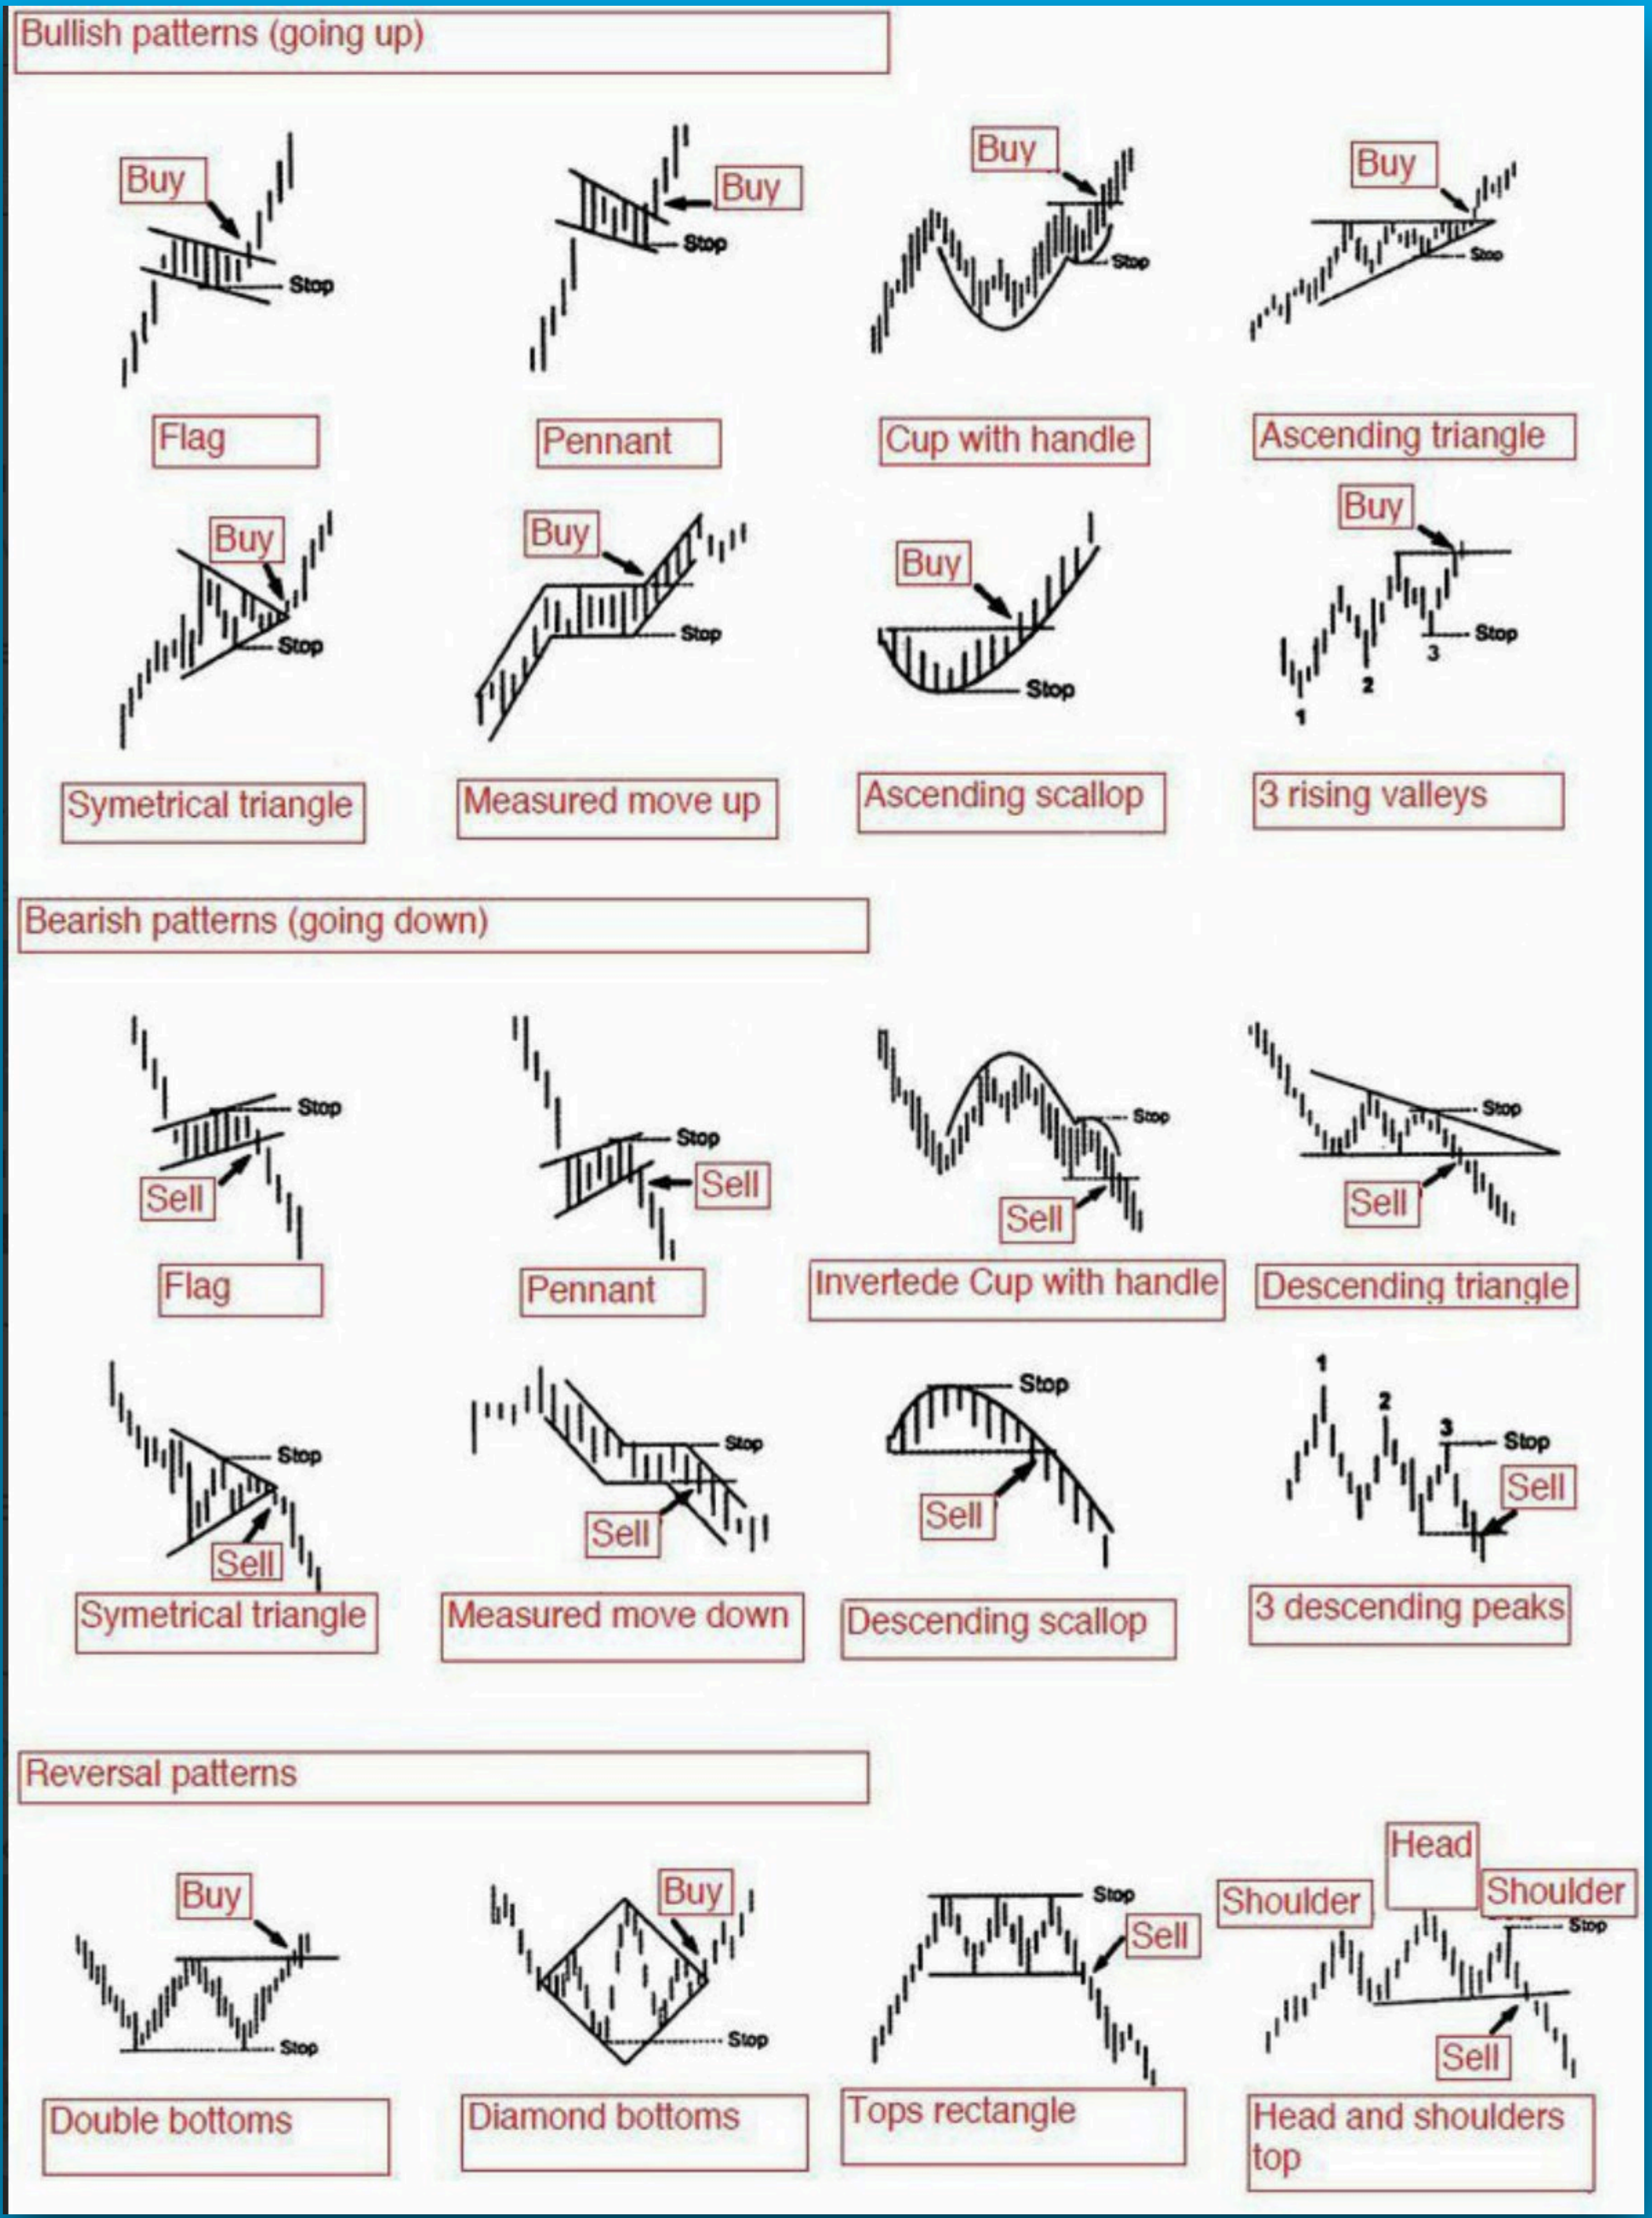
Bullish patterns (going up)
Flag
Pennant
Cup with handle
Ascending triangle
Symetrical triangle
Measured move up
Ascending scallop
3 rising valleys
Bearish patterns (going down)
Flag
Pennant
Invertede Cup with handle
Descending triangle
Symetrical triangle
Measured move down
Descending scallop
3 descending peaks
Reversal patterns
Double bottoms
Diamond bottoms
Tops rectangle
Shoulder
Head
Shoulder
Head and shoulders top
Buy
Sell
Stop

Financial Markets

x

y

"non-linear, complex, changing"

Finance History

f(•)

f(x) ≠ y

"brain-driven & beauty myth"

AI in Finance = finaince

x

m(•, a, b)

m(x, a*, b*) ≈ y

"data-driven & AI-first"

Financial Markets

x

y

"non-linear, complex, changing"

Finance History

"normative economics = assumptions, axioms, etc."

(too) "simple and elegant theories"

f(•)

"hardly any supporting empirical evidence"

f(x) ≠ y

"brain-driven & beauty myth"

AI in Finance = finaince

"positive economics = data, relationships, etc."

x

"general, parametrizable, trainable algorithms"

m(•, a, b)

"might show good performance, but black box"

m(x, a*, b*) ≈ y

"data-driven & AI-first"

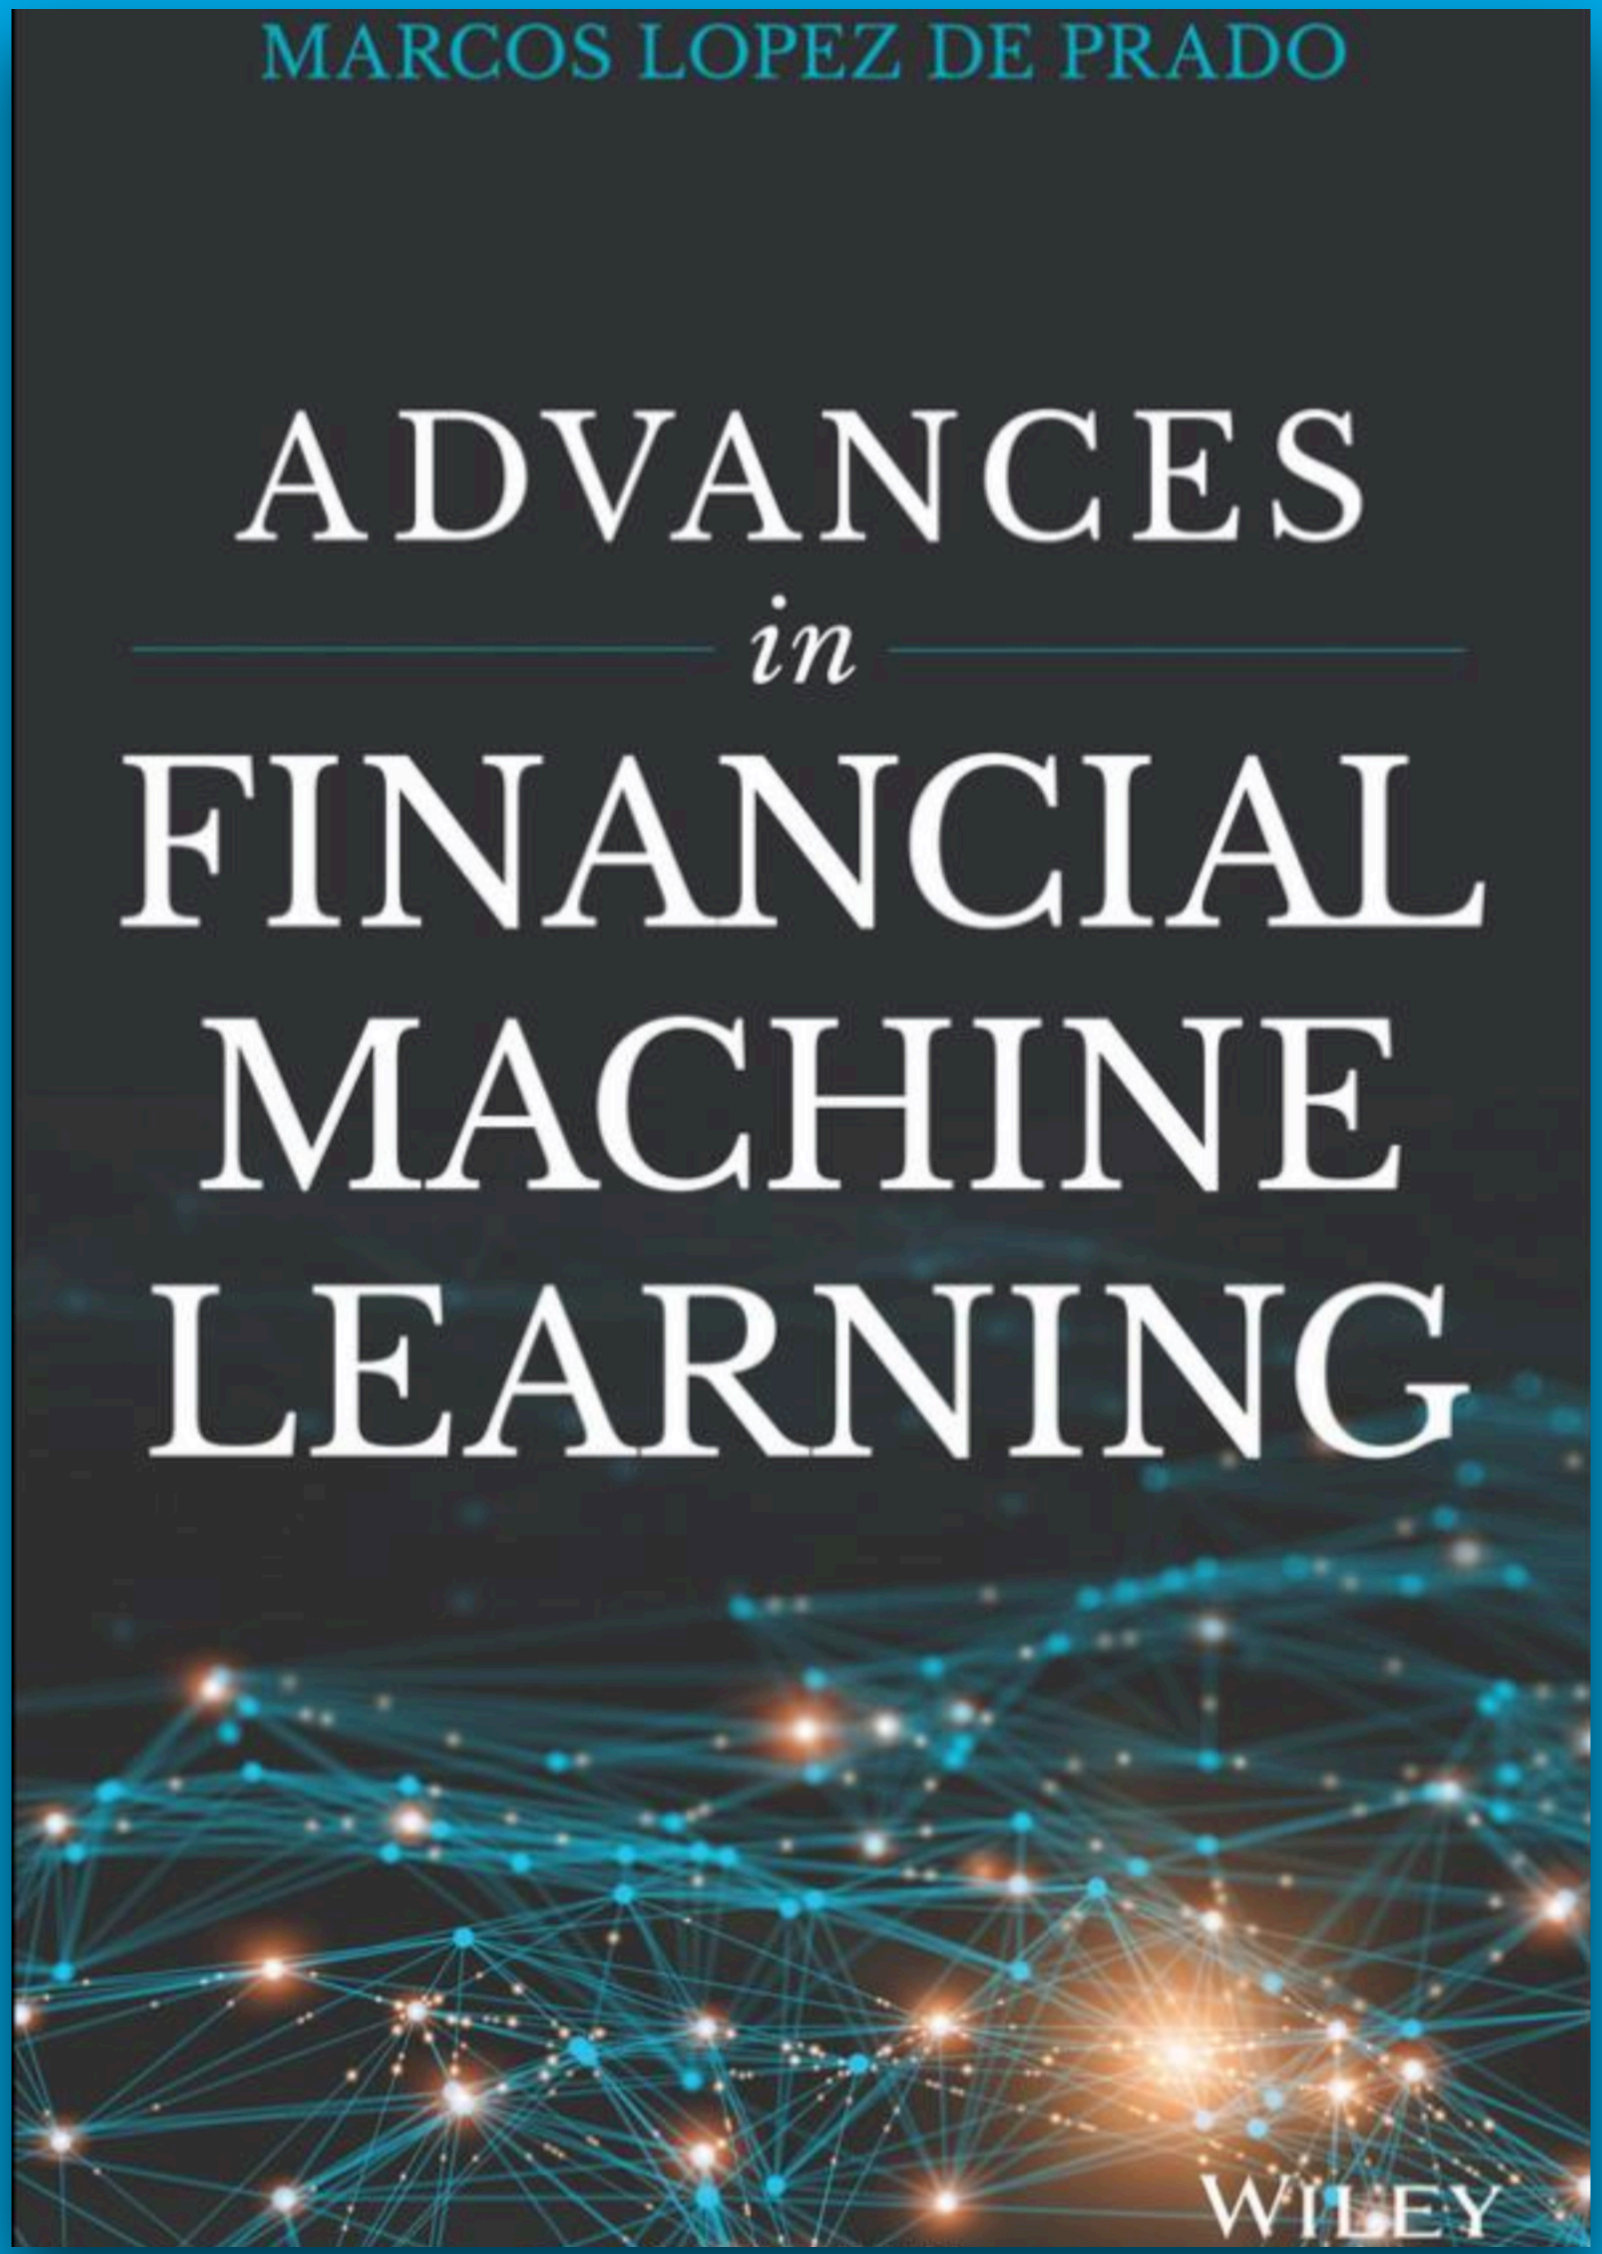

"The essential tool of econometrics is multivariate linear regression, an 18th-century technology that was already mastered by Gauss before 1794 … It is hard to believe that something as complex as 21st-century finance could be grasped by something as simple as inverting a covariance matrix."

"… what if economists finally started to consider non-linear functions?"

"An ML algorithm can spot patterns in a 100-dimensional world as easily as in our familiar 3-dimensional one."

"Econometrics might be good enough to succeed in financial academia (for now), but succeeding in practice requires ML."

Marcos López de Prado (2018)

scikit learn

TensorFlow

NVIDIA®

Algorithms

Artificial Intelligence

Machine Learning
(LogReg, Gaussian NB,
Decision Trees, SVM)

Deep Learning
(DNN, CNN, RNN)

Reinforcement Learning
(Simple, Q-Learning, DRL)

Unsupervised Learning
(Clustering, Dim Reduction)

Classification

Supervised Learning

Estimation

Online Learning

Policies

Supervised Learning.

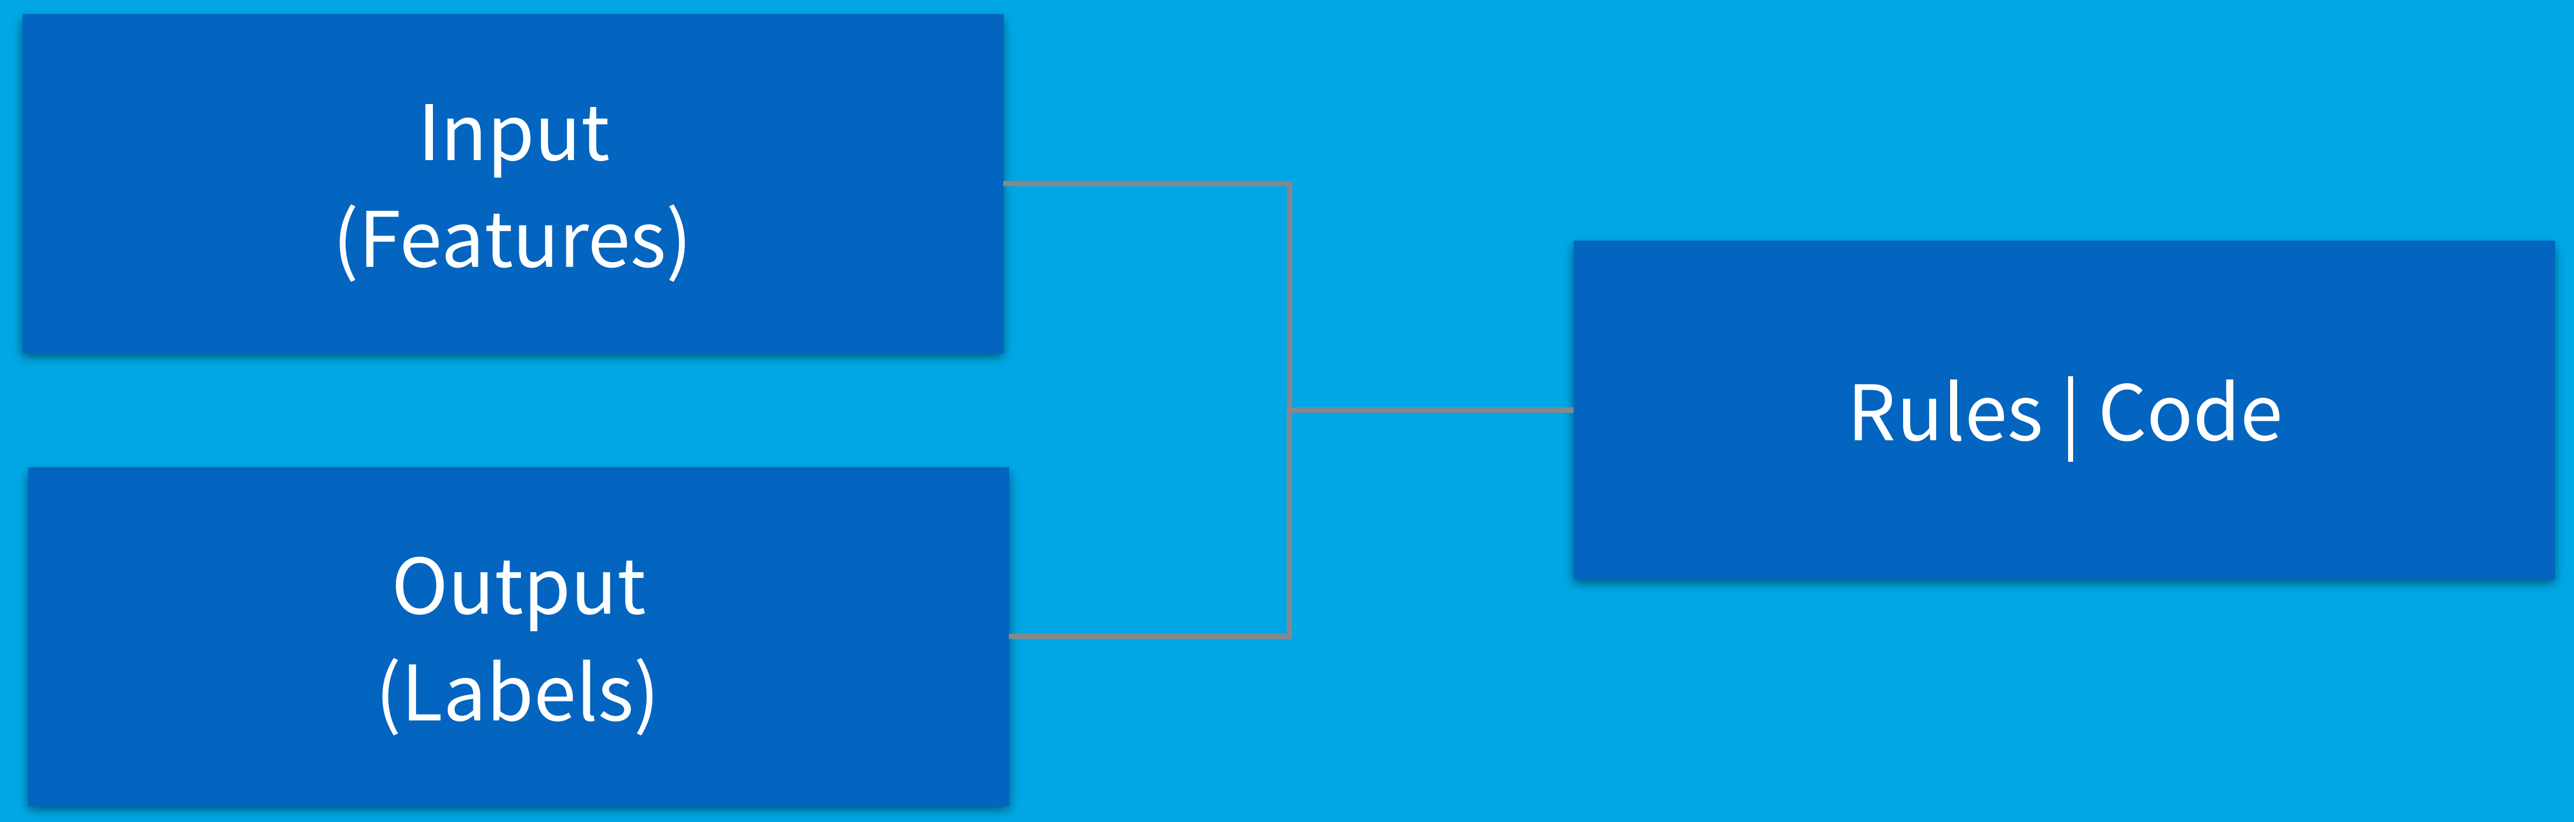

Unsupervised Learning.

Online Learning.

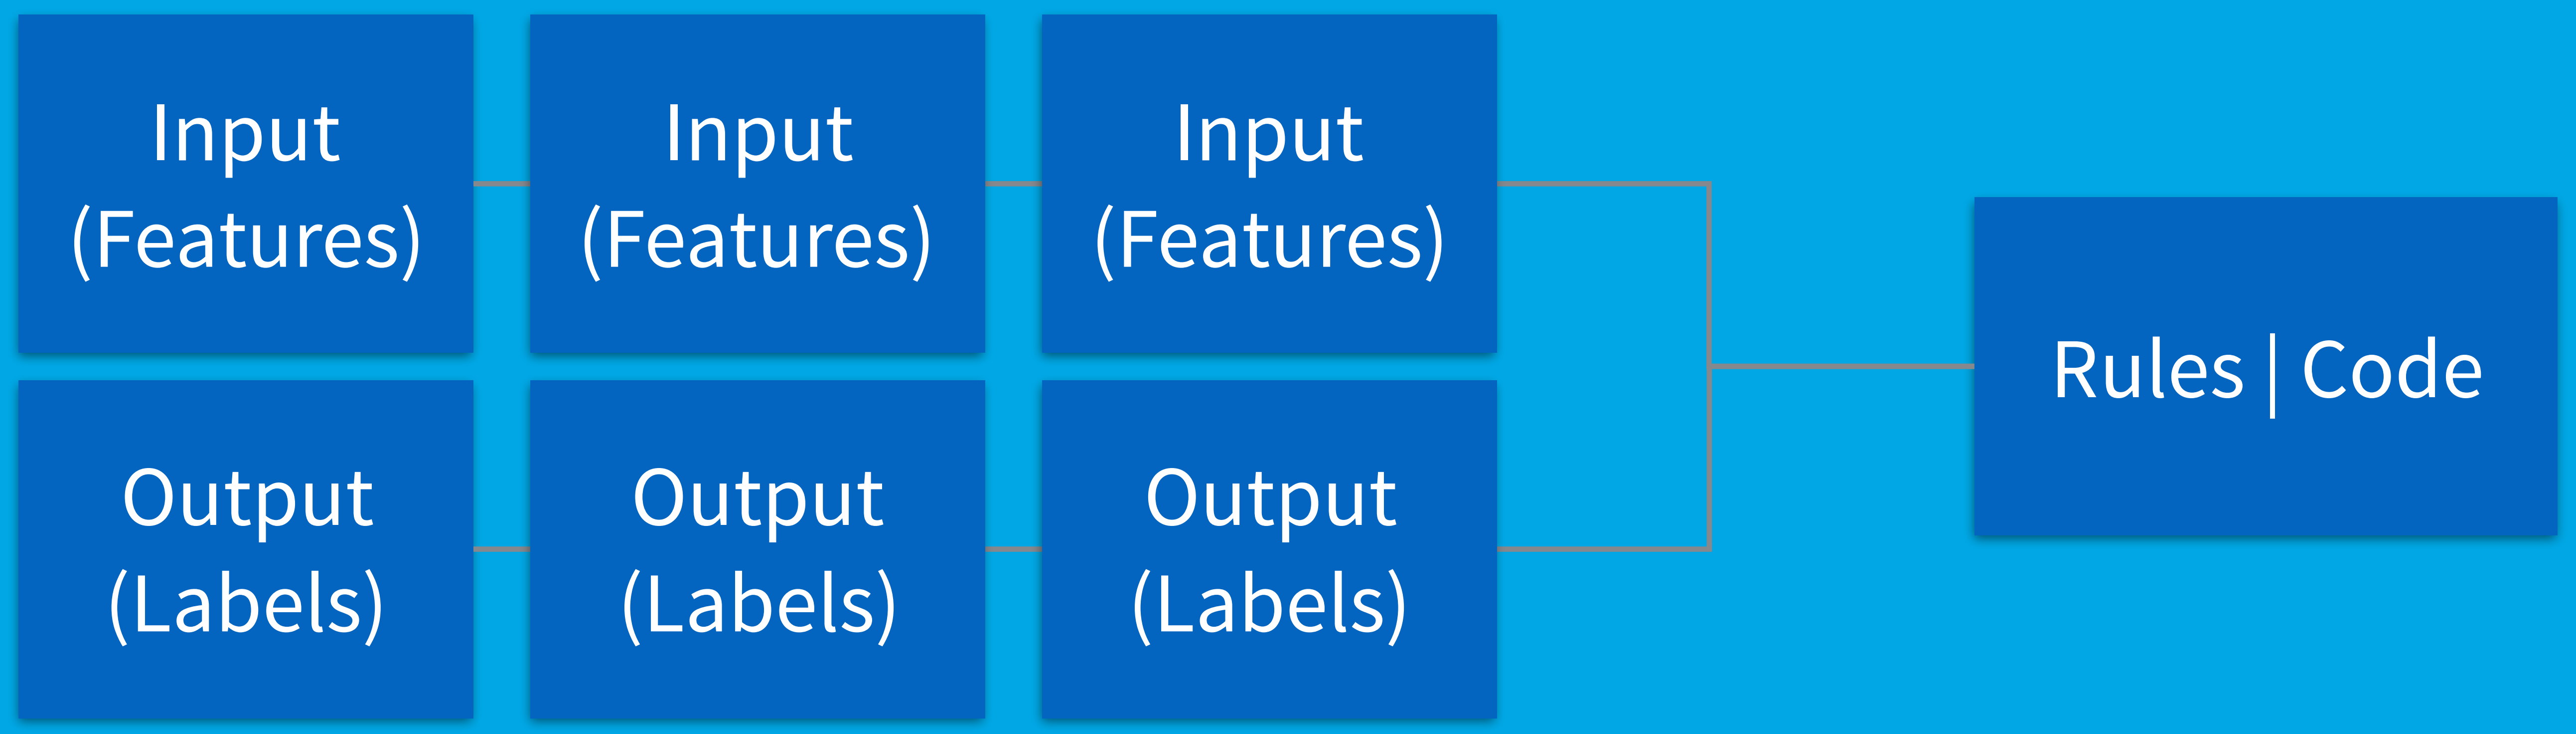

Some specifications and explanations:

1. **supervised learning**: input data (features) and output data (labels) are given; the algorithm learns from observed patterns
2. **unsupervised learning**: only input data (features) are given; the algorithm identifies patterns, cluster, etc.
3. **online learning**: both input data (features) and output data (labels) arrive incrementally (over time); the algorithm updates its parameters (policies) incrementally
4. **classification**: the problem of learning about and predicting labels as two or more discrete categories (e.g. `{0, 1}` or `{A, B, C}`)
5. **estimation**: the problem of learning about and predicting labels as continuous values (real numbers, floating point numbers, e.g. `1.435`)
6. **policies**: the problem of learning about and applying action policies (e.g. `if (x=1, y=0.5, z='low') then take action B2`)

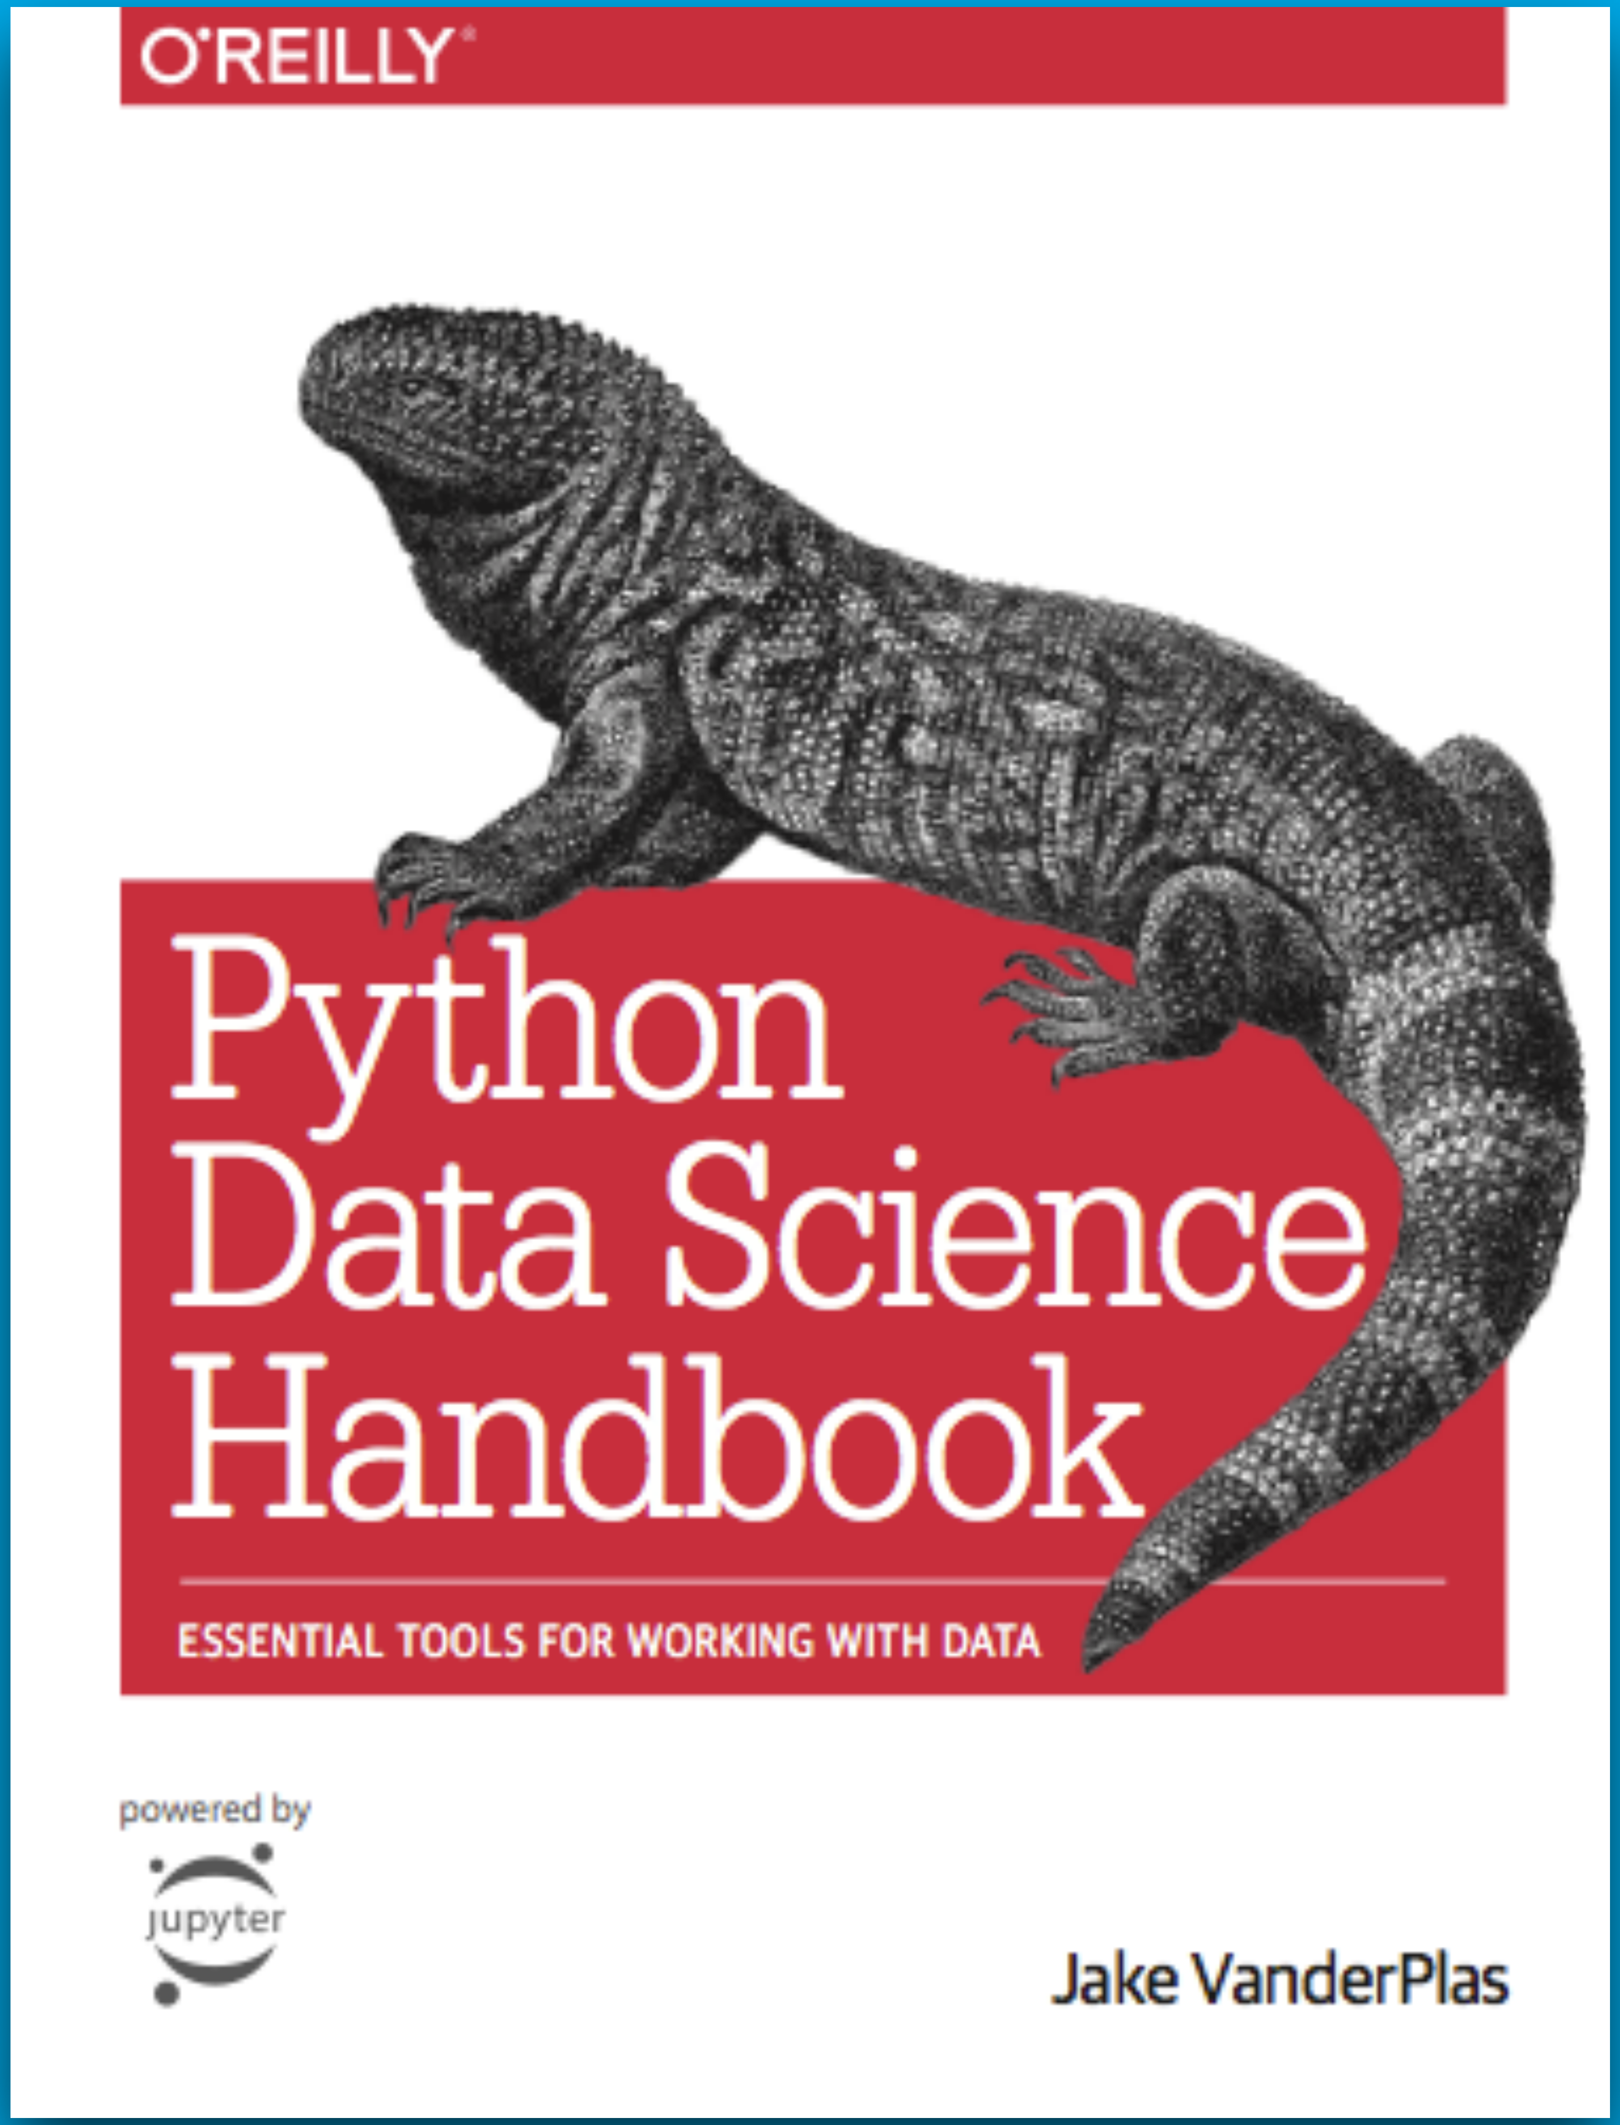

Practical Introduction to ML with Python:

- IPython: Beyond Normal Python
- Introduction to NumPy
- Data Manipulation with Pandas
- Visualization with Matplotlib
- Machine Learning (ca. 180 pages)

Deep Learning

Deep Learning
—Some Background

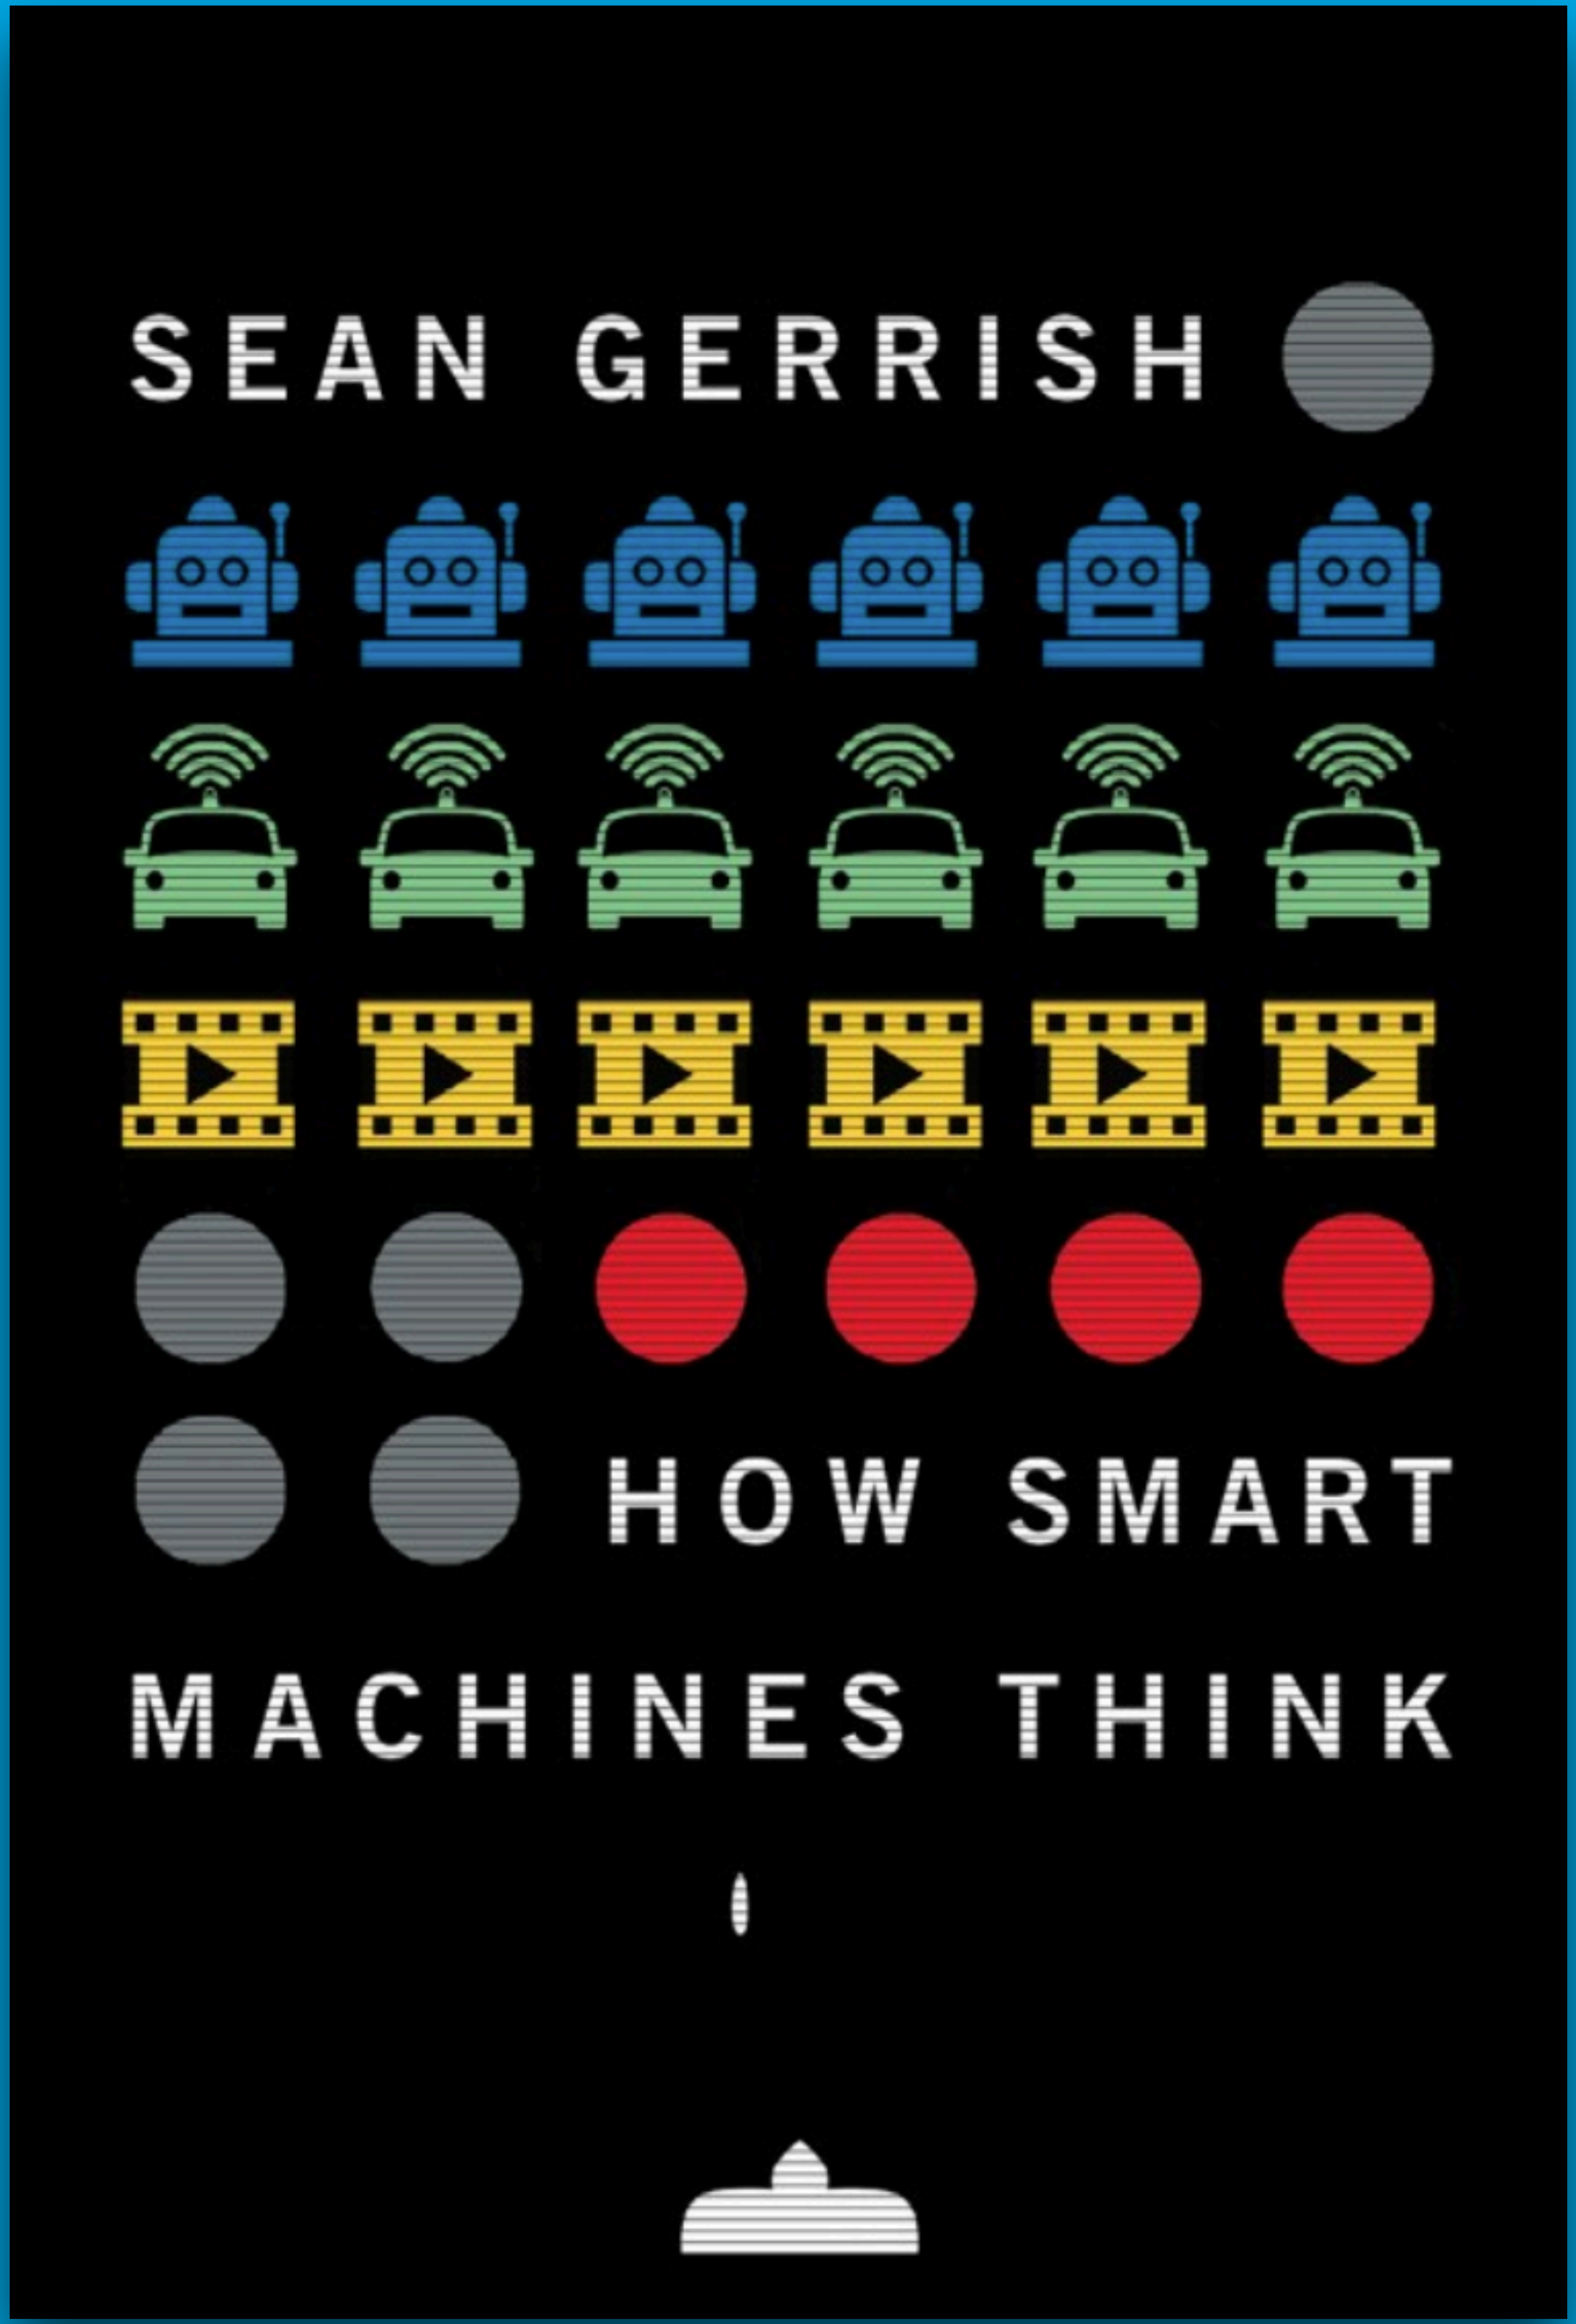

Success Stories about Deep Learning and Deep Reinforcement Learning:

- Self-Driving Cars
- Recommendation Engines
- Playing Atari Games
- Image Recognition & Classification
- Speech Recognition
- Playing the Game of Go

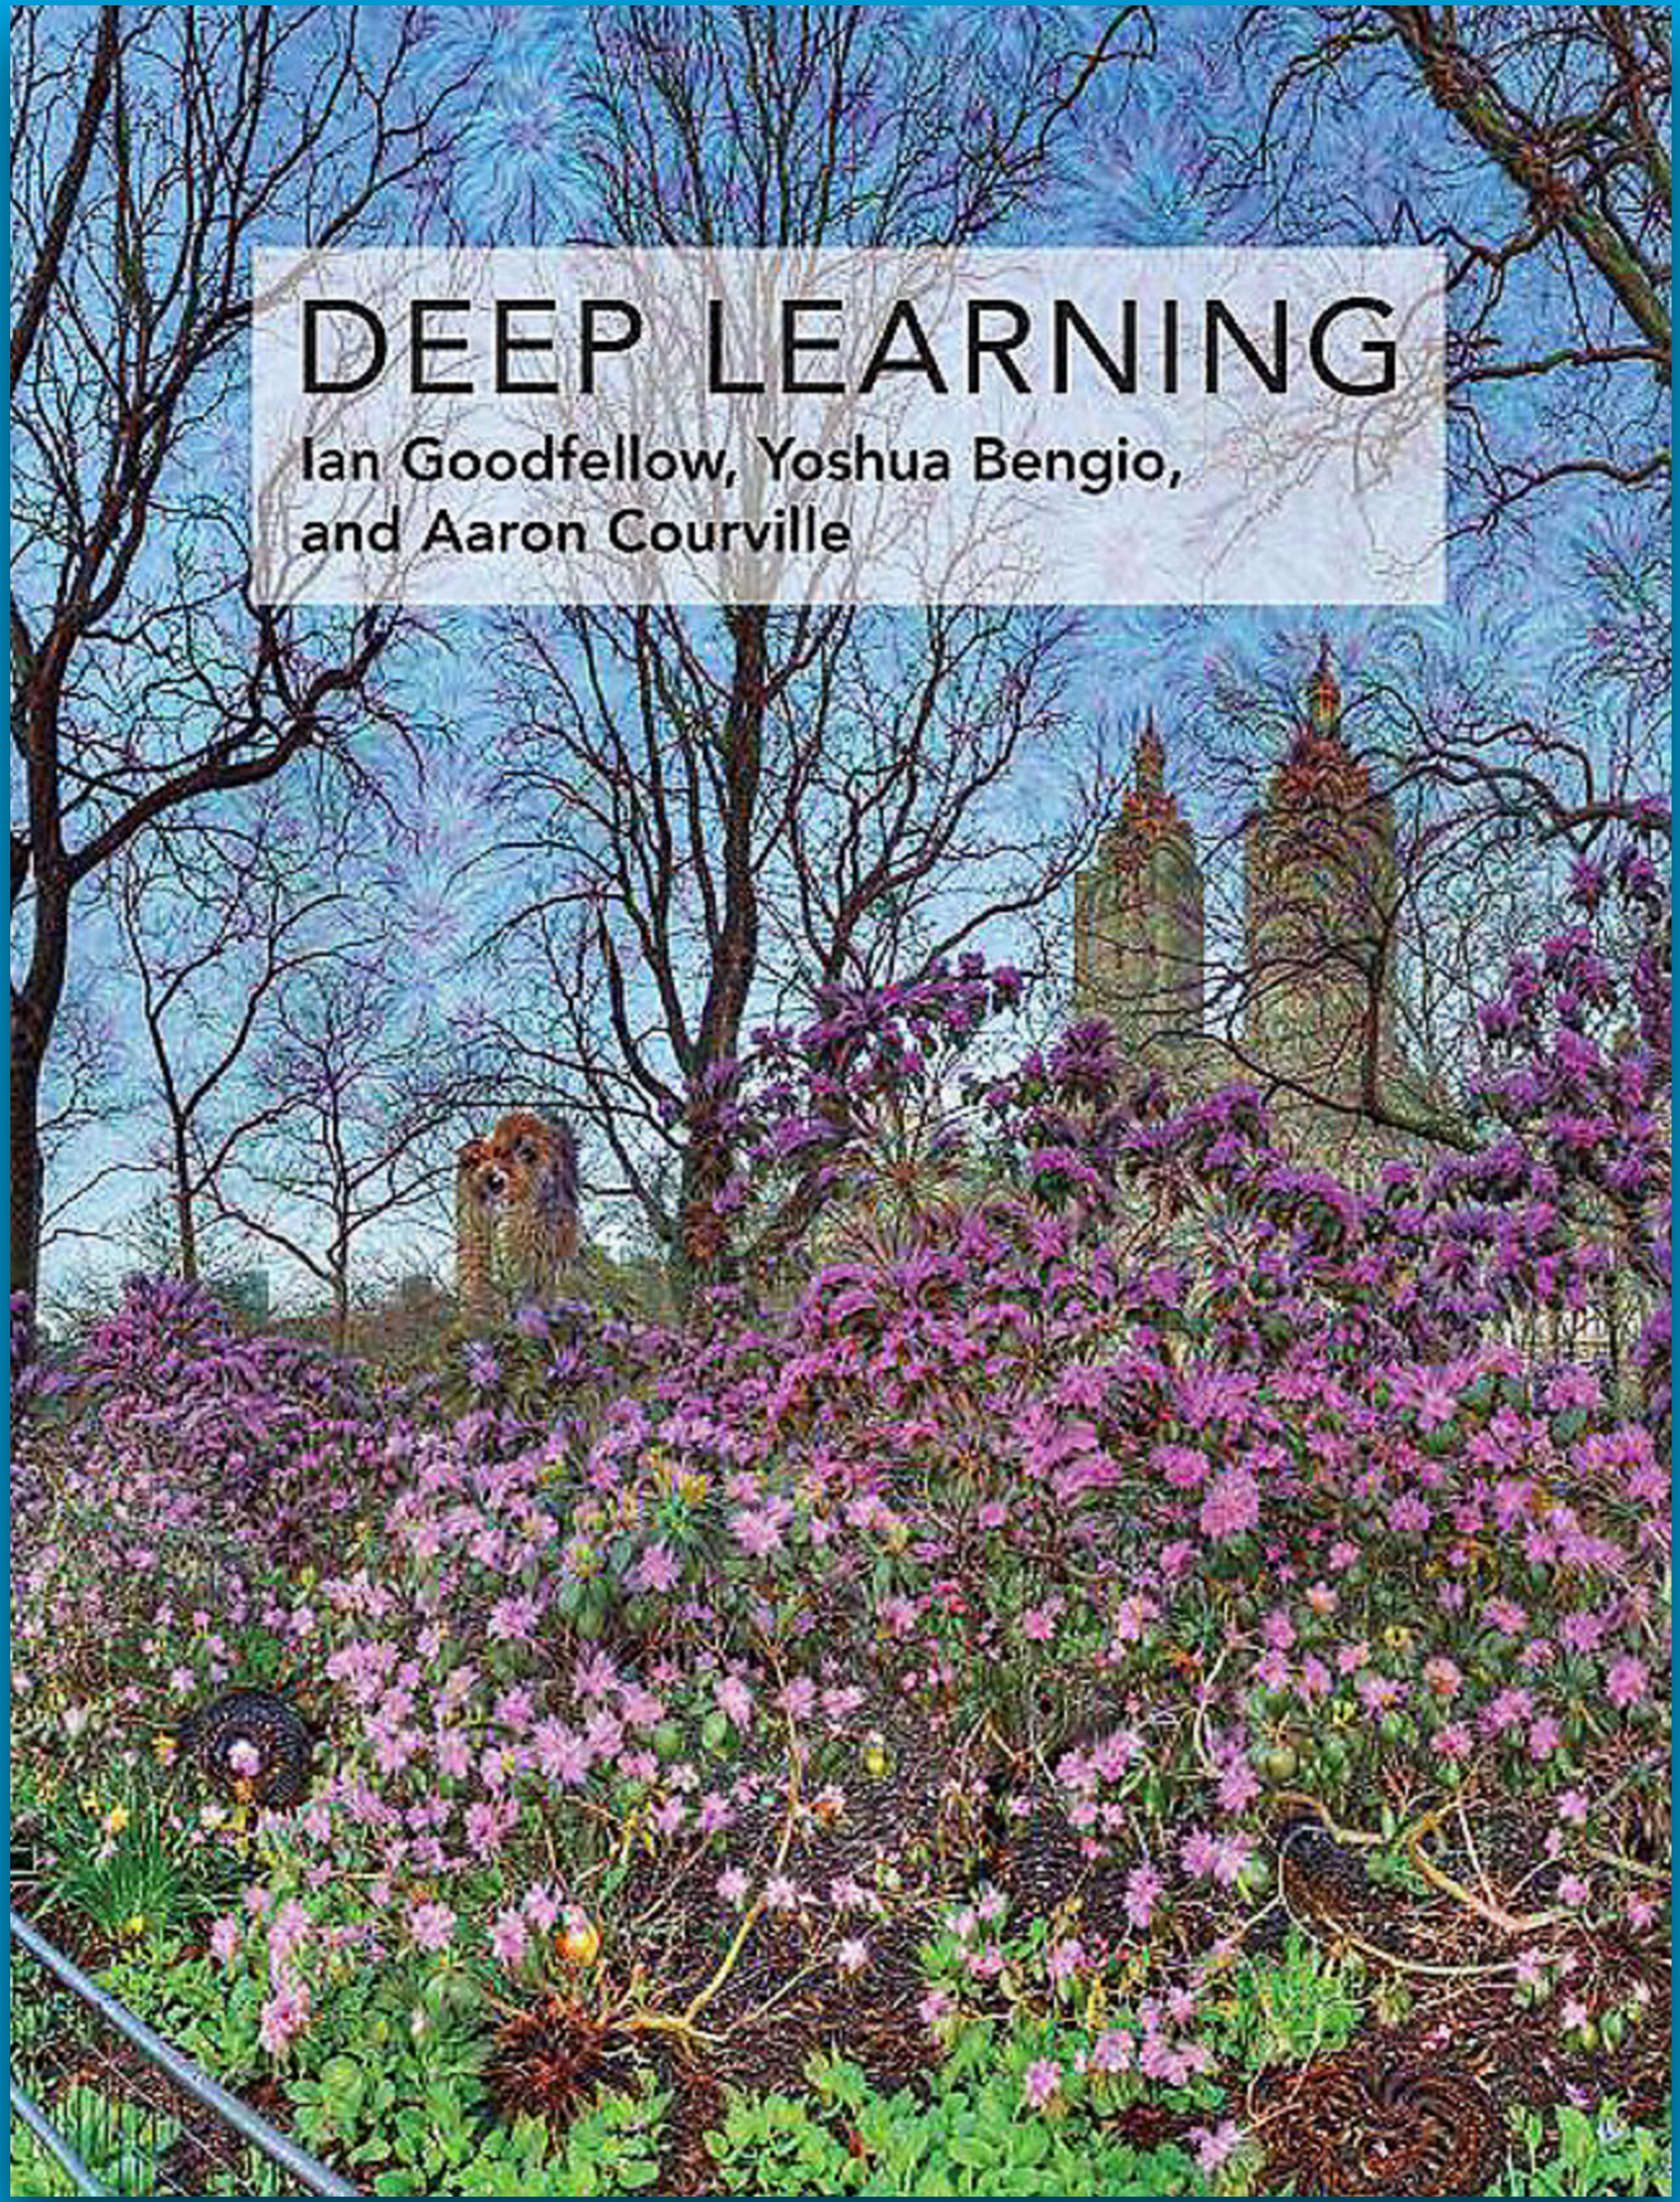

Mathematics of Deep Learning:

- Applied Mathematics
- Machine Learning Basics
- Deep Feedforward Networks
- Regularization for Deep Learning
- Optimization for Training Deep Models
- Convolutional Networks
- Recurrent & Recursive Nets
- Monte Carlo Methods
- …

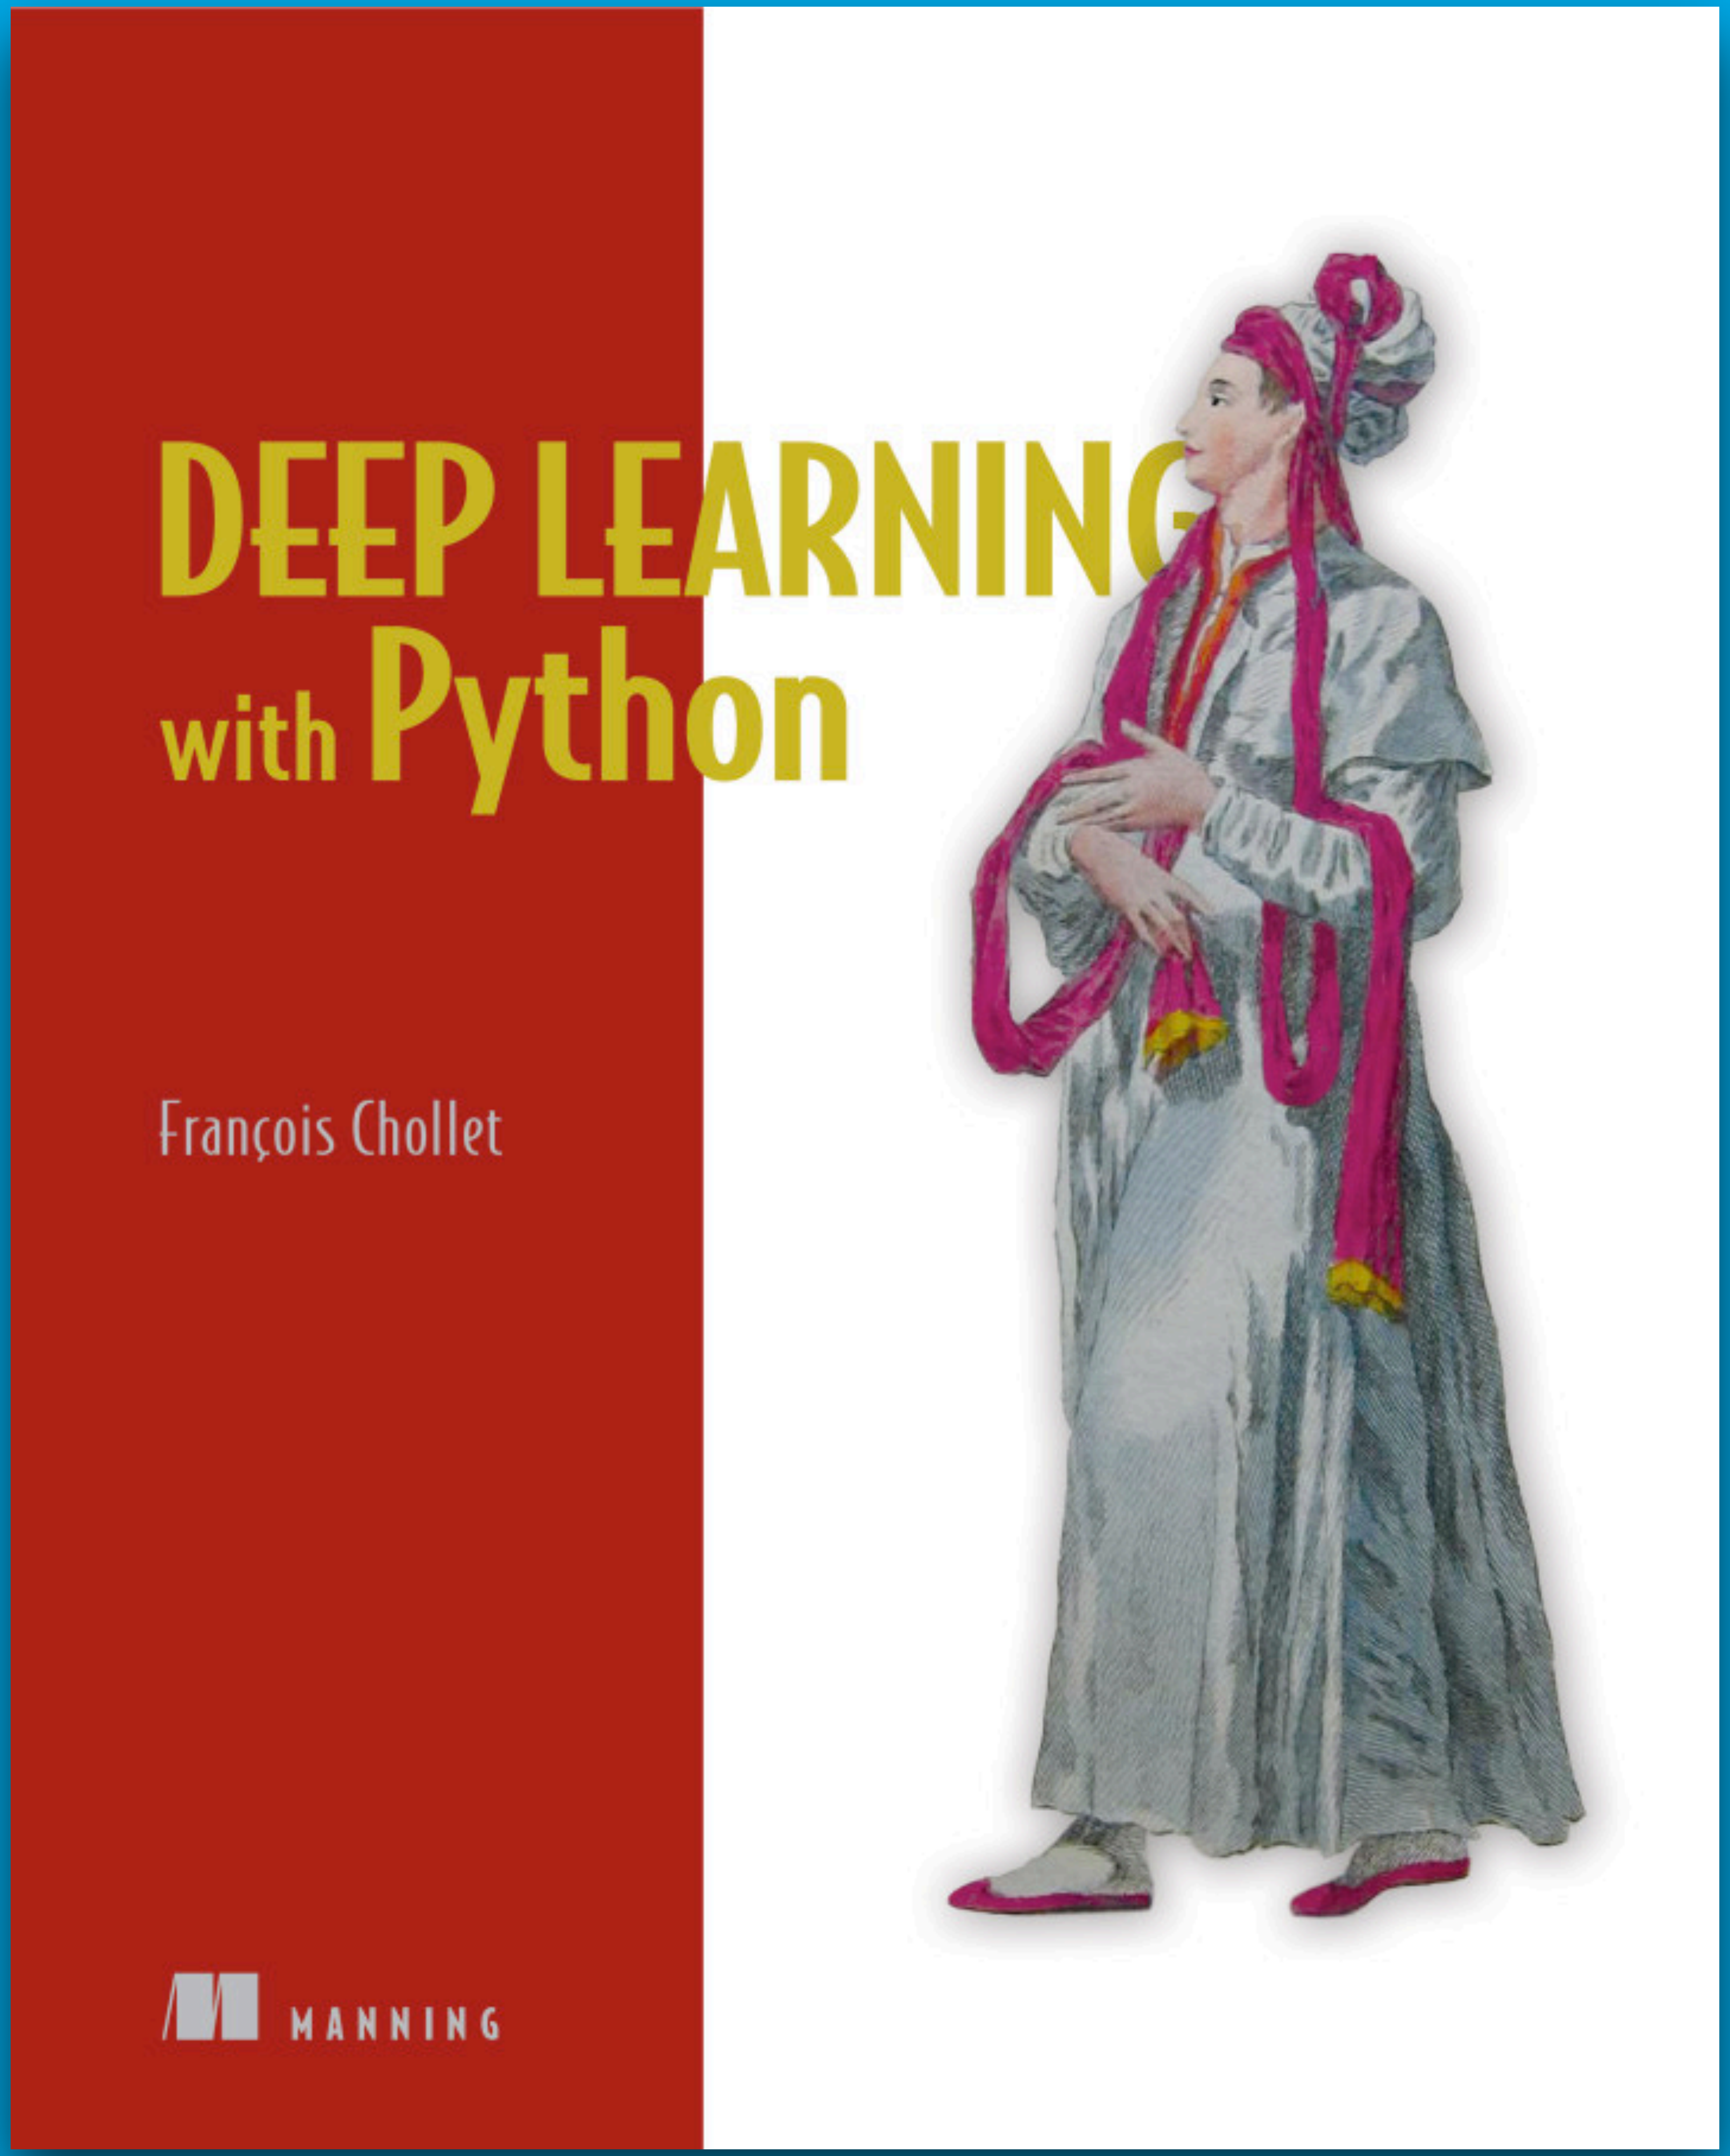

Practice of Deep Learning (with Python and Keras):

- What is Deep Learning?
- Mathematical Building Blocks
- Getting Started with Neural Networks
- Fundamentals of Machine Learning
- Deep Learning for Computer Vision
- Deep Learning for Text and Sequences
- Advanced Deep Learning Best Practices
- Generative Deep Learning

Deep Learning
—Building Blocks

Neural Network
0 Hidden Layers

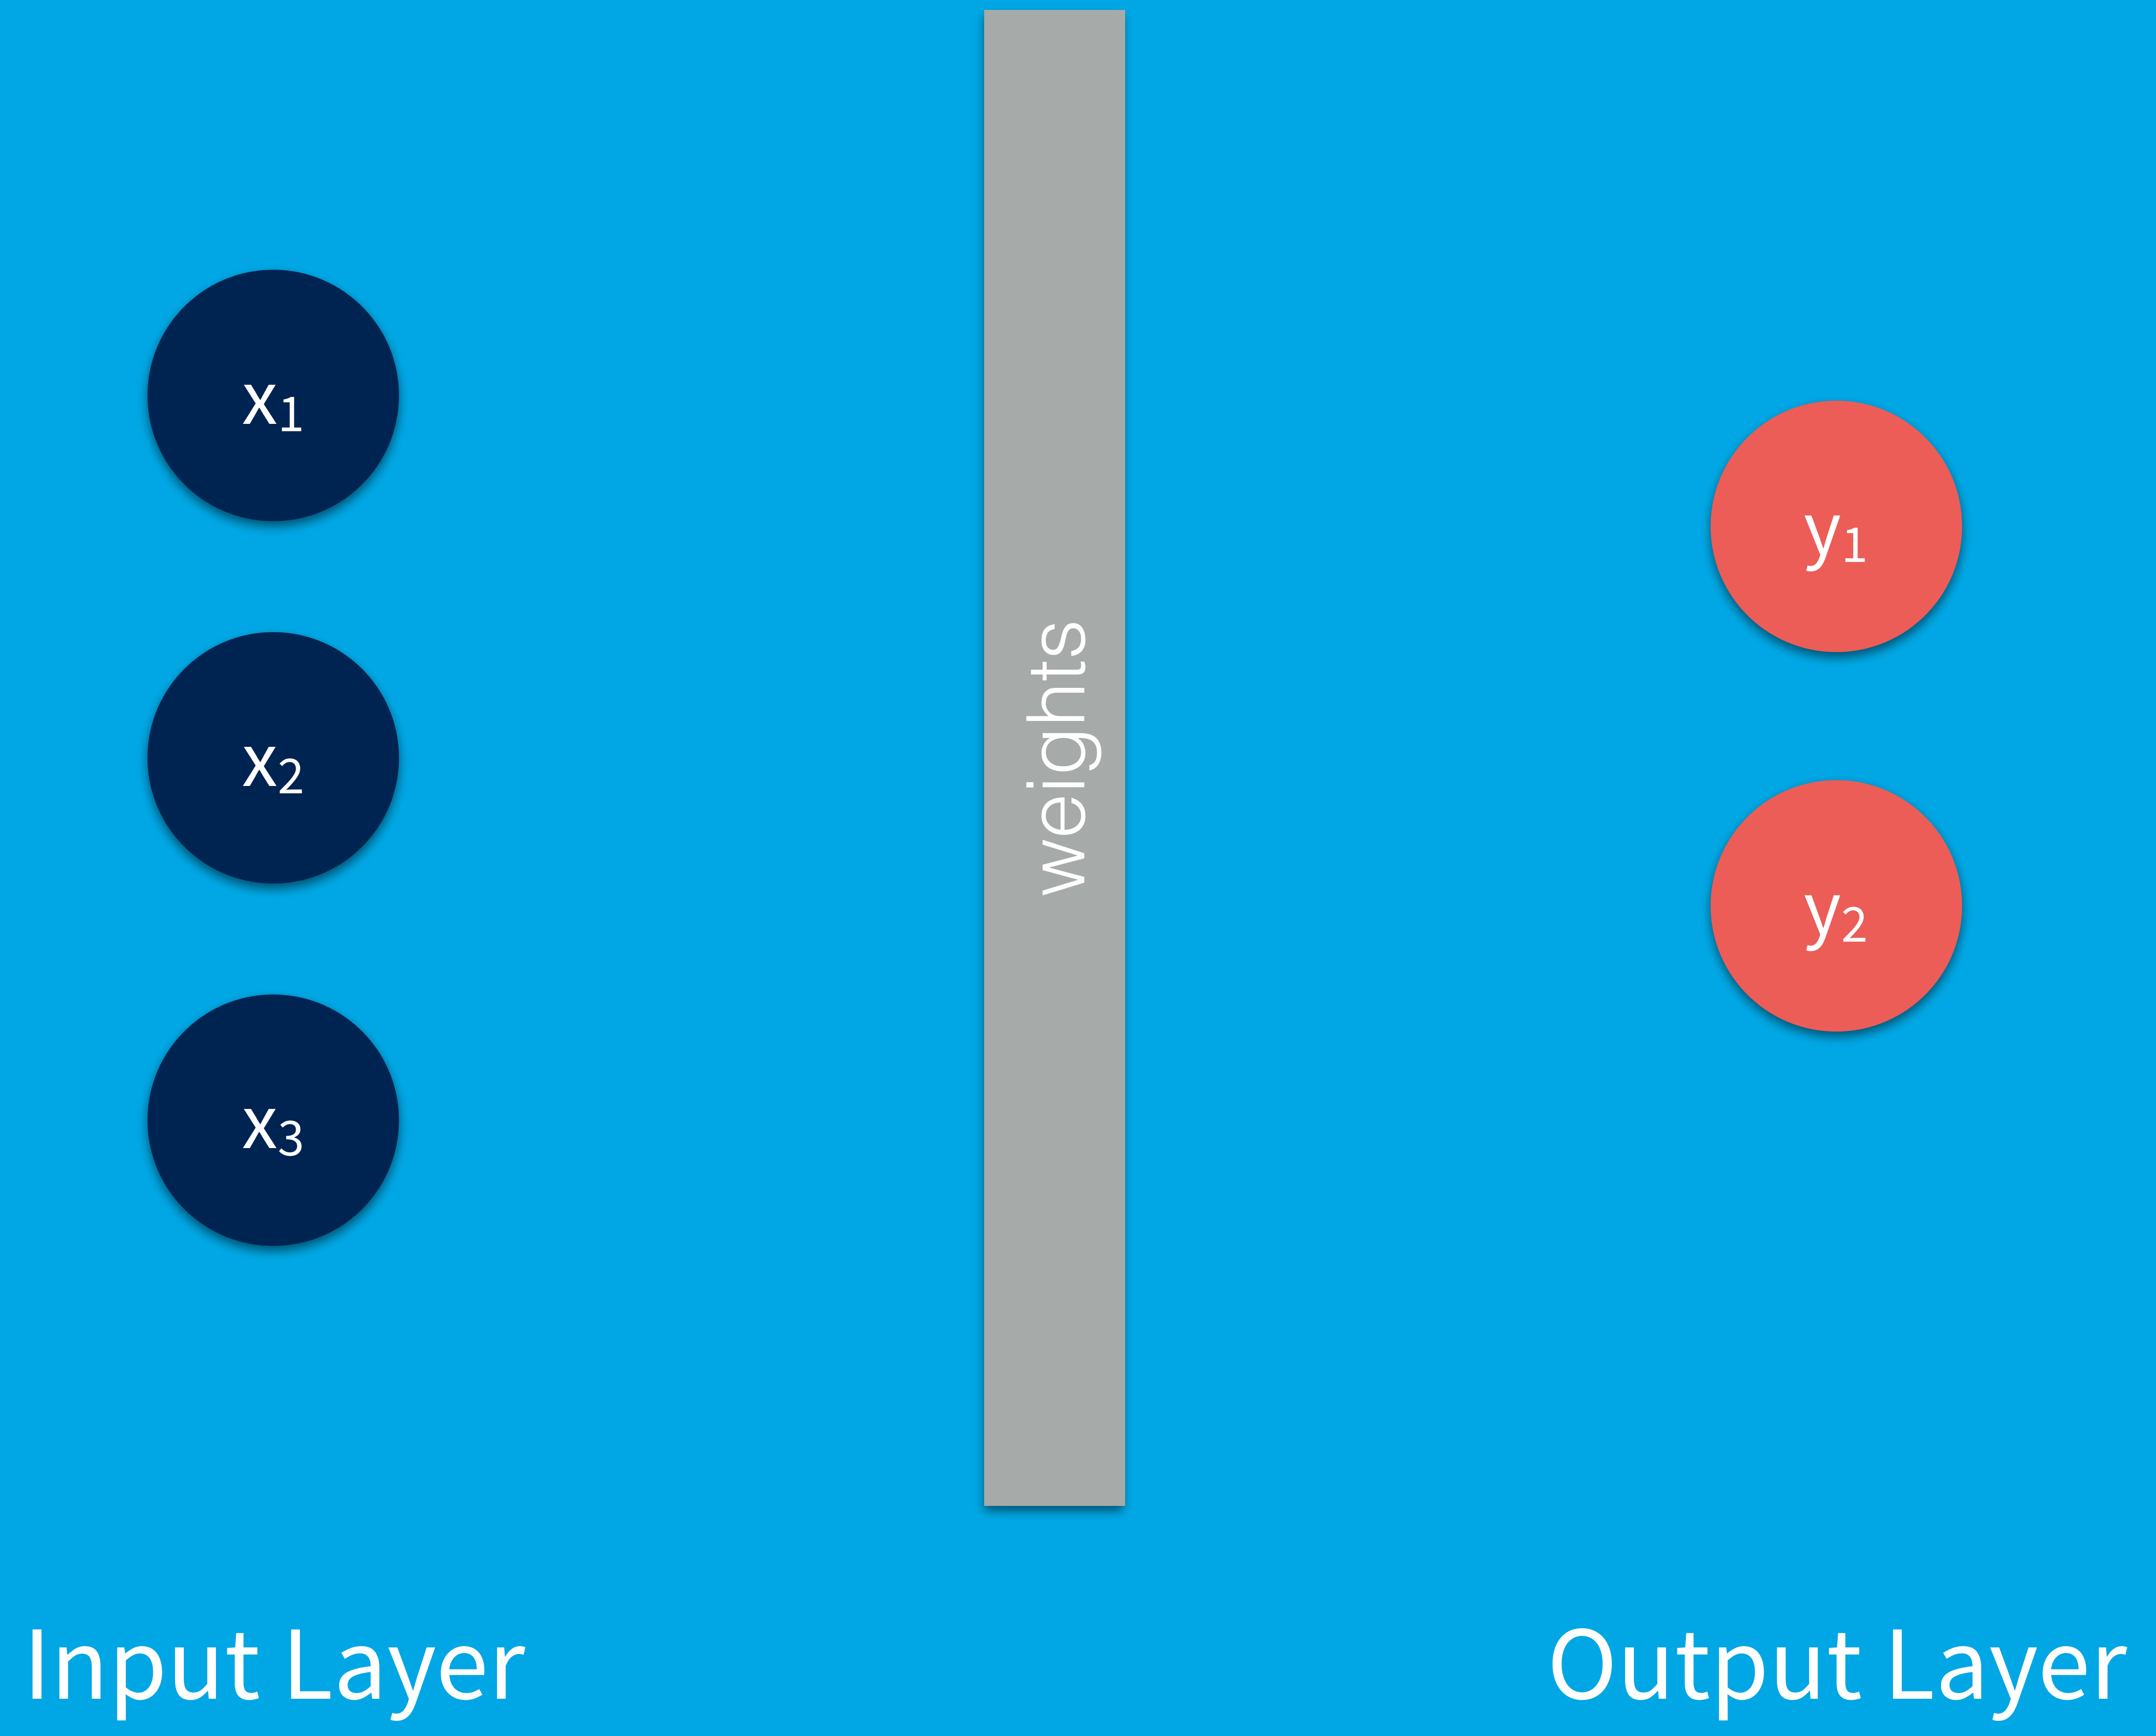

Neural Network
1 Hidden Layer

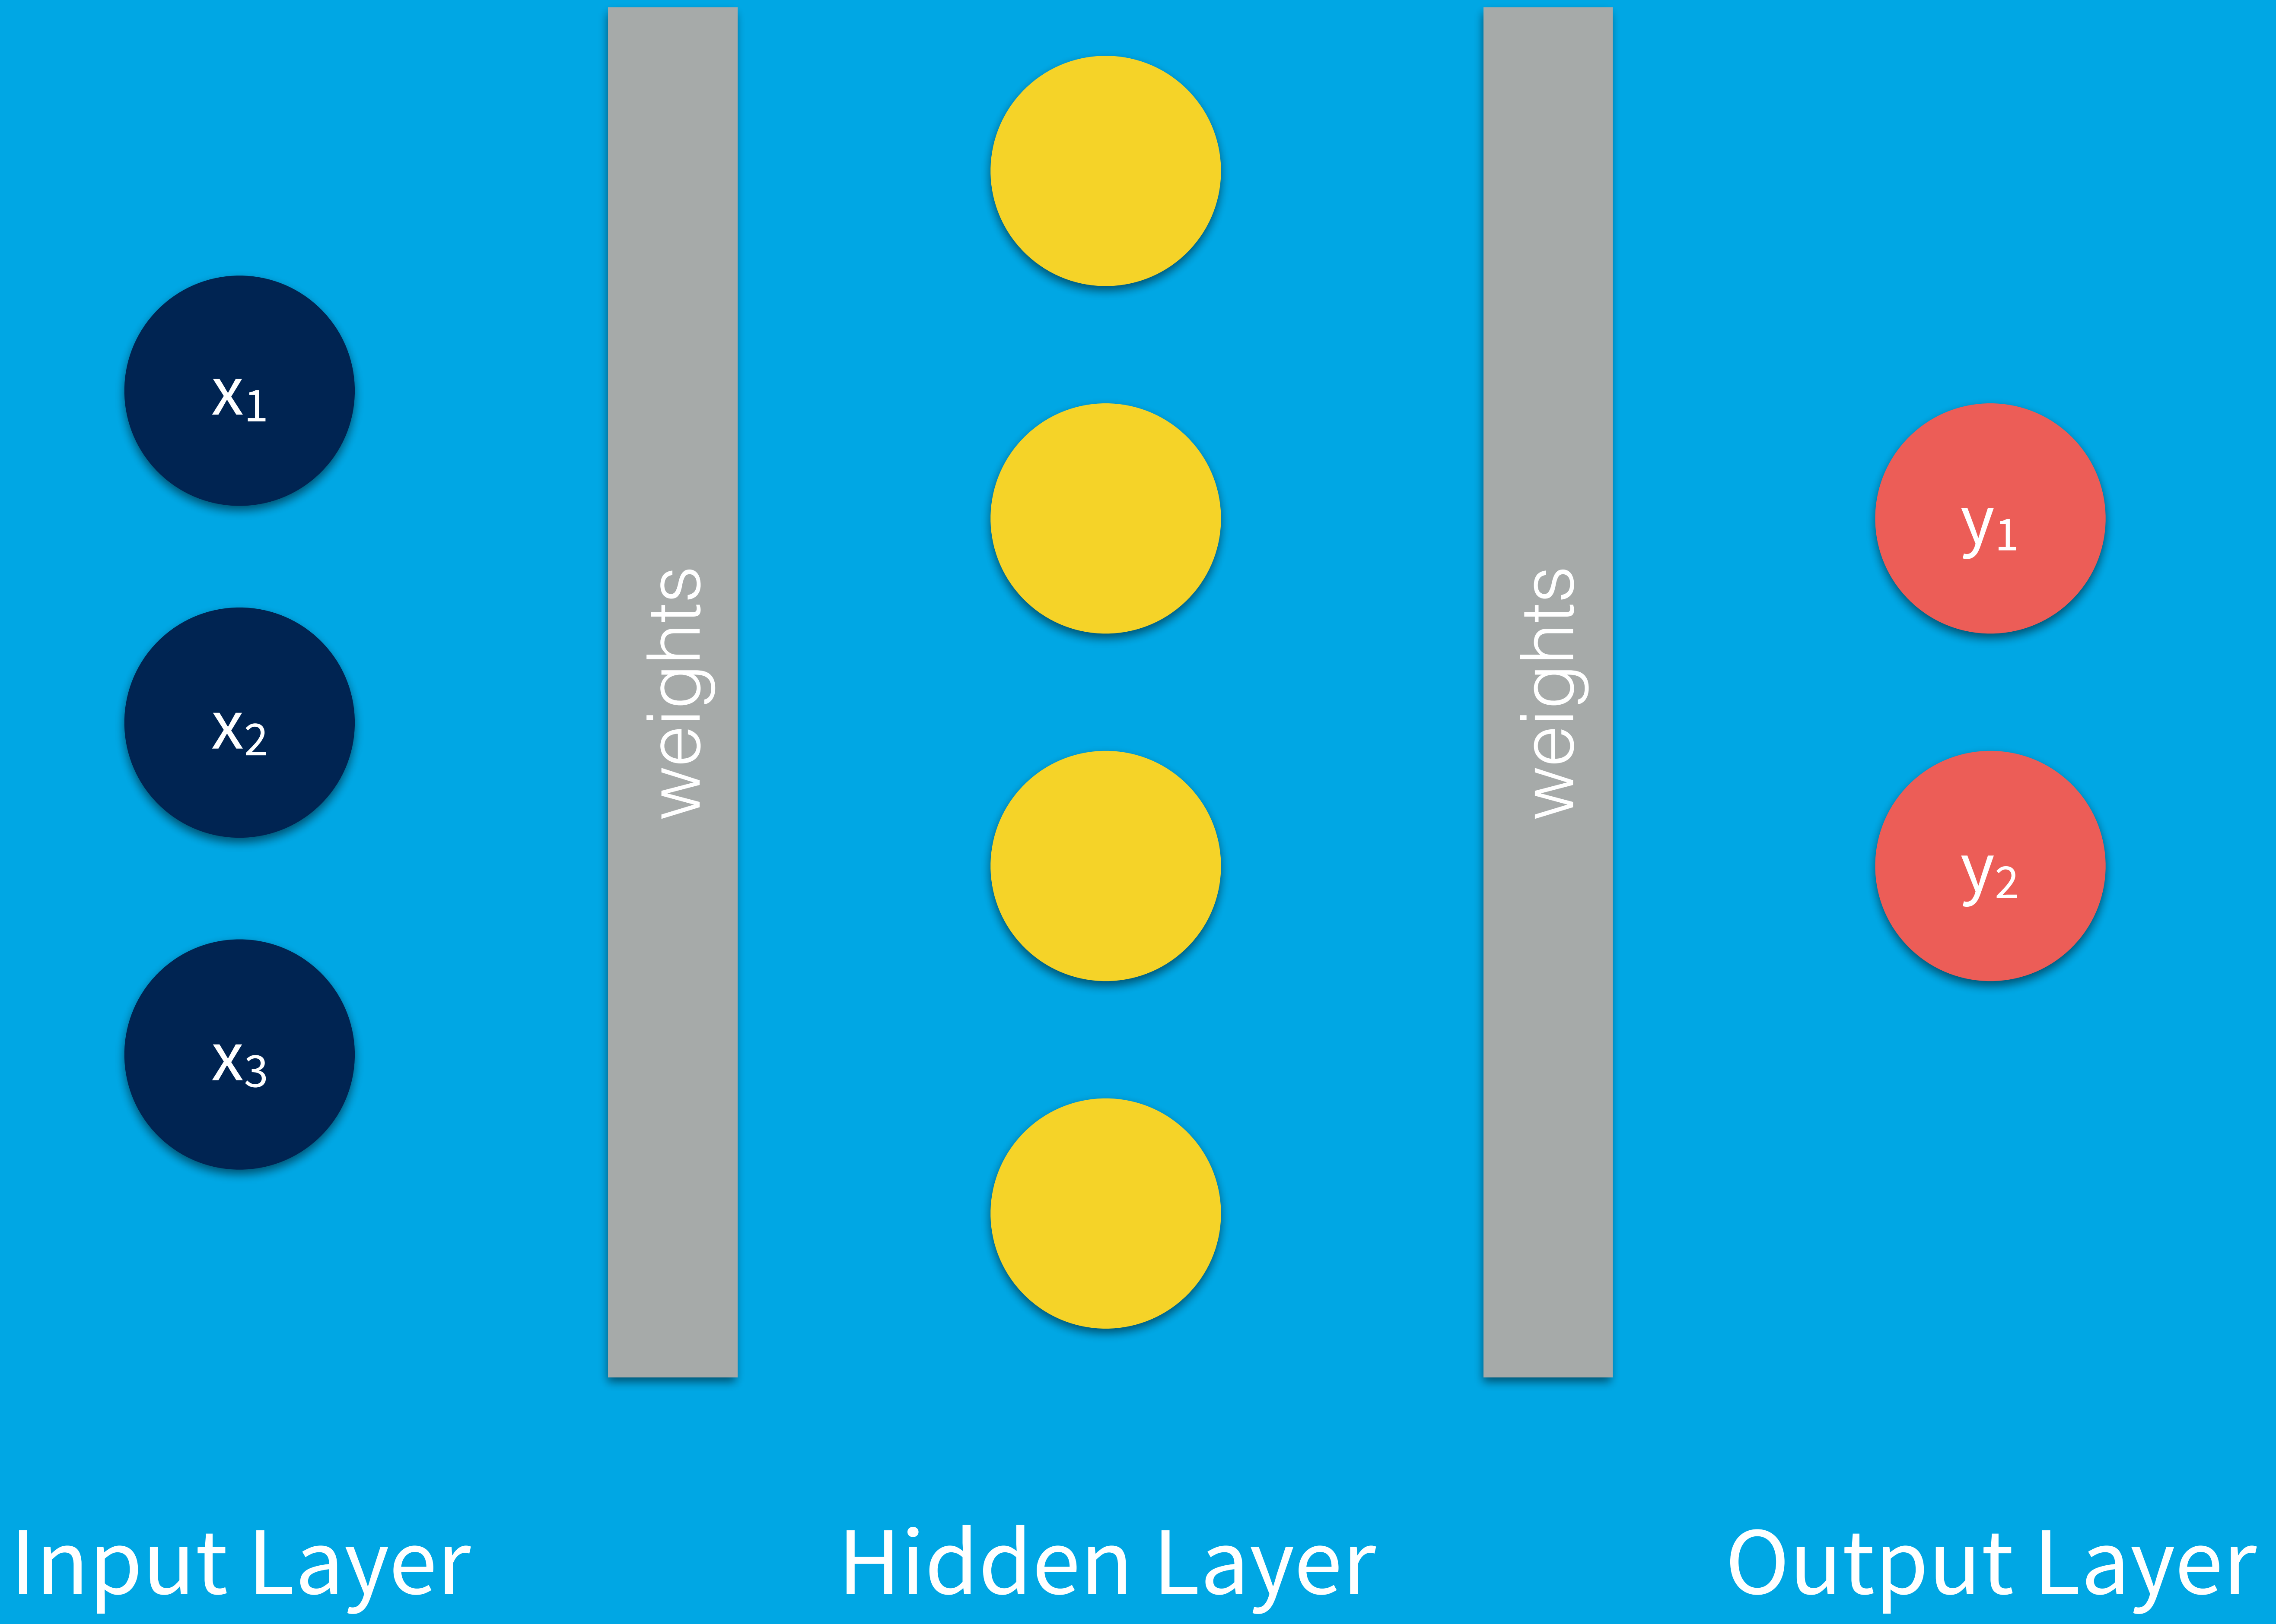

Neural Network
2 Hidden Layers

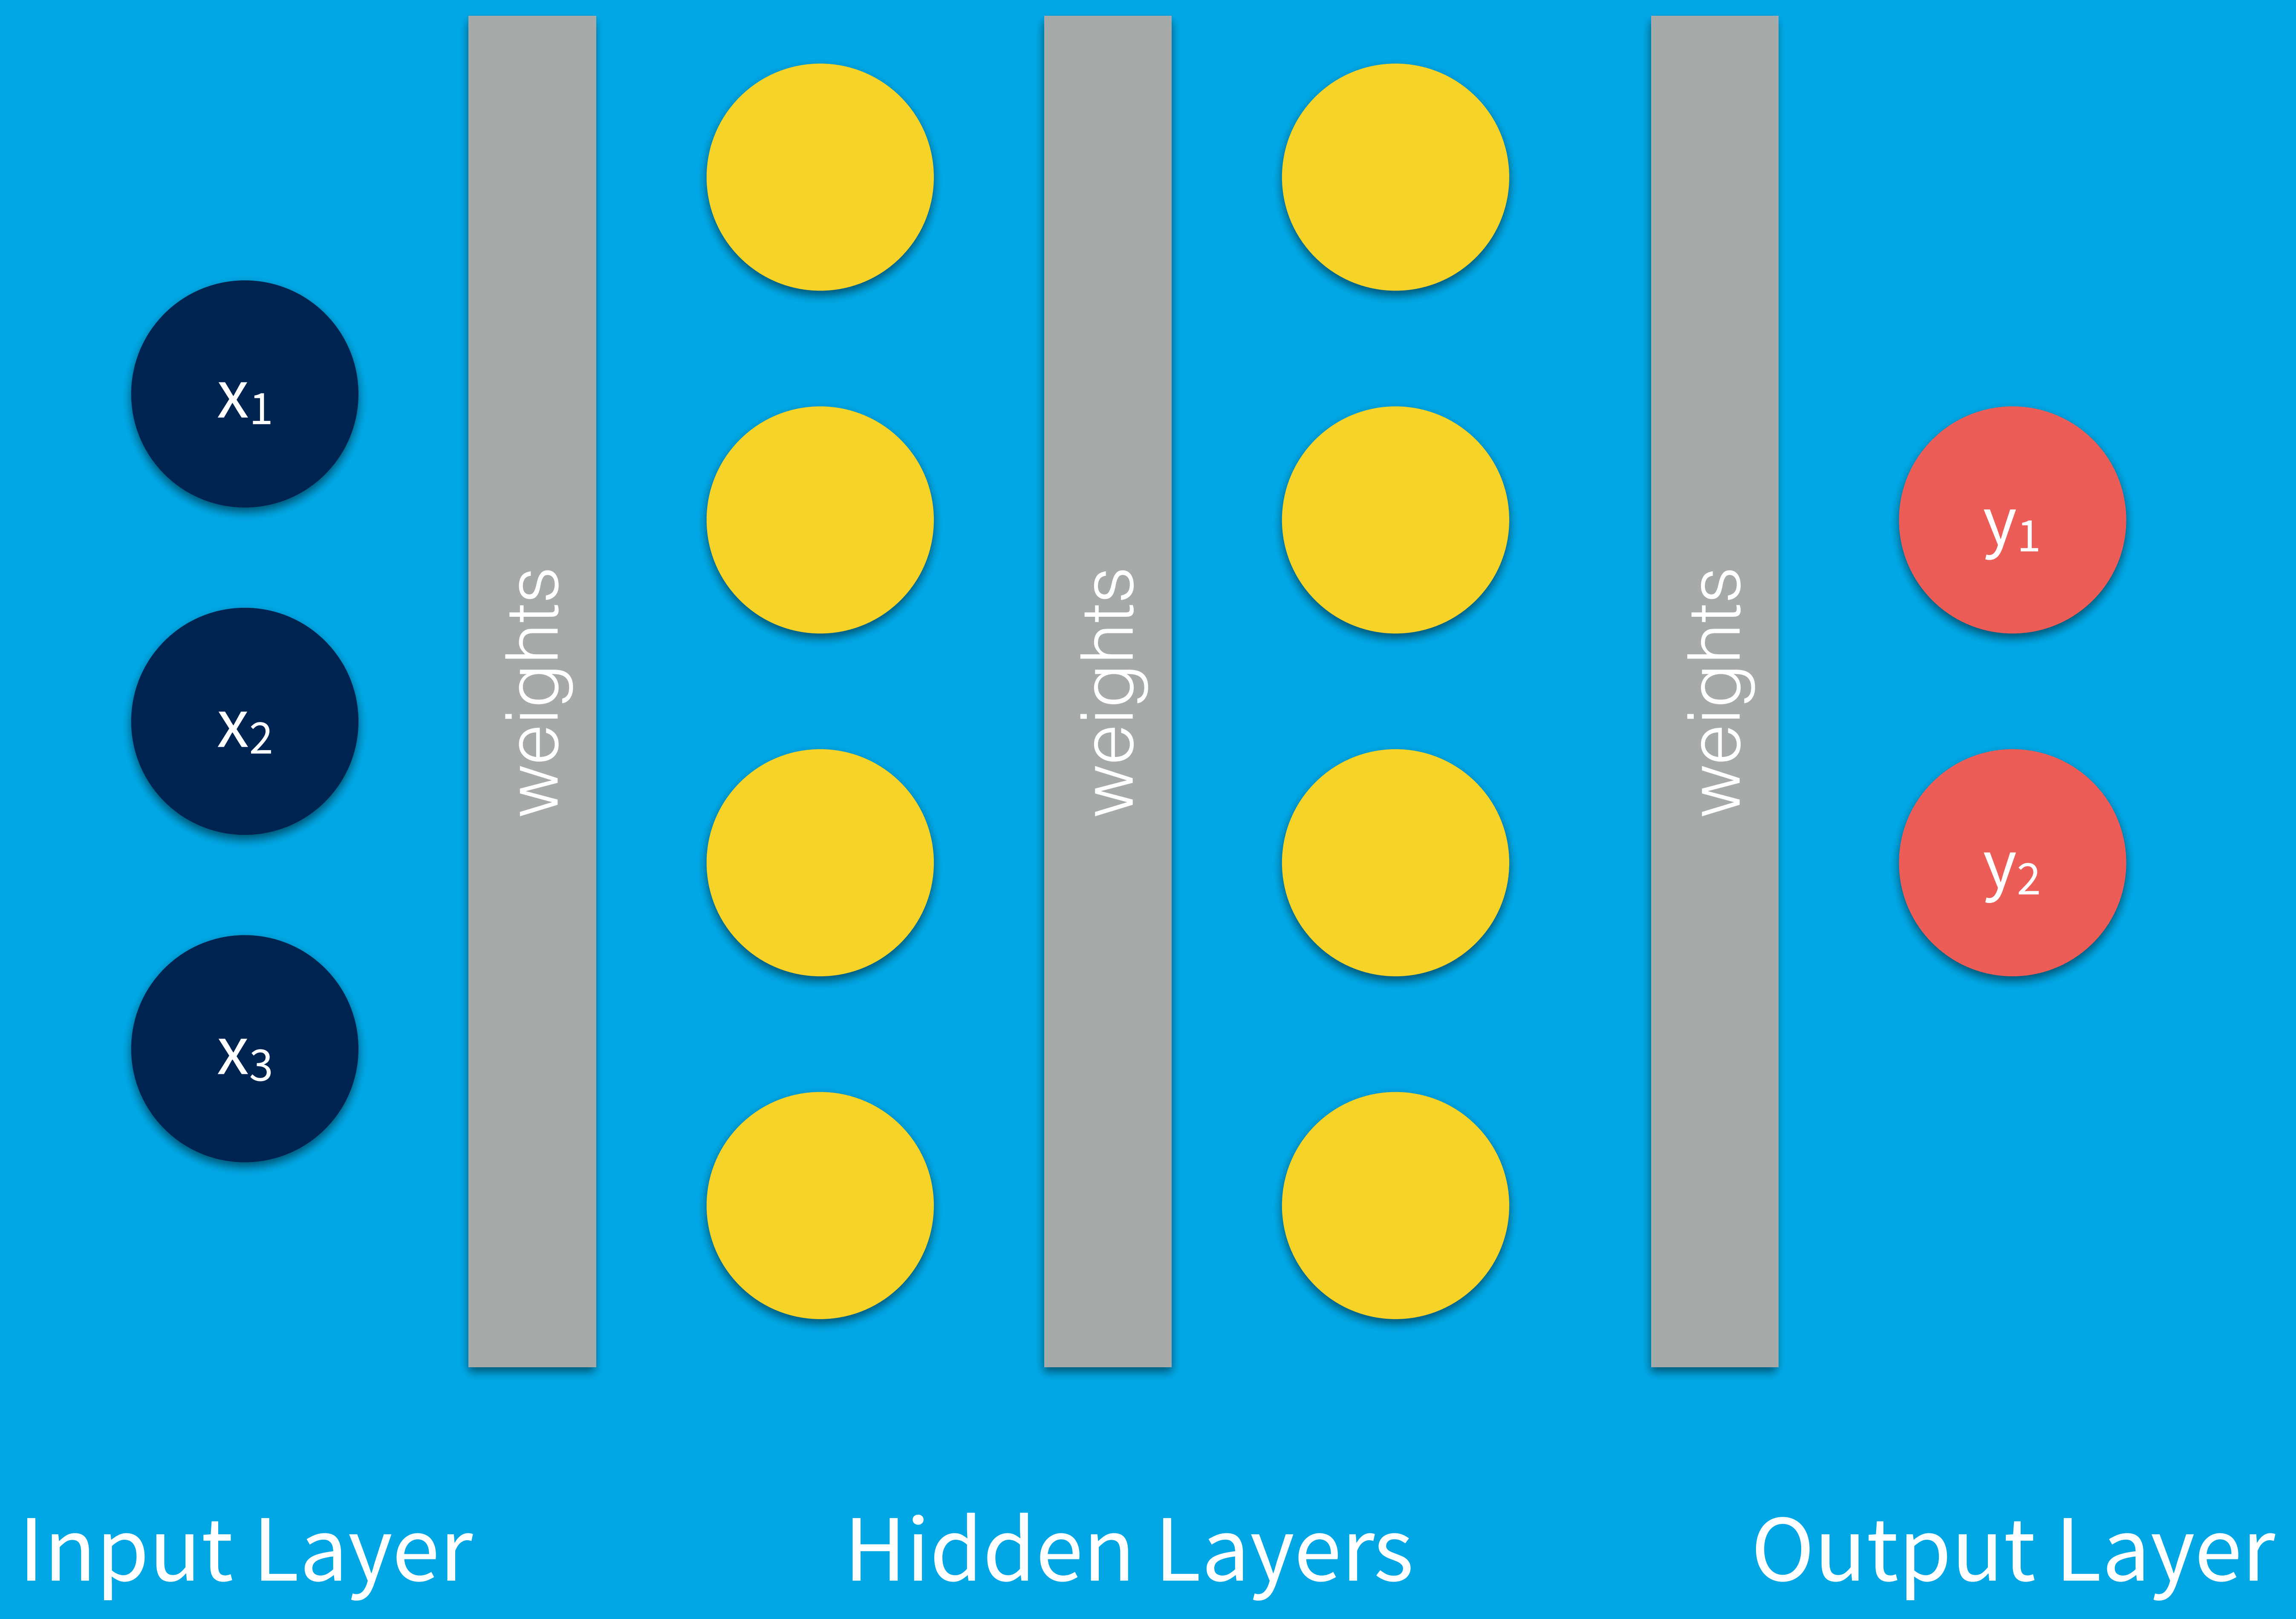

Deep Learning
—Universal Approximation Theorem

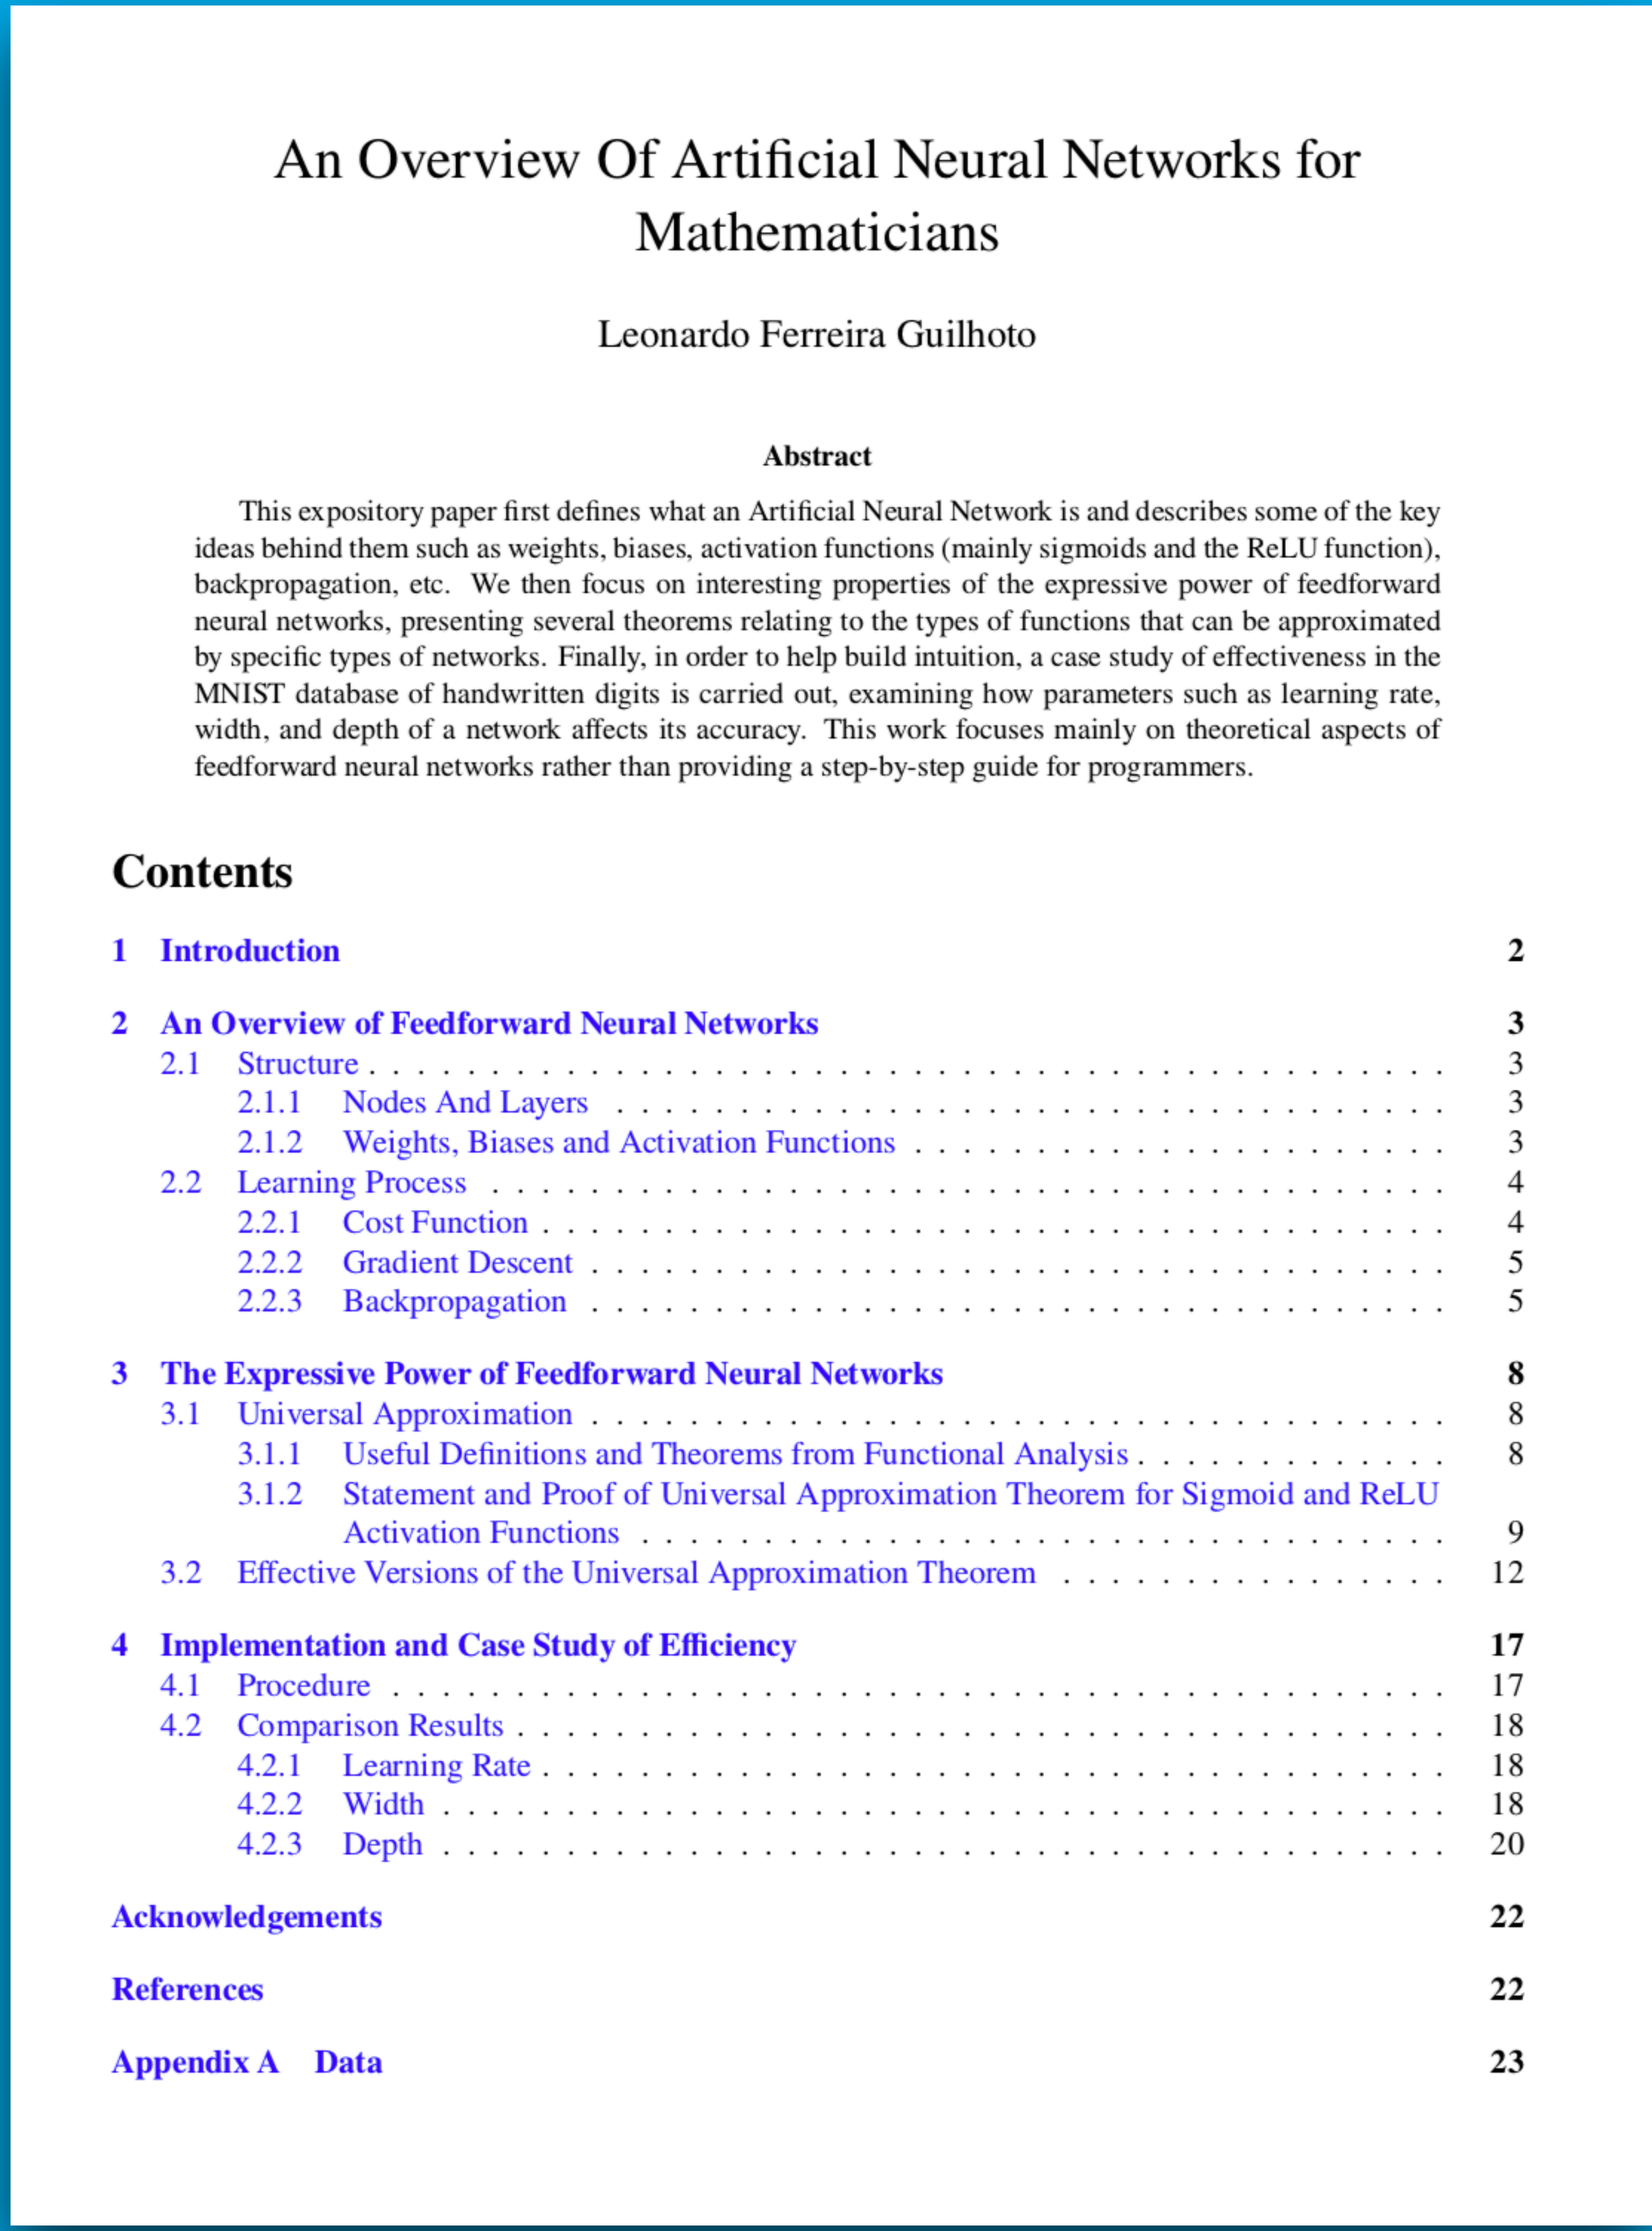

An Overview Of Artificial Neural Networks for Mathematicians

Leonardo Ferreira Guilhoto

Abstract

This expository paper first defines what an Artificial Neural Network is and describes some of the key ideas behind them such as weights, biases, activation functions (mainly sigmoids and the ReLU function), backpropagation, etc. We then focus on interesting properties of the expressive power of feedforward neural networks, presenting several theorems relating to the types of functions that can be approximated by specific types of networks. Finally, in order to help build intuition, a case study of effectiveness in the MNIST database of handwritten digits is carried out, examining how parameters such as learning rate, width, and depth of a network affects its accuracy. This work focuses mainly on theoretical aspects of feedforward neural networks rather than providing a step-by-step guide for programmers.

Contents

"In the mathematical theory of artificial neural networks, the universal approximation theorem states that a feed-forward network with a single hidden layer containing a finite number of neurons can approximate continuous functions on compact subsets of R^n, under mild assumptions on the activation function. The theorem thus states that simple neural networks can represent a wide variety of interesting functions when given appropriate parameters; however, it does not touch upon the algorithmic learnability of those parameters."
—https://en.wikipedia.org/wiki/Universal_approximation_theorem

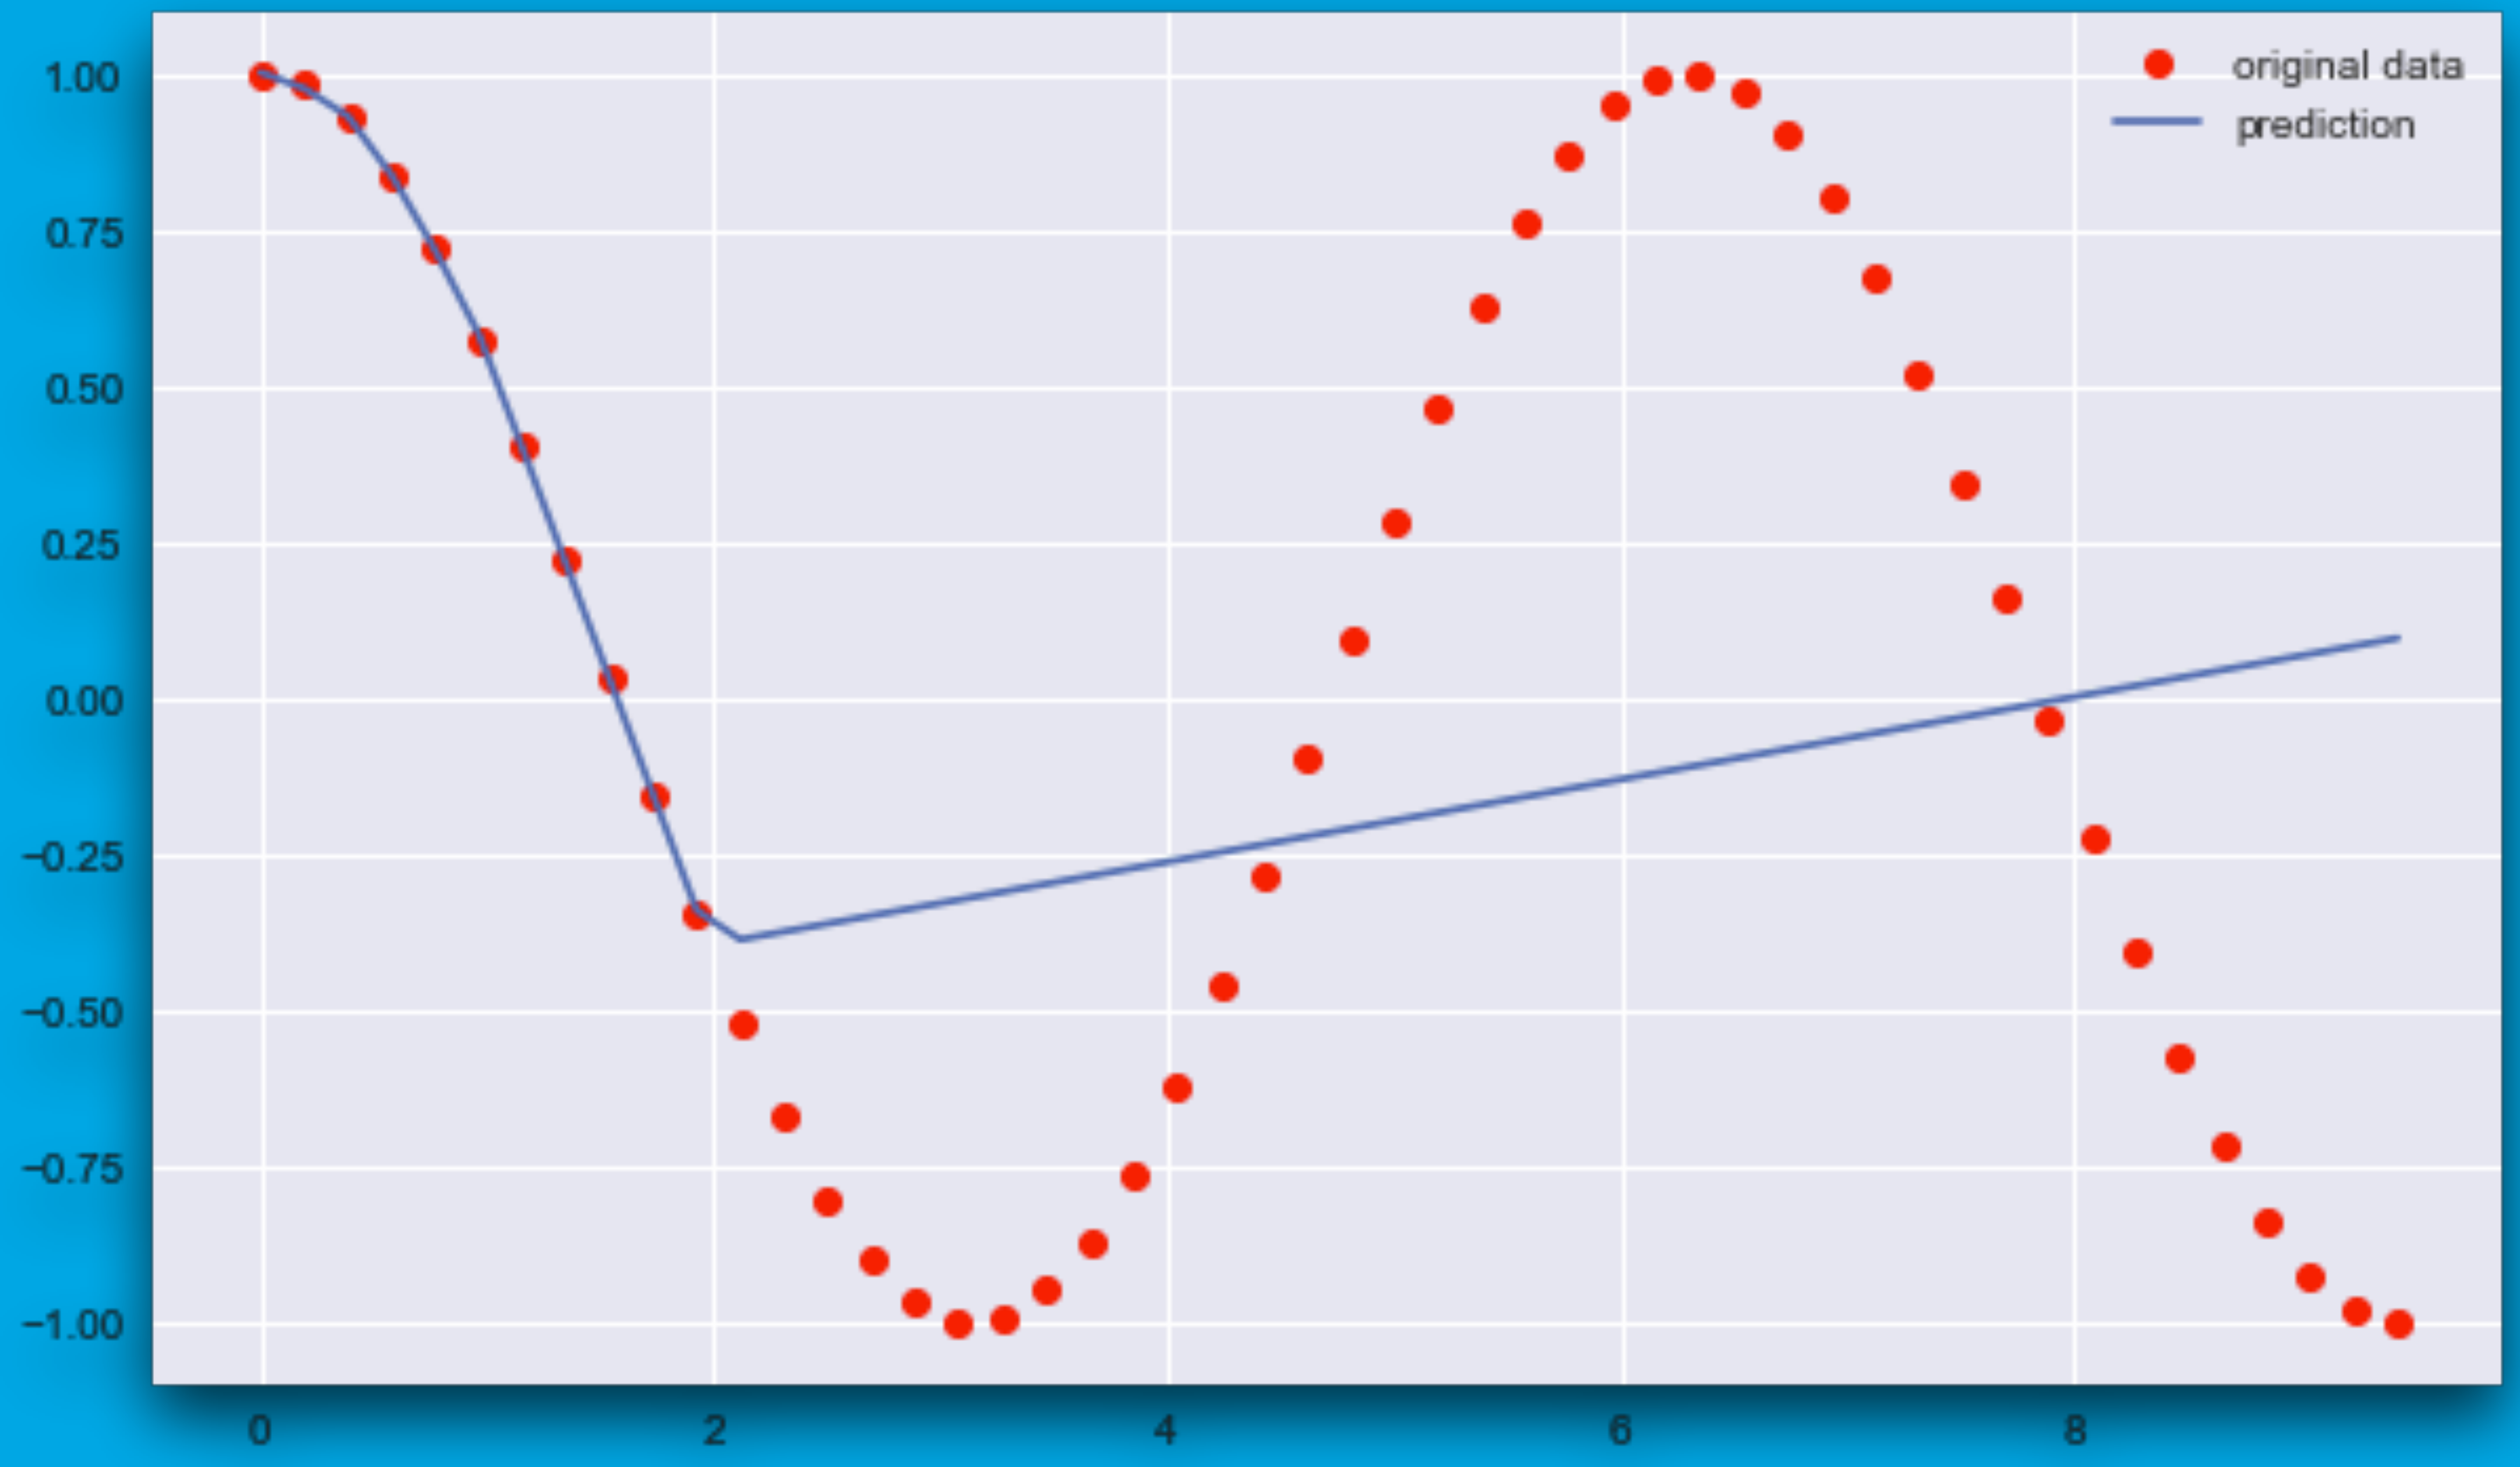

First Illustration with Keras & Tensorflow

Market Prediction

Market Prediction —Scikit-Learn

With Scikit-Learn there are Deep Neural Network (Multi Layer Perceptron, MLP) models available both for estimation ...

```
from sklearn.neural_network import MLPRegressor
model = MLPRegressor(hidden_layer_sizes=1 * [1024,],
                     activation='relu', solver='adam',
                     learning_rate_init=0.001,
                    nesterovs_momentum=False,
                     shuffle=False, max_iter=10000
                    validation_fraction=0.1)
model.fit(x, y)
pred = model.predict(x)
```

… and classification.

```
from sklearn.neural_network import MLPClassifier
model = MLPClassifier(hidden_layer_sizes=1 * [1024,],
                      activation='sigmoid', solver='adam',
                      learning_rate_init=0.001,
                      nesterovs_momentum=False,
                      shuffle=False, max_iter=10000
                     validation_fraction=0.1)
model.fit(x, y)
pred = model.predict(x)
```

Market Prediction
—Keras

Keras, with e.g. TensorFlow as its backend, allows the sequential building of Deep Neural Networks.

```
from keras.layers import Dense
from keras.models import Sequential
from keras.optimizers import Adam

model = Sequential()
model.add(Dense(128, input_dim=1, activation='relu'))
model.add(Dense(48, activation='relu'))
model.add(Dense(1, activation='linear'))   # estimation
# model.add(Dense(1, activation='sigmoid')) # classification
adam = Adam(lr=0.001, beta_1=0.9, beta_2=0.999,
            epsilon=None, decay=0.0, amsgrad=False)
model.compile(loss='mse', optimizer=adam,
              metrics=['mse', 'accuracy'])

model.fit(x, y, epochs=2000, verbose=False)
pred = model.predict(x)
```

Conclusions

1. Finance has long been driven by the **“beauty myth”** — elegant but too simplistic models, equations and approaches.
2. The availability of **big financial data** (historical—streaming, structured—unstructured) gave rise to data-driven finance.
3. It might be assumed that the **“unreasonable effectiveness of big data”** holds true in the financial domain as well.
4. Due to the availability of big data (e.g. billions of hours of virtual car driving), **Artificial Intelligence** (AI) is changing almost every area of our lives.
5. It is to be assumed that in the same way the **combination of data-driven and AI-first finance** will change the field for good.

1. **Deep Learning** approaches “make us hopeful” that we can overcome the main corollary of the Efficient Markets Hypothesis, i.e. that the analysis of historical data is useless (for the creation of alpha).
2. Furthermore, there are **alternative algorithms** available that might also be useful (better) in predicting market movements:
   A. recurrent neural networks
   B. convolutional neural networks
   C. deep reinforcement learning

1. However, so far we have ***only*** considered the **prediction part of algorithmic trading** (i.e. the signal generation).
2. Two important topics have been left out:
   A. **market microstructure** elements (e.g. transaction costs) have not been considered in any meaningful way.
   B. In addition, **execution rules** play an important role (sizing, resizing, stop loss, profit capture, etc.) for the trading performance.

After all, working with AI algorithms — based on Python — and applying them to financial problems is fun, intellectually stimulating and might finally lead to the “holy grail” of finance:

Being able to consistently outperform others and the markets.

This naturally raises questions regarding the future of the finance domain, the eduction of people working in it, the ways companies compete in the field and also regarding ethics and governance.

The Python Quants GmbH

Dr. Yves J. Hilpisch
+49 3212 112 9194
tpq.io | ai@tpq.io | @dyjh