Dawn of Behavioural Finance, 1688

Dawn of Behavioural Finance, 1688

By

Vasileios Kallinterakis

Cambridge Scholars Publishing

Dawn of Behavioural Finance, 1688

By Vasileios Kallinterakis

This book first published 2024

Cambridge Scholars Publishing

Lady Stephenson Library, Newcastle upon Tyne, NE6 2PA, UK

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN: 978-1-0364-1340-8
ISBN (Ebook): 978-1-0364-1341-5

To my beloved Parents, who taught me to think for myself

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

PREFACE

Behavioural finance was born during the 1980s as the result of a cross-fertilisation between finance and (primarily) psychology, with other disciplines (including sociology and decision science) also contributing to its foundations. The role of psychology in financial decisions was largely downplayed during much of the 20th century, with finance scholars mainly relying on the neoclassical paradigm; drawing on the concept of the Economic Man, this paradigm involved a more mathematical representation of decision-making, aiming at calibrating financial reality via quantitative models. To the extent that this paradigm ignored the human nature of economic agents, it gradually began to be challenged since the 1950s (when Herbert Simon proposed the notion of bounded rationality) and, more so, since the 1970s (when Daniel Kahneman and Amos Tversky denoted the role of prospect theory and biases/heuristics in decision-making). From then on, the trajectory of behavioural finance is, more or less, known; initially challenging market efficiency (a cornerstone of neoclassical finance research), researchers began exploring deviations from efficient pricing (in the form e.g., of return anomalies and predictability), before availing themselves of datasets of more depth (such as e.g., higher frequency data and transaction data/portfolio holdings of retail/institutional investors) to examine various behavioural trading patterns (e.g., herding, disposition effect etc.).

Although the above description of the evolution of behavioural finance is largely accurate, it suggests that the role of psychology in financial decisions was only appreciated in the postwar decades–something not accurate. Evidence from prewar scholarship suggests that the role of psychology in economics, in general, was raised in a series of treatises going as far back as the 18th century; key names in economics, such as Adam Smith, Thorsten Veblen and Jeremy Bentham touched upon behavioural issues in economic life (Camerer et al., 2004; Zak, 2004; Sapra and Zak, 2008), while such issues were even raised by authors of literature aimed at non-academic audiences (Mackay, 1841; Selden, 1919; Clason, 1926). However, it is important to note that none of those authors essentially came close enough to the spirit of modern-day behavioural finance; this was accomplished in the Netherlands long before their time, when Amsterdam hosted the world's first-ever organised stock exchange.

In the year 1688, the Sephardi merchant Joseph de la Vega published a book aiming at describing aspects of the investment life in the then-nascent Dutch stock market. Titled *Confusion of Confusions*, the book analysed technical aspects of trading operations, besides offering a discussion of various forms of investors' behaviour. What is astounding–as shall be depicted in ample detail later on in this book–is that the issues Vega discussed and the behavioural phenomena he described, correspond to concepts systemised by behavioural finance research over three centuries after his time. *Confusion of Confusions* is often cited as (and, indeed, is) the first book ever written on stock market trading and investors' behaviour. However, although it has motivated some academic attention in terms of the trading practices/tools it describes[1], very little discourse exists as regards its importance to behavioural finance, with the book's actual relevance to behavioural finance being extremely under-researched to date.[2]

The book at hand aims at remedying this situation, by conducting a comprehensive study of *Confusion of Confusions* from the perspective of the behavioural finance paradigm. Its crux lies on critically assessing the behavioural trading beliefs/patterns described by Vega in terms of modern behavioural finance theory. To offer some context, this book first embarks on a critical introduction to behavioural finance and how it challenges the neoclassical finance paradigm (Chapter 1). Chapter 2 provides a detailed analysis of *Confusion of Confusions*' historical context, in terms of the author, his Sephardi community and its role in the Dutch maritime trade, the establishment of the Dutch East Indies Company and how this ultimately motivated the launch of the Amsterdam Stock Exchange. The assessment of *Confusion of Confusions* based on the behavioural finance paradigm takes place in Chapter 3, drawing on ample comparisons between issues surfacing in Vega's descriptions/discussions and contemporary behavioural finance concepts. Chapter 4 outlines a series of implications of this study for various stakeholders (researchers; investment community; regulators and policymakers; the Sephardi community).

The book ultimately shows that there is nothing new under the sun; biases, heuristics, and collective behavioural phenomena extensively investigated by behavioural finance scholars were pretty much present in

[1] Examples of such studies include De Marchi and Harrison (1994), Muprhy (1997) and Szpiro (2011).

[2] The only work that has discussed *Confusion of Confusions* from a strictly behavioural finance perspective (albeit with narrow a scope, drawing on a limited number of quotes from it) was the paper by Corzo et al. (2014); a short discussion on the relevance of Vega's book to behavioural finance is also included in Economou et al. (2023).

the 17th century Dutch market. This is rather disturbing, considering that, despite our institutional and technological evolution, we commit errors of judgement remarkably similar to those of investors almost three-and-a-half centuries ago. This clearly shows that, whereas our material civilisation has evolved, the same cannot be argued about our nature; whether we fail or refuse to learn, we still make the same mistakes. *Confusion of Confusions* leaves no doubt about this, thus allowing us to concur that it constitutes the precursor to modern-day behavioural finance.

This further helps elevate Vega's position in behavioural finance scholarship and history; beyond simply treating the writing of *Confusion of Confusions* as simply a pastime for him (which it may also have been)[3], it is shown that its relevance to behavioural finance is far deeper than popularly acknowledged. Whereas his book has often been treated as a literary narrative of his time, it is clear that it contains very precise allusions to concepts that behavioural finance researchers have been working on during the past at least three to four decades. Far from forming a simple popular finance writing (as has often been dubbed), a detailed examination of it reveals that it forms the first actual treatise on behavioural finance issues by way of narration (in his own very unique style, which shall be elaborated on in Chapter 2). To that end, Vega deserves to be treated as the earliest founder of behavioural finance-on whom contemporary researchers can draw for inspiration. Furthermore, this should also suffice to secure him a position within the broad circle of his Sephardi community's scholars who helped shape scientific discourse in several branches of economics and finance during the Renaissance era and its aftermath. The present book helps raise awareness of an aspect of Vega's contribution to scholarship largely underexplored–if not outright downplayed, to date. It is my hope that the reader enjoys it and that it endows him with the motivation to learn more about behavioural finance–which Joseph de la Vega (unknowingly) helped found.

Durham, June 16th, 2024

[3] See Kellenbenz's comments in De la Vega (1957, xi).

ACKNOWLEDGEMENTS

This book was originally motivated via a series of online lectures on *Confusion of Confusions* I delivered in the Spring semester of academic years 2021-22 and 2023-24 for the Continuing Education Department of the University of Liverpool. I would like to thank Dr Glenn Godenho, Academic Director of Continuing Education for the University of Liverpool, for embracing these lectures (as well as any other lectures I have offered to deliver for his department since 2019); I would also like to thank the Department's members of the student support team for their invaluable assistance over the years. Special thanks also go to the participants of those lectures for the support, questions, and comments; ultimately, they helped me realise that expanding on the subject in the form of a book was the right thing to do.

CHAPTER 1

BEHAVIOURAL FINANCE: AN INTRODUCTION

The role of psychology in financial decision-making had been long acknowledged by early classical economists (e.g., A. Smith; W.S. Jevons), who nevertheless harboured reservations as per whether the impact of behavioural/emotional factors could ever be satisfactorily captured in standard economic models (Camerer et al., 2004; Zak, 2004; Sapra and Zak, 2008). This gradually began to change following the evolution of cognitive psychology and decision science, which allowed researchers by the 1980s to develop tools and measures of investors' behaviour. To better understand how Vega's book constitutes the earliest treatise on behavioural finance concepts, this chapter will discuss how behavioural finance evolved as a response to the previously established neoclassical finance paradigm and outline its key principles.

1.1 Neoclassical finance

1.1.1 The Economic Man

The neoclassical (also, known as rational) finance paradigm views financial decision-making largely as a process carried out with (some) mathematical precision, based on quantitative models/formulae. The archetypal representative agent of human behaviour in this paradigm's context is the "Economic Man" (*homo economicus*); first coined as a term by John Stuart Mill in 1836, it entailed two key features, namely rationality and utility-maximisation.[4]

The notion of rationality encompasses several aspects, which will now be presented in some detail. To begin with, rationality initially presupposes investors having complete knowledge and understanding of all information relevant to their decisions. In other words, they know all factors comprising

---

[4] For more on the genealogy of the term, see Bee and Desmarais-Tremblay (2023).

the true model of (i.e., the one that most accurately describes) an economic variable, possessing structural knowledge of the latter. Examples of such knowledge include knowledge about demand/supply functions of that variable, knowledge of how to estimate present and future general equilibrium prices for that variable, as well as knowledge about its stochastic properties over time (Brav and Heaton, 2002). In this context, if an investor, for example, aims at valuing (or forecasting the price of) a stock, the information she has on the factors affecting its price is sufficiently complete to allow her to build the correct valuation/forecasting model. For this to be possible, however, "sufficiently complete" information is not enough.

Enter here the second element of rationality: investors are assumed to be completely rational information processors, who make optimal statistical decisions with zero distortions. Assuming the above holds, investors would then be expected to combine all relevant information the right way to come up with an accurate valuation of an asset (or forecast of its value). This means that they: i) know all relevant information; ii) interpret it correctly; and iii) combine it correctly. Up to this point, the investor is said to bear *consistent beliefs*, which means that the subjective distribution she uses to value a variable (or forecast its future values) is the actual distribution reigning that variable's price-formation process (Sargent, 1993). The next logical step is for the investor to put together an empirical specification (i.e., a model) to capture the formation of the variable's prices (or to generate forecasts of its future values); to ensure the model is the correct one, she needs to enter any relevant factor in that model with the appropriate weight and ensure the model's functional form is correct.[5] In view of the above, this model will then be the true/correct one, whose coefficients can be estimated with no biases (Friedman, 1979) and allow her to reach the correct valuation/forecast for her variable of interest.

An issue arising from the above is whether investor's rationality implies that her expectations are also rational; after all, even the most rational investor with strong and unbiased processing capacity may encounter limitations in her valuations/forecasts, if the information available is bounded. Brav and Heaton (2002) outline this issue as one of information-exploitation versus information-availability. Assuming the rational investor resides in a rational expectations world, she is capable of exploiting the complete set of relevant information perfectly rationally. If the investor, however, resides in a non-rational expectations world, she will be unable to possess knowledge of the complete set of relevant information; in that case,

[5] E.g., whether the model is linear/non-linear, whether the factors entering it are scaled, standardised, in logarithmic form etc., to mention but a few.

she will have to resign herself to applying her rationality to the processing of only the subset of information available to her.[6]

Further aspects of rationality include Bayesian updating and superior foresight (Barberis and Thaler, 2003). The former suggests that investors update their beliefs with any newly incoming information–assuming that the processing of any new information signal is free from distortion. Superior foresight does not mean that the investor's valuations/forecasts are 100% accurate. Errors are not precluded; however, they are treated as a reflection of a lack of performance (i.e., they are due to fatigue or not enough effort)[7] and are assumed not to repeat themselves.

The second feature of the Economic Man is utility-maximisation; according to this, investors compare investment options and choose the one that allows them to enjoy the maximum expected utility. To generate the latter, investors identify each option's possible outcomes and the probability of each outcome's occurrence, before calculating the weighted average of all possible outcomes. In response to the argument that an investor can/need not know the exact probability of each outcome, Savage (1964) proposed the subjective expected utility (SEU) theory, according to which, each investor calculates his own subjective probabilities. Even if she knows that these probabilities entail a higher level of ambiguity, her decisions are presumed not to be affected, since she is assumed to possess consistent beliefs (in which case, her subjective probabilities are presumed to approximate the actual ones).[8] Expected utility theory further predicts that investors are risk-averse, with utility increasing with wealth at a decreasing rate; beyond a certain wealth-threshold, larger amounts of wealth are associated with greater risk. As a result, a risk-averse individual will only accept a risky gamble, if its expected utility exceeds the utility one derives

---

[6] Such conditions would imply structural uncertainty and, in their presence, the investor will be unable to infer the true distribution of their variable of interest and would likely have to learn it from the available data (or experience). See Brav and Heaton (2002) for more on this.

[7] In this case, the possibility of errors being due to lack of full understanding of information processing (which would render them "competence errors" and imply irrationality) is precluded; no recourse also exists to limitations in the computational abilities of the individual, despite the complexity surrounding algorithmic optimisation. For more on the distinction between performance and competence errors, see Stein (1996).

[8] Although it is hard to fathom the subjective probabilities of each individual to satisfy this, the mere identification of those probabilities allows, technically, for the calculation of expected utility; this suggests that, in the context of SEU it is the probabilities of outcomes that influence choices, not our level of confidence in those probabilities.

from a riskless gamble–i.e., as long as gambling makes sense on a risk-adjusted basis.

Expected utility theory relies on the axioms proposed by von Neumann and Morgenstern (1944), which include the properties of completeness, transitivity, independence, continuity, cancellation, dominance and invariance (see the discussion in Tversky and Kahneman, 1986). Completeness refers to the case whereby an individual's available preference options are all comparable; transitivity denotes that choices are consistent[9]; independence suggests that, when comparing two alternatives, we base our choice on their distinct, not common, aspects[10]; continuity implies that no option is infinitely more/less desirable than any other option[11]; cancellation denotes that, if two options yield similar outcomes when a state holds (but not when it does not hold), then that state can be ignored when comparing the two options; dominance is the condition whereby, if option A is better than other options in one state and at least as good in all other states, then A constitutes the dominant option; finally, invariance is about different representations of the same choice problem yielding the same preference (i.e., framing should/does not matter).

As a result, the Economic Man was assumed to be of unbounded cognitive capabilities:[12] he has access to all relevant information, which he can process unbiasedly and which he uses to update his beliefs in a Bayesian fashion. Even when his forecasts are mistaken, his mistakes are not systematic (they are not repeated twice), nor do they confer any emotional impact over him. Being in possession of unbounded cognitive capacity, he is able to handle large computational tasks; this allows him to be familiar with the expected utility of each risky option (since he is aware of all its

---

[9] This condition implies that each option's value is evaluated individually and does not depend on alternative options. Assume the utility of option A is greater than that of B, and B's greater than C's; in this case, you choose A over B, B over C, and A over C. If, for some reason, option C possesses features that render its rejection regrettable for you, then factoring regret into the consequences of choosing option A or B versus C will distort the transitivity principle. In this case, the choice will not be strictly based upon a common measure (utility) but also depend on non-utility related features (regret).

[10] E.g. if you have three options, let A, B and C, and A + C is preferred over B + C, then A will be preferred over B, irrespective of whether one accounts for C or not.

[11] If e.g. option A is preferred over B, and option B is preferred over C, it is possible that there exists a combination A + C viewed by the individual as indifferent to B.

[12] This has been eloquently outlined by Thaler and Sunstein (2009, 7) as follows: "If you look at economics textbooks, you will learn that homo economicus can think like Albert Einstein, store as much memory as IBM's Big Blue, and exercise the willpower of Mahatma Gandhi."

possible outcomes and each outcome's probability) and, to that end, chooses among different options with the aim of maximising his utility.

1.1.2 Market efficiency: Assumptions and criticisms

If the Economic Man dominates capital markets, it is reasonable to assume that traded assets will be priced largely in line with their intrinsic value[13], the latter varying with the flow of information arriving at the market. This market setting is formally deemed to be efficient, in line with Fama et al. (1969) and Fama (1970), who posited that efficient markets are those in which prices respond instantly to any new information and incorporate all available information at any point in time over time.[14] As per this definition, once information is incorporated into prices, it is impossible to realise abnormal returns by trading on this information; the notion here is that, once information arrives, its impact over prices alters the market and creates a new market valuation (different from the one prior to the arrival of the information in the past and also different from the one to arise once new information arrives in the future). As a result, since prices respond instantly to new information and the arrival-rate of new information obeys no specific pattern (i.e., news arrives randomly), the market cannot produce consistent/predictable patterns of prices (Farmer, 2002).[15] Taking together the randomness of news' arrival with the fact that prices respond instantly to news, suggests that, in an efficient market, prices follow a random walk in their formation process.[16]

---

[13] A stock's intrinsic ("fundamental") value is calculated by discounting the sum of the expected cash flows from the stock (e.g. its future dividends, earnings etc.) to the present at a given rate.

[14] The earliest known writings on the role of information in shaping prices can be identified with James Maitland and Yamagata Bantō; see the discussion in Takatsuki and Hisamatsu (2023).

[15] Two things are important to note here: first, even if a pattern were to arise, it would be quickly exploited by investors (or arbitraged away)–and eventually disappear (Farmer, 2002); second, the presence of transaction costs would further contribute to the dissipation of the profitability of any such pattern (Jensen, 1978).

[16] The concept of "random walk" in stock prices was first described by Jules Regnault in 1863 (Jovanovic and Le Gall, 2018); as Gujarati (2003) eloquently described it: "The term random walk is often compared with a drunkard's walk. Leaving a bar, the drunkard moves a random distance *ut* at time *t*, and, continuing to walk indefinitely, will eventually drift farther and farther away from the bar. The same is said about stock prices. Today's stock price is equal to yesterday's stock price plus a random shock." (Gujarati, 2003, 798). For a discussion of early evidence on the random walk, see Woo et al. (2020).

Although, as mentioned above, this means that one cannot beat (i.e., outperform) the market systematically, this does not mean that this can never happen. It is possible, for example, that an asset is temporarily over/underpriced (in which case, you can profit from exploiting its mispricing) or that an asset is riskier than the market average (in which case, by assuming higher risk when trading on it, you can realise an above-market return). As a result, you are rewarded with a return proportionate to the risk you take ("no free lunch").[17]

Contingent on the type of information incorporated in securities' prices, market efficiency can assume three forms (Fama, 1991). A market is weak-form efficient when prices fully reflect all information contained in past price movements. If this holds, then looking at past prices (or past trading data, in general, including for example past volume, turnover etc.) is assumed to contain no relevant information about future prices, thus rendering any attempt at extrapolating from historical prices[18] an exercise in futility.[19] A market is semi-strong form efficient, if prices fully reflect both all relevant publicly available information and all information contained in past price movements (thus suggesting that a market cannot be semi-strong form efficient without being weak-form efficient). In this case, the instant incorporation of new information into prices suggests that any attempt by investors/analysts to compute a stock's fundamental value following the arrival of news will lag behind the market's reaction to the

[17] In reality, "no free lunch" need not coexist with fundamentals-based prices; as Barberis and Thaler (2003) stated, "While both are true in an efficient market, "no free lunch" can also be true in an inefficient market: just because prices are away from fundamental value does not necessarily mean that there are any excess risk-adjusted average returns for the taking" (Barberis and Thaler, 2003, 1057).

[18] The practice of extrapolating from historical prices is more formally known as feedback trading and can entail a host of trading strategies (momentum; contrarian; technical analysis) and practices/tools (stop-loss orders; take-profit orders; portfolio insurance). For a detailed discussion of the literature on feedback trading, see Economou et al. (2023).

[19] Neoclassical finance traditionally tended to view trading on historical prices with ambivalence. Lo et al. (2000, 1705) mention that "It has been argued that the difference between fundamental analysis and technical analysis is not unlike the difference between astronomy and astrology. Among some circles, technical analysis is known as "voodoo finance." And in his influential book A *Random Walk down Wall* Street, Burton Malkiel (1996) concludes that "[u]nder scientific scrutiny, chart-reading must share a pedestal with alchemy."". Nevertheless, there exists a host of evidence (see Hoffmann and Shefrin, 2014 and the discussion therein) suggesting the wide popularity of technical trading strategies among investors.

news–in turn, casting doubt on the usefulness of fundamentals' analysis.[20] Finally, a market is deemed strong-form efficient, if prices reflect all privately held information, all relevant publicly available information and all information contained in past price movements.

Market efficiency presupposes that the trading process is dominated by rational investors who compete against each other using their information to predict securities' future prices (Fama, 1965). Given their competition, securities' prices will, at any point in time, reflect their information and tend to hover around their actual (intrinsic) values. Of course, to the extent that a company's intrinsic value depends upon its earnings' potential (and given that future earnings are impossible to know with certainty), the intrinsic value will always entail an element of uncertainty (Black, 1986); as a result, the actual price will always be an estimate of (and wandering randomly about) the intrinsic value.

In the context of the efficient market hypothesis, it is further possible that some investors do not subscribe to the paradigm of the Economic Man, or indeed any of its features; these investors are deemed irrational and can be accommodated in the neoclassical finance paradigm in two ways. First, even if they are of good size, their trades are assumed to be random (some may under- and others overvalue a stock; some may under- and others overreact to information) and thus, cancel each other out (Fama, 1998). Second, if their trades are correlated (i.e., there exists systematic irrationality) and motivate mispricing, rational investors can eliminate the latter via arbitrage (Friedman, 1953; Fama, 1965).[21] As a result, although

---

[20] The following quote from Mathiopoulos (1994) summarises this eloquently: "Based on the theory of efficient markets, namely that the price of a stock in the market reflects all available information relevant to the course of the underlying corporation, it is worth wondering whether it is the analyst that comes up with the forecast or the stock itself–and which purpose the analyst serves in the end" (author's translation from the Greek original).

[21] Arbitrageurs are investors who earn riskless profits at no cost. As per Hirshleifer (2014, 6): "Arbitrage is the purchase or sale of goods to profit from differences in effective prices across trading venues. The term is used broadly to refer to the exploitation of profit opportunities whenever some assets are overpriced relative to others, based on the idea that buying cheap assets and selling similar but expensive ones can yield a relatively low-risk return." In this case, investors go long on cheap and short on expensive similar assets. Going short implies short selling, namely selling something you do not currently own with the obligation of buying it back at a given price in the future. You obviously do that if you expect the price to fall, the expectation being that overpriced stocks will see their prices fall and underpriced ones their prices rise. The practice of short selling can be traced virtually since the birth of capital markets, as De la Vega (1957) documents (and as shall be discussed

market efficiency does not refute the possibility of non-rational/less-than-perfectly rational investors, it views their impact over the price-formation process to be of no import.[22]

An intriguing view on this issue was advanced by Friedman (1953), who argued that in an efficient market, investors act "as if" they are trying to rationally maximise expected returns with full knowledge of necessary data. "As if" means that the behaviour of individual investors matters little; rather, it is the aggregate outcome of their actions that will have to be rational for the market to be efficient. The presence of "as if" in the efficient markets' paradigm may bring the latter closer to reality (after all, it is hard to presume homogeneity among investors, much less homogeneity in their rationality), yet also complicates things, as it implies that valuing a stock can be an uncertain and extremely vexing process (as it assumes the interaction of traders of various rationality levels). With "as if" the focus now shifts from procedural rationality (which examines whether the strategies and rules underlying behaviour are rational or not) to theoretical rationality (which judges behaviour only by its outcomes, regardless how they were achieved). In this case, an individual reaching a rational decision by chance or imitation (e.g., of his rational peers) is also perceived as

---

later on in this book); for more on this, see also Marco and Van Malle-Sabouret (2007) for arbitrage on the bonds issued by the Dutch East India Company and Dempster et al. (2000) and Bell et al. (2016) for evidence from the 18th century London and Amsterdam markets. Beyond the finance-related literature, very early references to situations reflective of differential pricing of similar items are included in ben Maimon's Ma'aser Sheni 4:21 ("When a person transports produce from a place where it is expensive to a place where it is inexpensive, he should redeem it according to its value in the place he is redeeming it.") from the 2nd-3rd centuries AD; and Wu Ch'eng-en (1973, 304)'s 16th century Chinese novel "Monkey" ("[…] and he told his servants to saddle six horses, but not to let Tripitaka's horse be saddled for the present. Then they all rode down to the river to look. It was true enough that people were walking across the ice. "Where are they going?", asked Tripitaka. "They are traders", said Mr Ch'en, "from the country on the other side of the river. Things that here are worth a hundred strings of cash are worth ten thousand on the other side. And things that are there worth only a hundred are here worth two thousand. So great are the profits to be made that they are willing to undertake the journey even at the risk of their lives.").

[22] Rational or not, investors may well exhibit disagreement as per their interpretation of information due to behavioural factors (biases/heuristics; see next section) as well as persistent differences in opinions (Carlin et al., 2014). The potential of investors' disagreement was not viewed by Fama (1970) as evidence against efficiency "…unless there are investors who can consistently make better evaluations of available information than are implicit in market prices" (Fama, 1970, 388).

rational, thus rendering possession of the formal properties of rationality less important in the attainment of rational outcomes.[23]

Despite early empirical support for market efficiency (Fama, 1965; Fama and Blume, 1966; Ball and Brown, 1968; Jensen, 1968; Fama et al., 1969; Waud, 1970; Scholes, 1972; Keown and Pinkerton, 1981), criticism began to mount shortly thereafter. To begin with, although the neoclassical finance paradigm prided itself over the mathematical calibration of its empirical designs, there exists to date no singular "correct"/"right"/"true" model of asset pricing. This is an issue most prominently reflected in the so-called joint-hypothesis problem (Fama, 1970; 1991); according to this, it is not possible to *ad hoc* test for market efficiency in isolation without considering an asset pricing model. Since, by definition, prices in an efficient market properly incorporate all information, an asset pricing model is needed to quantify how "proper" this incorporation is. As a result, if an anomaly (i.e., a deviation from efficiency) surfaces in the return-generation process, it is hard to assert whether this is due to the market being inefficient or whether this is due to the asset pricing model used being the wrong one. The problem is further aggravated by the established (Cochrane, 2011) "factor zoo", since the huge number (in the hundreds) of common risk factors[24] proposed during the past few decades clearly suggests that it is practically impossible to come up with a single model that works equally well across all markets, time periods and economic conditions/market states. Lo and MacKinlay (1999) further argued that market efficiency needs more hypotheses (related to investors' preferences, information structure, and business conditions) to be added when testing for it. If this were to happen, they argued, that would render testing for market efficiency a test of multiple hypotheses simultaneously, in which case a rejection of

[23] See the discussion in Merkle (2007) for more on the issue of theoretical-versus-procedural rationality.

[24] Feng et al. (2020) assessed 150 factors, Harvey et al. (2016) and Neuhierl et al. (2023) investigated over 300 factors, while Harvey and Liu (2019), Hou et al. (2020) and Bartram et al. (2021) explored over 400 factors. Bryzgalova et al. (2023) drew on 51 factors to test for 2.25 quadrillion models based on those. Beyond the factor zoo, an additional problem relates to the choice of the proper proxy for market performance, since this will be needed as a benchmark to calculate excess market and market-adjusted returns in asset pricing models. If dealing with 300-400 factors sounds onerous enough, one should probably brace themselves before the number of possible benchmarks which can be employed in asset pricing: as a 2018 Financial Times article by Authers (2017) showed, there exist over 3 million equity indices. The number becomes even more astounding considering that the total number of publicly listed companies worldwide is less than fifty thousand.

the efficiency hypothesis would render it impossible to assert in which of those sub-hypotheses the root of inefficiency lies.

Another issue surfaced with the Grossman-Stiglitz paradox (Grossman and Stiglitz, 1980); according to this, if the market is efficient and all information (including insider information) is reflected in the price, then no one has an incentive to expend resources to gather information and trade on it. If this holds, the question then arising is how/whether it is possible for all information to be reflected in the price. If, for instance, we assume a trader in possession of costly inside information (costly in terms of both its collection and its legal repercussions–in case he is prosecuted for trading on it), it is hard to fathom that he will incur that information's costs without the possibility of being compensated by profitably trading on it; this is one example, whereby market efficiency (in this case, its strong form) cannot apply. As a result, Grossman and Stiglitz (1980) showcased that informationally efficient markets constitute a practical impossibility.

While the above criticisms were based on analytical-theoretical frameworks, research also saw the emergence of a series of empirical criticisms levelled out against market efficiency. Much of this research focused initially on calendar/seasonal anomalies[25], as well as on the presence of anomalies in stocks of specific features (e.g., size and value effect).[26] Ball and Brown (1968) documented the presence of a post-earnings-announcement drift in stock returns (supportive of underreaction in securities' prices). Shiller (1981) and LeRoy and Porter (1981) showcased that stock market prices are far more volatile than could be justified by a rational model (assuming constant discount rates), thus suggesting that prices entail noisy components, something further elaborated on by Black (1986). Mehra and Prescott (1985) argued that equity returns command a premium over risk-free investments over and above what risk-aversion would warrant, while further evidence denoted that stocks exhibit short-term continuation (momentum; Jegadeesh and Titman, 1993) and long-run reversals (DeBondt and Thaler, 1985).

---

[25] Examples of such anomalies include the January effect, the Monday effect, the weekend effect, the holiday effect and the turn-of-the-month effect, to mention but a few key ones. For more on those, see Lakonishok and Smidt (1988) and Plastun et al. (2019).

[26] Size effect was initially documented by Banz (1981), who reported higher expected returns for small versus large capitalisation stocks; Lakonishok et al. (1994) demonstrated that high (low) book-to-market stocks tend to out(under)perform in future periods.

1.2 Behavioural finance

The systematic unearthing of empirical anomalies/regularities in the 1980s focused researchers' attention to explanations beyond the realm of neoclassical finance, with the focus shifting to evidence from cognitive psychology and decision science from previous decades.[27] This gradually gave birth to a new interdisciplinary research domain, behavioural finance, which primarily challenged neoclassical finance's key assumptions of rationality and arbitrage.

1.2.1 How behavioural finance challenges rationality

1.2.1.1 The 2-System paradigm

As per the debate on rationality, behavioural finance has largely relied on a two-system representation of the human mind (Camerer et al., 2005; Kahneman, 2012); contrary to the neoclassical finance paradigm, where rationality was assumed to be key, the two-system representation assumes that humans can entertain two distinct decision/thinking modes, one that is more "intuitive" (reflexive) and one that is more "rational" (reflective).[28]

The intuitive system (also known as "System 1") is largely uncontrolled (residing outside awareness/consciousness) and involves automatic, reflexive responses, generated with little/no effort.[29] System 1 resorts to automatic processes, which are relevant to solving evolution-related problems (e.g., issues of survival), without necessarily adhering to logical principles (Ardalan, 2018). These processes are rapid, operate in parallel, and, hence allow for multitasking. Relying on patterns retrieved from associative memory, System 1 hinges heavily on habits and impressions to find shortcuts for problems and is prone to emotions. As it relies on heuristics, rather than analytical processing, it aims at passing judgement about people/things/situations, rather than engage into deeper reflection of those (hence being easy to jump to conclusions). Examples of tasks performed by the reflexive mind include driving a car and avoiding a food associated with unpleasant past side-effects to us.

---

[27] For more on how this cross-fertilisation of finance and psychology/decision science research took place initially, see Kahneman (2012).
[28] See Camerer et al. (2005) for a discussion of earlier literature involving alternative versions of similar 2-system models.
[29] Since automatic processes rely on long-established habits/reflexes, we are unable to scrutinise them (they are embedded in our psyche and performed unconsciously) or develop introspective insight into how we decide/judge based on them.

The reflective system (also known as "System 2"), is largely controlled and deductive, processing information analytically/slowly and with effort (following a series of rules), without emotions interfering with this process. Unlike System 1, it treats every thought, concept, person or object with no assumptions, thus maintaining self-control. System 2 uses controlled processes (Camerer et al., 2005; Ardalan, 2018), which are serial (i.e., follow clearly demarcated steps), deliberate (they are used when encountering a novel/special challenge/situation) and are subject to introspection (being aware of controlled processes, we can describe them). Examples of tasks performed by the reflective mind include building a portfolio of stocks and calculating mortgage payments under alternative rates.

Overall, human behaviour in this setting is viewed (Kahneman and Frederick, 2002) as the outcome of an interaction between controlled and automatic processes[30]; whereas controlled processes are the products of more recent biological evolution in humans (they involve computational/analytical skills), automatic processes rely on more primitive brain strata (as they involve elements of affect and intuition).[31] [32] [33] System 1 is the brain's

[30] The relatively nascent interdisciplinary domain of neurofinance has proposed an expansion of the 2-system model into a 4-system model; according to the latter, controlled and automatic processes intersect with cognitive and affective ones. For a detailed presentation of this, see Camerer et al. (2005).

[31] This hinges on what is known as the "triune brain"; according to this, the human brain is divided into three regions that correspond to different stages of its evolutionary development, namely the reptilian (focusing on the organism's basic survival functions), the mammalian (focusing on social emotions) and the hominid (focusing on more cognitive/analytical processes) brain. For more on this, see Camerer et al. (2005), Peterson (2014) and Desmoulins-Lebeault et al. (2018).

[32] Controlled (System 2) processes can often transition into System 1; deliberative tasks that are performed repeatedly can prompt our brain to automate their processing. See Ardalan (2018) for more on this.

[33] Systems 1 and 2 need not necessarily be at odds. If both systems generate the same answer to a problem, System 2 will actively endorse System 1's judgement; if, however, System 1 is wrong and System 2 fails to detect it, then System 2 will passively endorse System 1's judgement. In this setting, both types of endorsement (active; passive) imply a connection between error-detection and error-correction. An interesting example of a third perverse interaction between the two systems is presented by Risen (2016). Exploring how humans respond to superstitions, she proposes the acquiescence model. I will present it here in some detail to showcase how the interaction between the two systems is anything but simple. To begin with, superstitious intuitions hail from System 1, as long as System 2 says so; if System 2 is (not) activated, those intuitions can be tackled (prevail). The question therefore is whether System 2 will intervene to correct the false beliefs emanating from superstitions and motivate the individual to reject them. For this correction to take

"default mode", since deliberation can be very taxing for the brain, requiring concentration and, hence, effort. To the extent that it competes for attention and cognitive resources with the rest of the brain's tasks, thinking taxes the mind, more so because it is neurologically "expensive" (about a quarter of our caloric intake is needed to keep the brain running; Forsythe, 2009). As this suggests that thinking is costly, this adds to the costs of decision-making, prompting humans to apply thinking only up to the point where its costs do not exceed its benefits. In turn, this suggests that we do not incorporate all available information into our decisions, which, as a result, are not optimal.[34] This forms a first serious challenge to rationality (and, *in tandem*, market efficiency): if investors do not utilise all signals they can from their informational environment, this denotes that their trades will not reveal all available information. Consequently, market prices will not incorporate all relevant information either, thus compromising a market's efficiency.

System 1 is particularly influential in our understanding of the world, a process mediated by several tools. To begin with, it is important to bear in mind that we live in a signal-rich world, which we cannot cognitively capture/process in its objective entirety;[35] what we do perceive is a truncated, selective version of it, filtered by a series of perceptional factors

---

place, System 2 has to recognise that superstitions are wrong, i.e., detect an error in our thinking mode. If this error is (not) detected, then there exists (no) error-correction. In this case, System 2 can i) correct a false belief (superstition) or ii) endorse it, either because it thought it to be correct, or because it failed to notice that it is wrong. Risen suggests that acquiescence can introduce a third form of "endorsement": System 2 detects the error of System 1's superstition, but does not correct it, thus acquiescing to superstition, even though it knows it to be false. In this case, error-detection and error-correction are decoupled, thus casting doubt over the assumption that detection of an error leads people to correct it; in this specific case, we can maintain superstitious beliefs, even though we realise they are irrational. See Risen (2016) for more on this.

[34] Of course, attaining the optimal outcome need not always be the case in the first place; it may be that optimising is not possible or that the computational costs of optimising are disproportionately high, leading to a negative net cost-benefit outcome. As a result, there exist cognitive limitations, in terms of both knowledge and computational capacity, that introduce boundaries to the rationality of individuals–a concept that came to be known as "bounded rationality" (Simon, 1959). According to this, people prefer "satisficing" over optimising, since what is important is to satisfy the drives of our actions, not to attain an optimal outcome.

[35] Although the human brain can capture around 11 million pieces of data per second, only about 40 of them will be consciously processed (Konnikova, 2013).

(Konnikova, 2013).[36] As a result, much of our attention is motivated, depending largely on which aspects of the world trigger our perception's focus. Even if, however, our attention allowed us improved powers of reception, a second obstacle arising would be one of storage: our memory is subject to flaws that hinder us from proper recollection–or even accurate recollection. Memory fails in two domains, namely content (what recollections I have managed to store) and structure (how I have stored them), and this leads to errors in memory-performance (Konnikova, 2013); it also leads to motivational encoding (the situation whereby we remember things better when we are interested and motivated) and motivation to remember (when we remember what we want/think is important to remember). Taken together, the above suggest that both attention and memory are motivated, thus implying that each of us possesses their own map of reality; as we are not all equally motivated to see the same things in the same way, what we end up observing is not THE reality, but OUR reality. System 1 uses a host of tools that shape our motivated perception of our environment; some of these tools affect the reception of signals from our environment (they are known as "biases") and others affect the processing of these signals (they are known as "heuristics").

1.2.1.2 Biases and heuristics

Biases are psychological forces that distort our perception/reception of signals from our environment. A bias resembles a tinted glass: you know there is something behind it but you cannot see it clearly. This may be either because of limitations in human perception or the way in which a situation is presented. Assume for example a company that has been posting positive earnings for the past eight quarters and that in this quarter posts negative earnings. How would you react to this? There is no singular right (or wrong) answer here. One possibility would be to sell the stock, either for fundamentals-related reasons (if you believe the negative earnings to signify adverse information about the prospects of the company) or behavioural ones (e.g., because you panic). It is, however, equally likely that you will hold onto that stock, since the current quarter of negative earnings will be

[36] These factors include (Konnikova, 2013): coincidence (if a stimulus reminds you of a past experience involving a similar stimulus, hence capturing your attention; you then extrapolate the outcome of the past experience to the present); relevance (a stimulus that is relevant to you for some reason, will capture your attention); subjectivity (a situation can capture your attention, if you believe it to match your beliefs); exclusivity (not paying attention to information because of neglect); and engagement (when "flow" makes you immerse in a task with full attention).

viewed within the frame of the eight previous quarters' positive earnings (the current negative earnings will appear as something close to a blip in an established trend of positive earnings). The latter would be reflective of *conservatism bias* (more on this shortly) and would be a function of how you would be framing your picture of the company's earnings. Let us now add a bit more information into this example, by assuming that earnings in each of the previous eight quarters hovered within the £50,000-£100,000 region and the losses in the current quarter are equal to a negative £1,000,000. Clearly, this magnitude of losses would be salient enough to capture your attention more vehemently, increasing the likelihood of you selling the stock[37]; as a result, how a problem is framed can affect our responses to it, by motivating the surfacing of different biases into the decision-making process.[38]

Heuristics, on the other hand, represent innate processes that affect how we process information. They constitute rules of thumb that have been acquired via (largely reinforcement) learning over the years and form cognitive algorithms operating automatically, without us being aware/conscious of them (Hirshleifer, 2015, 136). Primarily deployed when dealing with obfuscated choice problems, they offer shortcuts towards easy/satisfactory (though not optimal) solutions. Imagine, for example, that you want to build a portfolio of 10 stocks from the S&P500 index-universe; choosing 10 from among 500 stocks is obviously not an easy task and clearly merits some rule that will justify the selection-process. To facilitate this process, the human brain can come up with various shortcuts; in this case, you can: focus on the 10 largest S&P500 index-constituents in terms of capitalisation; pick up the 10 names from the index-list you recognise; pick up the first 10 stocks from that index mentioned in the Wall Street Journal today; pick up yesterday's top 10 performing S&P500 index-constituents; ask your financial advisor; ask your wife/husband etc. Although none of these shortcuts are likely to be ideal in helping you construct the optimal portfolio, they are each capable of simplifying the selection-process, thus denoting that heuristics help provide satisficing, yet not optimal, solutions.

The following example will help demonstrate how biases and heuristics can lead individuals to deviate from rationality. Assume stock ABC and suppose that its price has been rising for a few weeks in a row; you treat this as a trait representative of a good stock and, hence buy it. This is reflective of the *representativeness heuristic* (Kahneman and Tversky, 1972): drawing on a limited sample of observations (a few weeks' prices), you generated

[37] Conversely, if the current quarter's earnings were a meagre negative £10,000, you would be likely to pay far less attention to it.
[38] This, in turn, violates the invariance principle mentioned earlier.

inferences about the population of observations (the stock is good).[39] Let us now assume that the stock's price keeps rising and after some time it experiences a couple of days of negative returns; it is likely that after such a long spree of positive returns, you may view one or two down-days as a blip, downplaying their importance. If so, this is reflective of the *conservatism bias* (Edwards, 1982), which implies slow updating of your beliefs in view of newly incoming information.[40] There is, of course, the possibility that the large (small) number of price rises (falls) prompts you to think that it is more likely that the stock will eventually start falling and, hence decide to sell the stock in advance. This will be motivated by an extrapolative belief in mean-reversion, a belief more formally known as *gambler's fallacy* (Barberis and Thaler, 2003).[41] Of course, before any of the above-mentioned psychological forces kick in, you will first have to identify your window of reference; in other words, which sample of past prices (past one, two, three etc. weeks) do you choose to focus on? This suggests the presence of *anchoring* (Tversky and Kahneman, 1974), which is about extrapolating from some reference point when having to form a judgement. *Availability bias* (Tversky and Kahneman, 1974) can further impact your judgement: if you have been through an event salient enough in your past investment history that looks similar to the stock's recent price movements, this will prompt you to judge stock ABC based on your

[39] As Hirshleifer (2015, 147) argued with respect to representativeness, "People assess the probability of a state of the world on the basis of how typical of that state the evidence seems to be. This is reasonable if typicality is a proxy for the conditional probability of the evidence given the state of the world." In the context of the specific example mentioned here, the stock's surge in price for a few weeks does not have to be a typical proxy for the stock being a good pick. In reality, a good stock does not necessarily have to document such a spree for a few weeks–nor does a stock rising for a few weeks have to qualify as good. An additional issue, of course, here is defining what constitutes a "good stock". As a result, this helps showcase the key implication of the representativeness heuristic, namely that it prompts people to see patterns–and overreact to recent information.

[40] The implication of the conservatism bias is underreaction (in effect, momentum) in the market, since it will take longer for new information to be incorporated in prices. It also distorts Bayesian updating, since, by downplaying information incongruent with the perceived trend, it hampers investors' learning process.

[41] In gambler's fallacy, people view short sequences of events as representative of long ones; as a result, if the stock's price has been rising for a few weeks, you may believe that a reversal is imminent, since in a longer sequence (e.g., a 10-year price-series), rising prices are followed by falling ones.