



Department of Economics and

Business Administration

Master's Degree in Finance

**From Optimization Models to
Real Choices: Portfolio
Construction Between
Financial Theory and
Investor Psychology**

Supervisor

Chiar.mo Prof. Paolo Giudici

Master's Thesis

of Sebastiano Augusto Matteo Merlini

Matr. n. 558130

Academic Year 2025-2026

TABLE OF CONTENTS

Introduction	4
CHAPTER 1 Theoretical foundations	6
Saving and Investing: A Distinction Not to Be Taken for Granted	6
The Different Forms of Risk	8
The General Principle: Time as a “buffer”	10
CHAPTER 2 Financial Education	13
Why People Don't Invest?	13
What the Data Tells Us	15
A Personal Perspective: My Own Experience	18
CHAPTER 3 Modern Investment Tools, Delegation and Autonomy	22
Now we analyse some real cases	28
The Other Side of Accessibility: Risks Without a Safety Net	31
Toward a Balanced View	33
CHAPTER 4 Empirical part/model construction	37
Data source and sample	37
The ordinal logistic regression: results	38
The simple neural network: results	42
Conclusions	53

LIST OF FIGURES

Table 1. Ordinal logistic regression, coefficient estimates	39
Figure 1 Ordinal logistic regression	40
Table 2. Confusion matrix (test set, n = 2,411)	41
Table 3. Simple neural network, classification performance (test set, n = 2,411)	42
Figure 2. Simple neural network architecture (1 hidden layer, 9 nodes)	43
Table 4. Permutation importance.....	44
Figure 3. Permutation importance, ranked	45
Table 5 Real households Examples	46
Table 6. Confusion matrix (test set, n = 2,411)	47
Table 7. Deep neural network, classification performance (test set, n = 2,411)	48
Figure 4. Deep neural network architecture (3 hidden layers: 5-3-2 nodes)	48
Table 8. Permutation importance.....	49
Figure 5. Permutation importance, ranked	50
Table 9. Confusion matrix (test set, n = 2,411)	52

Introduction

Keeping money in a bank account has long been considered, by most Italian households, the safest possible choice. This thesis begins from a simple but consequential observation: that perceived safety is largely an illusion. Money left idle does not merely fail to grow it loses real value over time, eroded by inflation and by the opportunity cost of returns never earned. Yet despite this, a large share of Italian household wealth remains parked in low- or zero-yield deposits, while genuine investment, even in simple, low-cost, well-diversified instruments remains an activity perceived as complex, risky, or reserved for those with specialist knowledge or substantial capital.

This thesis challenges that perception. Its central argument is that investing has become considerably more accessible than it is commonly believed to be, and that this accessibility, combined with a minimal but genuine foundation of financial understanding, is sufficient for an ordinary person to begin building meaningful wealth over time. Developing this argument requires, first, a clear grasp of the basic vocabulary of personal finance: the distinction between saving and investing, the relationship between risk and return, the logic and purpose of diversification, and the often underestimated role played by time horizon in determining which financial choices are appropriate for which investor. These concepts, while individually simple, are too often blurred together in everyday language and popular belief; much of the confusion and hesitation surrounding investing can be traced back to this lack of basic clarity.

Yet understanding these concepts intellectually is not, on its own, sufficient to explain why so many people continue to avoid investing even when doing so would clearly serve their long-term interests. This thesis therefore also engages with the behavioural and cultural dimension of the problem: loss aversion, status quo bias, and a cultural inheritance, particularly pronounced in the Italian context, that has long treated real estate and bank deposits as the only legitimate forms of family wealth. Financial education, in this sense, is treated throughout this work not merely as a matter of individual knowledge, but as a form of protection: a well-informed investor is better equipped to recognise unrealistic promises, understand the true cost of financial products, and avoid the emotionally driven decisions that so often turn temporary market volatility into permanent, realised losses.

Alongside this theoretical and behavioural grounding, this thesis pays close attention to the practical tools that have, in recent years, made investing genuinely accessible to the ordinary retail investor, a transformation driven largely by financial technology. Two distinct paths are examined and compared: delegated investing through automated robo-advisory platforms, which construct and manage a portfolio on the investor's behalf based on a stated risk profile; and self-directed investing through low-cost digital brokerage platforms, which place the decisions, and their consequences, directly in the investor's hands. Neither path is presented as inherently superior; rather, each is shown to carry its own trade-offs between cost, control, and the degree of financial understanding it realistically demands of the person using it. This discussion is grounded not only in theory but in real, concrete illustration: four actual investment portfolios, belonging to real individuals at different stages of their investing journey, are presented and analysed throughout this work, ranging from the simplest possible starting point: a single diversified ETF accumulated at a modest monthly contribution, to considerably more complex and aggressive allocations involving individual equities and cryptocurrency.

At the centre of this thesis lies an empirical case study that shaped its overall direction: Giudici and Polinesi's (2019) research into the statistical methodologies used to classify investor risk profiles on robo-advisory platforms, and to verify whether the risk actually proposed to an investor matches the risk they are willing and able to bear. Taking this study as both inspiration and methodological starting point, this thesis extends the exercise with an original empirical analysis of its own: three statistical models of differing structure and complexity—an ordinal logistic regression and two feed-forward neural networks—are built and compared on real microdata from the Bank of Italy's Survey on Household Income and Wealth.

Taken as a whole, this thesis is intended to serve a dual purpose: to offer a methodologically sound empirical contribution to the study of investor risk profiling, in direct dialogue with the existing academic literature in this field; and, at the same time, to make the case through theory, real data, and real examples—that investing is neither as complex nor as inaccessible as it is often perceived to be, and that the most important step any investor can take is simply the first one.

CHAPTER 1

Theoretical foundations

Saving and Investing: A Distinction Not to Be Taken for Granted

In everyday language, "saving" and "investing" are often treated as synonyms. In reality, these are two profoundly different financial behaviours, both in their objectives and in their consequences.

- 1) Saving** is the act of setting aside a portion of one's income without consuming it immediately. The natural home of savings is the current account or deposit account: safe, liquid instruments (meaning immediately available), but offering a null return. Saving responds to a logic of prudence and availability: it serves to build a reserve ready for unexpected expenses, for short-term goals, or simply to provide a sense of security.
- 2) Investing**, on the other hand, is the decision to deploy that saved sum into a financial instrument with the expectation that, over time, it will generate a return. Unlike saving, investing always involves a certain degree of risk: the invested capital is not guaranteed, and its value can fluctuate over time.

This distinction, however it may seem obvious once framed in these terms, has solid bases: it is rooted in a deeply ingrained cultural attitude, particularly pronounced in Italy. Italian households have historically been among the most "saving-oriented" in Europe, yet also among the least inclined to move those savings toward investment. This behaviour has understandable explanations: the memory of past financial crises, a widespread distrust of financial markets perceived as distant, complex, or reserved for "the experts of the sector" and a preference for the perceived safety of cash or a bank account, even at the cost of forgoing a return. The problem is that this apparently prudent choice is far from consequence-free, it simply shifts the risk from a visible form (the fluctuation of a financial market) to an invisible and silent one, no less real for being unseen. The central point, often overlooked by those approaching personal finance for the first time, is that keeping savings in a bank account is not, in fact, a cost-free operation. It is an operation with zero nominal risk, the figure on the bank statement never decreases, but it carries an implicit cost: that of inflation, which over time erodes the purchasing power of money. To grasp this concept without resorting to calculations, it is enough to think about what a given amount of money could buy twenty years ago compared with what the same nominal figure can buy today: the number on the account has remained identical, but its real purchasing power has diminished, because in the meantime the prices of goods and services

have risen. Capital left idle for years silently loses value in real terms, a less visible risk than a market crash, but no less concrete, precisely because it acts constantly and cumulatively over time.

There is also a second aspect, distinct from inflation but equally relevant: what economics calls opportunity cost. Every time a sum of money remains idle rather than being invested, one implicitly forgoes the return that sum could have generated over time. This is not a loss that appears on a bank statement, which is precisely why it tends to go unnoticed, yet it is nonetheless a real trade-off, one that accumulates year after year and that, over a long horizon, can make a substantial difference to the growth of one's wealth. These two elements, the silent erosion caused by inflation and the opportunity cost of unused liquidity, lead to a conclusion that will run through this work: once a prudential liquidity threshold has been set aside (a reserve for emergencies and short-term goals), leaving the remaining capital idle is not the "safe and best" choice it is often believed to be, but a choice that still carries risk, simply a less visible one than that of a financial market.

Investing, from this perspective, should not be understood as an activity reserved for few experts or for those who already hold substantial wealth, but as a tool, accessible understandable today more than ever, within reach of anyone thanks to the digital platforms, discussed later in this work, for protecting and growing the real value of one's savings over time. Moreover one of the foundational principles of finance is the relationship between risk and return: in practice, there is no such thing as a higher return without a correspondingly higher risk. In finance is not possible to have a low-risk investment with high possible returns.

- 1) **Return** is the gain (or loss) that an investment generates over time, usually expressed as a percentage of the invested capital. It can stem from different sources: the appreciation of the instrument's price (for example, a share that increases in value relative to its purchase price), or periodic flows such as dividends (in the case of shares) or coupons (in the case of bonds).
- 2) **Risk**, in finance, should not be understood merely as "the possibility of losing money", but more precisely as the degree of uncertainty attached to an investment. That is, the variability of its possible future outcomes. A government bond from a stable country carries much lower risk (and therefore a lower expected return) than a share in a company, whose value can fluctuate considerably in the short term.

It is worth explaining why the relationship between risk and reward exists. No rational investor would accept a higher risk without expecting, in return, a higher

reward: if a risky bond offered the same return as a safe government bond, no one would choose the former. It is precisely in order to attract capital that riskier instruments must offer, in theory, a higher expected return. This is what in finance is called "risk premium". This does not mean that a risky investment *will* necessarily perform better: it simply means that, on average and over the long run, the market tends to reward those who accept greater uncertainty.

The Different Forms of Risk

It is important to distinguish between different forms of risk that an ordinary investor encounters:

- **Market risk:** the possibility that the value of an instrument changes due to general factors affecting the entire economy or the entire market: interest rate movements, a recessionary phase, geopolitical events. This is a risk that affects all financial instruments simultaneously, in fact this risk can not be completely avoid.
- **Specific risk** (or idiosyncratic risk): tied to a single issuer or company. For example, an internal management crisis, a scandal, a failed product launch. Unlike market risk, this type of risk can be significantly reduced through portfolio diversification, as discussed in the next section.
- **Liquidity risk:** the difficulty of selling an instrument quickly, when needed, without having to accept a penalising price. A heavily traded instrument (such as shares of a large listed company) is generally very liquid; a real estate investment or a stake in a small unlisted company is far less so.
- **Credit risk** (or counterparty risk): the possibility that whoever issued the instrument – a government, a company, a bank – is unable to honour its commitments, meaning to repay the capital or pay the promised interest.

For completeness of sake we have to say that a common investor faces prevalently the first, second and third risks elenced. The one remaining is mostly for bigger investors than the retails.

Understanding this distinction is already an important first step in financial education. Many people fear "risk" in the abstract, as a single undifferentiated block, without understanding *which type* of risk they are actually taking on in a given situation and end up avoiding any form of investment. Knowing, for instance, that a government bond from an economically solid country carries mainly market risk (its price can fluctuate) but very limited credit risk, allows for far more informed choices than simply thinking, in general terms, that "financial markets are dangerous". If risk cannot be eliminated entirely, it can nonetheless

be managed, and partly reduced. The main tool for doing so is diversification: not concentrating one's capital in a single instrument, but spreading it across multiple instruments, sectors or geographic areas, so that the negative performance of one does not jeopardise the entire portfolio. This is the principle behind the popular saying "don't put all your eggs in one basket", which finds solid theoretical grounding in finance and it is, not by chance, one of the concepts we will return to when discussing the robo-advisor case study, where diversification is measured concretely through the correlations between instruments.

The intuition behind diversification is simple: if capital is invested in a single company, an investor's entire wealth depends on the performance of that company. If that company goes bankrupt, the entire investment fell off. If, instead, the same capital is spread across hundreds of different companies, belonging to different sectors and geographic areas, it becomes far less likely that all of them will run into trouble at the same time. The difficulties of one company or sector tend to be partly offset by the positive performance of others. This is what, in the previous section, we called specific risk, which diversification can significantly reduce, although it does not eliminate market risk, which by its nature affects the entire system. For an ordinary investor, the main financial instruments available for building a diversified portfolio are:

- **Government bonds:** loans that the investor grants to a government, in exchange for periodic interest (a coupon) and the repayment of capital at maturity. Generally considered low-risk, especially for economically solid issuers, although the level of risk can vary considerably from country to country. Italy, for example, has historically always paid back its bonds.
- **Corporate bonds:** function in a similar way to government bonds, but are issued by private companies. The level of risk depends on the financial soundness of the issuing company: bonds from large, solid companies are generally safer, while those from weaker companies offer higher returns precisely to compensate for the greater credit risk.
- **Shares-stocks-equities:** represent ownership stakes in a company. Whoever buys a share becomes, to a small extent, a co-owner of the company, sharing in both its profits (through dividends and price appreciation) and its risks. They generally offer higher return potential over the long term, but also greater short-term volatility.
- **Mutual funds:** instruments managed by professionals (fund managers) who pool the capital of multiple investors and allocate it across a diversified basket of securities, according to a defined strategy. They allow even investors with limited capital to access a level of diversification that would be difficult to build on their own.

- **ETFs (Exchange Traded Funds):** funds that passively replicate a market index (for example, an entire equity or bond market), listed and traded on the stock exchange like shares. Unlike actively managed mutual funds, they do not seek to "beat the market" but to faithfully replicate its performance, generally at much lower management costs. These are the instruments at the centre of the case study analysed later in this work, precisely because of their ability to offer diversification at low cost and with high transparency.
- **Cash and deposit accounts:** the safest and most immediately available component of a portfolio. Although generating little or no return, they play an important role as a reserve for short-term needs, even though, as noted in the previous section, they are largely ineffective at protecting purchasing power from inflation over the long term.

It is important to stress that diversification, on its own, is not sufficient to build a portfolio suited to a given investor: the choice and combination of these instruments, and the proportions in which they are blended, are not arbitrary. They depend on two variables examined in the following sections:

- 1) The investor's time horizon
- 2) Their risk profile

Which turn a simple list of available instruments into a coherent investment strategy suited to each person's situation and objectives. The time horizon is the period for which an investor is willing, or needs, to keep their capital invested before having to use it. It is one of the most underestimated variables among those approaching investing for the first time, yet it is arguably the one that, more than any other, determines which instruments are suitable and which are not often more so than simple psychological risk tolerance

The General Principle: Time as a “buffer”

The general principle is as follows: the longer the time horizon, the greater an investor's ability to tolerate the short-term fluctuations of riskier instruments. Equity markets can go through difficult years, with significant losses: economic crises, geopolitical shocks, recessionary phases, but historically they tend to recover and generate positive returns over long periods (ten, twenty, thirty years). An investor with a long time horizon ahead can afford to go through these negative phases without being forced to sell at a loss, simply by waiting for the market to recover over time. Conversely, an investor with a short time horizon (a few months or a couple of years) cannot afford the same risk: if the market goes

through a downturn right at the moment when the funds need to be liquidated, the loss becomes real and permanent. This is where one of the most common mistakes among inexperienced investors lies: investing in risky instruments (such as shares) sums of money that will in fact be needed in the short term, basically picking the wrong financial asset for their needs. This explains why riskier instruments are generally better suited to long time horizons, while short-term goals call for more stable and liquid instruments. A useful way to think about this is to mentally associate each financial goal with its natural time horizon: a sum that will be needed in six months for an unforeseen expense should be kept liquid, or nearly so; capital set aside for a home purchase in five years can tolerate a moderate degree of risk; savings intended for retirement, with a horizon of twenty or thirty years, can instead afford a more substantial allocation to riskier instruments such as shares, precisely because the time available acts as a buffer against short-term fluctuations.

The importance of time is not limited to how long an investment is held, but also to *how* it is entered into. Rather than investing a lump sum all at once, many investors choose to spread their contributions over time investing a fixed amount at regular intervals (for example, monthly), regardless of whether the market is up or down at that particular moment. This approach is called the SAVING PLAN (or PAC in italian) and it's one of the most powerful tool we can use in investments, It does not require any special market expertise or attempt to predict short-term movements; it simply relies on time itself to smooth out the effect of entering the market at a particularly high or low point, since purchases are made across many different price levels rather than a single one. Exploiting the compound interest rate effect, our investment grows bigger and bigger over time following an exponential curve, creating a fortune over decades. It is a practical illustration of how time horizon and investment discipline work together, rather than as separate considerations.

This is not a rigid, automatic rule other factors also come into play, such as each individual's psychological risk tolerance, but it is a common-sense criterion that should guide every informed investment decision, and one that is often ignored by those who let themselves be guided more by the emotion of the moment (the euphoria of a rising market, the fear during a crash) than by genuine planning based on their own time frame. Even if the theory seems so easy and obvious the real challenge is to have the consistency and the discipline of be invested by 20 or more years, trusting the plan and ourself.

Bringing together everything discussed so far risk, return, diversification, time horizon leads to a central concept in personal finance: the investor's risk profile, that is, the combination of objective and subjective elements that determine what level of risk is suitable, and sustainable, for a given person. It is useful to distinguish between two components that together define the risk profile, as they are often confused with one another:

- **Risk capacity** is an objective, measurable dimension: it depends on factors such as income, overall wealth, job stability, age and, as discussed, time horizon. A person with a stable income, substantial wealth and a long time horizon has objectively more room to withstand the fluctuations of a risky investment than someone with limited resources or short-term liquidity needs
- **Risk tolerance**, on the other hand, is a subjective, psychological dimension: it concerns how much a person is emotionally willing to withstand fluctuations in the value of their investment without panicking and selling at the wrong moment. Two people with the exact same objective economic situation can have very different risk tolerances, shaped by personality, past experiences, and their level of familiarity with the sector.

A well-constructed risk profile takes both dimensions into account: offering a highly risky portfolio to someone who objectively has the financial capacity to sustain it, but who psychologically cannot tolerate it, almost inevitably leads to impulsive decisions during periods of market downtrend, panic selling at precisely the worst moment, turning a temporary loss into a permanent one.

From this distinction follows a fundamental principle: there is no universally portfolio, valid for everyone. There is a portfolio suited to a given risk profile, built around a person's specific situation. This is precisely the principle underlying the European MiFID regulation (Markets in Financial Instruments Directive), which requires anyone offering financial advisory services whether a human advisor or an automated platform to gather information about the investor before proposing any product, precisely to ensure this correspondence between profile and proposal, and to protect the investor from choices unsuited to their situation.

CHAPTER 2

Financial Education

Why People Don't Invest

Having established, in the previous chapter, that keeping savings idle carries hidden costs of its own, a natural question arises: if this is true, why do so many people still avoid investing altogether? The answer is rarely a single cause, but rather a combination of psychological, cultural and informational factors.

- 1) **Lack of financial literacy** is the most immediate explanation, but it is worth being precise about what this actually means. It is not simply "not knowing what a stock is or how investments work". More often, it means not having the basic conceptual tools discussed in the previous chapter. The distinction between saving and investing, the relationship between risk and return, the role of time horizon, which makes the whole subject feel difficult and intimidating rather than approachable. Without these tools, any decision in this area feels like an impossible step to make and the natural human response to the unknown is avoidance.
- 2) **Loss aversion** plays an equally important role. Research in this field has consistently shown that people feel the pain of a loss more intensely than the pleasure of an equivalent gain. Studies show that in order to psychologically compensate a certain amount of money lost, an investor has to make 3 times it's loss. Losing 100 euros feels considerably worse than gaining 100 euros feels good. Applied to investing, this asymmetry means that the mere possibility of seeing one's capital temporarily decrease in value can outweigh, psychologically, the prospect of a much larger gain over time, even when the numbers would suggest otherwise. This is one of the deepest reasons why so many people prefer the appearance of safety offered by an idle bank account over the appearance of risk associated with investing, even when that risk is statistically well understood and manageable.
- 3) **Status quo bias** explain why once an habit is established in this case, simply leaving money in a current account, changing it requires an active decision, a moment of friction that many people never quite get around to. Doing nothing feels safe abd easy by default, precisely because it requires no action and no confrontation with uncertainty.

- 4) **Cultural heritage** is a further, often underestimated factor, distinct from the psychological biases described above, though closely intertwined with them. In many countries, the relationship between households and financial markets is shaped less by individual reasoning than by inherited habits, transmitted across generations largely without question. Italy is a clear example of this pattern: the preference for real estate and bank deposits over financial instruments is not simply the sum of individual choices made independently by millions of households, but reflects a deeply rooted cultural template passed down from parents to children, in which "safety" is associated almost exclusively with owning a home and holding cash, while financial markets are culturally coded as something distant, speculative, or reserved for a narrow elite. This template took shape over decades in which financial markets genuinely were less accessible than they are today, but it has persisted well beyond the conditions that originally justified it. This pattern is not universal, and comparing it with other countries helps to show that it is cultural rather than simply a matter of income or education. In the Netherlands, Sweden, Denmark and other Northern European economies, participation in financial markets whether directly or through pension funds and long-term savings plans is a normalised, unremarkable part of financial life, embedded in the culture to the point that avoiding it would be seen as the unusual choice, rather than the default one. The same difference in cultural orientation, rather than in disposable income alone, is one of the reasons why the share of household wealth held in bank deposits varies so widely across European countries with otherwise comparable levels of economic development, as shown by the data discussed later in this chapter. This suggests that closing the gap in financial participation is not purely a matter of raising financial knowledge in the abstract, but also of shifting a cultural narrative that, in countries like Italy, has for decades treated investing as an exception rather than the norm.
- 5) **Distrust**, often rooted in memories of past financial scandals or crises, adds a further layer of resistance. For many people, particularly in countries with a history of banking crises or high-profile financial scandals, the financial sector as a whole can carry a reputation for being self-serving or opaque, which discourages engagement even with instruments, such as low-cost diversified ETFs, that have very little in common with the products at the centre of those past scandals.

The consequences of low financial literacy are not confined to the individual investor who misses out on potential returns; they extend to the wider economy. At the individual level, poor financial literacy is associated with a higher likelihood of falling into debt traps, making poorly informed borrowing decisions, being more vulnerable to fraud and mis-selling, and under-saving for retirement a particularly pressing issue in countries with ageing populations and increasingly strained public pension systems. At the collective level, when a large share of household savings remains parked in unproductive, low-yielding instruments rather than being channelled into productive investment, this represents a form of misallocated capital on a national scale. Capital that could, in principle, support the growth of companies, infrastructure or innovation instead sits idle, generating no economic value beyond the psychological comfort of its owner. This is one of the reasons why financial education is increasingly treated not merely as a matter of individual well-being, but as a matter of public interest, closely tied to a country's long-term economic competitiveness.

What the Data Tells Us

International comparisons consistently place Italy among the countries with comparatively low levels of financial literacy, particularly when measured against other advanced economies. Surveys conducted by organisations such as the OECD and by the Bank of Italy have repeatedly shown that a significant share of the Italian population struggles with basic financial concepts for instance, understanding compound interest, the effect of inflation on purchasing power, or the basic principle of risk diversification concepts that are directly relevant to the choices discussed in the previous chapter. This pattern is not unique to Italy, but it appears particularly pronounced there when compared with countries such as the Netherlands, Germany or the Nordic economies, where financial education is often integrated into school curricula from an early age. This gap is not without consequence: it correlates closely with the low share of Italian household wealth allocated to financial instruments beyond bank deposits, relative to other European countries with comparable income levels.

According to the triennial survey conducted by the Bank of Italy (IACOFI), based on the harmonised OECD/INFE methodology, the level of adult financial literacy in Italy increased slightly between 2020 and 2023, moving from 10.2 to 10.6 on a scale from 0 to 20, though remaining at low levels overall. Breaking this figure down into its components is particularly informative: the improvement was driven mainly by the financial behaviour dimension (from 4.2 to 4.7 on a scale from 0 to

9) and by attitudes (from 2.0 to 2.3 on a scale from 0 to 4), where the behaviour dimension considers the management of financial resources in the short and long term, while the attitude dimension assesses individuals' orientation towards saving and care in spending money. By contrast, the score relating to pure financial knowledge which includes concepts such as inflation, simple and compound interest rates and risk diversification remained broadly stable, moving only from 3.9 to 3.7 on a scale from 0 to 7. Encouragingly, there was a specific improvement in the understanding of inflation and its impact on households' purchasing power, partly attributable to the sharp rise in inflation experienced in the preceding two years, together with the financial education initiatives promoted by the Bank of Italy on this specific topic. [Bancaditalia](#)

This picture becomes even more revealing when placed in international context. The OECD/INFE 2023 International Survey of Adult Financial Literacy, which measures the percentage of adults reaching a minimum target score across a large sample of countries, shows Germany at the top of the ranking with 75% of adults reaching this threshold, followed by a long list of advanced and emerging economies; Italy appears very close to the bottom of the ranking, well below the OECD average and behind virtually all other major Western European economies. This gap is consistent with the findings of other sources: one OECD-based estimate suggests that only around 44% of Italian citizens demonstrated a level of financial literacy considered sufficient, a figure that closely echoes an earlier finding from the Standard & Poor's Global Financial Literacy Survey, which had placed the share of financially literate Italians at only about 37%, indicating that this is not a newly emerged problem, but a persistent structural gap that has changed only marginally over the better part of a decade. This low level of financial literacy correlates closely with the composition of Italian household wealth. According to the Bank of Italy's distributional wealth accounts, households in the bottom half of the wealth distribution hold more than 90% of their assets in dwellings (73.6%) and deposits (17.5%), while households in the upper wealth brackets hold considerably more diversified portfolios, with a significant share of financial instruments other than deposits. This pattern of modest wealth concentrated in property and idle bank deposits, with genuine diversification largely confined to wealthier households is broadly consistent with the picture across much of Southern and Eastern Europe: a European-wide analysis by EFAMA found that while households in Denmark, Sweden and the Netherlands hold less than 30% of their financial wealth in deposits, the share of deposits exceeded 70% in 2022 in countries such as Malta, Portugal, Lithuania, Bulgaria, Slovenia, Poland, Cyprus and Greece, and more broadly, European households as a whole increased their holdings of cash and bank deposits from €10,260 billion in 2015 to €13,944 billion in 2022, raising the deposit share of total financial wealth from 36.7% to 41.1% over the same period. Italy, while not the

most extreme case in this comparison, sits firmly on the "deposit-heavy" side of this European divide a pattern that mirrors, almost point for point, the financial literacy gap described above ([Oxford Academic](#)). Taken together, these figures support the argument outlined earlier in this chapter: the tendency to leave savings idle in Italy is not simply a matter of individual preference, but is closely tied to a measurable and persistent gap in financial knowledge one that recent, modest improvements in financial behaviour and attitudes have not yet been able to close.

It is worth stressing a point that is sometimes lost in discussions of financial education: its purpose is not only to encourage people to invest more, but also, and perhaps more importantly, to protect them from poor decisions both the decision to avoid investing altogether when it would clearly be beneficial, and the opposite mistake of investing carelessly, without understanding the risks involved. A well-informed investor is better equipped to recognise unrealistic promises (guaranteed high returns, "risk-free" opportunities), to understand the fees and structures of the products being offered, and to avoid emotionally driven decisions during periods of market turbulence. In this sense, financial education functions as a form of self-protection that complements, rather than replaces, the regulatory protections (such as MiFID) already in place at an institutional level. Despite these challenges, it is worth noting as a bridge toward the next chapter that the barriers to entry into investing have fallen dramatically over the past decade, largely thanks to digital technology. Retail investors can finally have a complete access to the financial market easier than ever before. Minimum investment amounts have shrunk, in some cases to just a few euros; transaction costs, once a significant deterrent for small investors, have in many cases been reduced close to zero; and the interfaces through which people access financial markets have become dramatically simpler and more intuitive than those of traditional banking channels. This can represent a great opportunity for the ones that want to know how it works, but can also represent a trap for the ones that believe into an "easy way to get rich". In fact this technological shift does not, by itself, solve the problem of financial literacy arguably, in some respects, it makes financial education even more necessary, precisely because it has become so easy to invest that one can do so without truly understanding what one is doing. But it does mean that the traditional argument "investing is not for people like me" is, today, considerably weaker than it once was on purely practical grounds. This is precisely the space occupied by robo-advisors and self-directed investment platforms, examined in the next chapter, which represent two different technological responses to the same underlying accessibility problem.

A Personal Perspective: My Own Experience

Beyond the theoretical and statistical arguments discussed in this chapter, it is worth grounding this discussion in direct, lived experience, my own path into investing, which began in June 2024 and has, over roughly two years, shaped much of the perspective behind this thesis.

How it started. My interest in investing was born out of curiosity, partly triggered by the kind of advertising for investment platforms that has become common online, and partly out of a more personal concern: an awareness, consistent with the argument made in the first chapter of this work, that money left idle in a bank account loses value over time because of inflation. This realisation pushed me to look for a way not just to protect the value of my savings, but ideally to grow it. What followed was a period of self-directed study and research, helped by the fact that today's investor has access to an overwhelming amount of information online. The real challenge, in fact, was not a lack of information but the opposite: learning to distinguish which sources and which concepts were actually worth the time to understand, and which were noises.

Choosing a strategy. After this period of research, I identified a small number of platforms that offered a reasonable combination of security, regulatory reliability and ease of use, and settled on a systematic investment plan (called SAVING plan) as the strategy that best matched both my risk tolerance and my objective. That objective was, and remains, explicit: building a personal retirement reserve independently, given the widespread uncertainty in Italy around the long-term sustainability of the public pension system (INPS). One of the platforms I used, Trade Republic, is worth briefly describing for context: "it is a German neobroker, headquartered in Berlin, operating under a full banking license granted by the European Central Bank and supervised by BaFin, Germany's federal financial regulator a regulatory profile that provided a meaningful degree of reassurance at a stage when trusting an unfamiliar digital platform with real money still felt like a significant psychological step".

Starting small and scaling up. I began with a deliberately small amount, €50, mainly to understand how the mechanics of the platform and the investment process actually worked before committing more seriously. I then set myself an initial strategy of contributing roughly €100 per month, with a time horizon of thirty years. As I gained confidence and a deeper understanding of both the instrument and the broader market, I gradually increased my monthly contribution, eventually reaching €350 per month. As of today, this has resulted in a total invested capital of just over €10,000, with a net profit of just under €3,000 generated over roughly two years. A result that, while modest in absolute terms given the short time

frame, has been personally meaningful in demonstrating that the theoretical mechanics described earlier in this thesis do translate into real outcomes.

Testing other instruments. Alongside this core strategy, I also experimented, cautiously and with limited amounts, with individual stocks and with cryptocurrencies, largely as a way of directly experiencing the difference in risk and behaviour between these instruments and a diversified ETF-based plan.

Lessons from volatile periods. This period was not without turbulence. I experienced, first-hand, episodes of significant market stress including: the market disruption caused by the Trump administration's tariff announcements in 2025, and more general pullbacks of around 10%, which are a normal, if uncomfortable, feature of equity markets. Living through these episodes, rather than only reading about them, reinforced a conclusion that I consider central to this thesis: a systematic investment plan built around a diversified ETF remains, in my experience, the most practical, reassuring and manageable solution available to an ordinary investor. This is true even though the underlying instrument a globally diversified equity ETF is, by its nature, a volatile asset; it is precisely the discipline of the systematic, long-horizon approach, combined with genuine diversification, that made these periods of volatility manageable rather than alarming.

A comparison across instruments, based on direct experience. Having tested all three approaches personally, I can offer a first-hand comparison that complements the more theoretical distinctions drawn earlier in this thesis. For an investor who is hesitant about investing, uncertain about how much risk they are willing to accept, or simply new to the subject, an ETF-based systematic investment plan with a time horizon beyond twenty years remains, in my view, the most suitable starting point. Individual stocks, by contrast, proved considerably more volatile in practice and demand a level of ongoing study and attention that goes well beyond what a systematic ETF plan requires. Cryptocurrencies are, in my assessment, a different matter entirely: without a genuine periodic strategy and disciplined plan behind their use, i would not recommend they be given serious consideration by an inexperienced investor.

Why people hesitate and a suggestion for families. Having shared this experience with people close to me, some of whom I encouraged to start investing themselves, I have come to understand more concretely why so many people, particularly in Italy, find this first step difficult. Trusting an unfamiliar platform, and watching a fixed sum leave the bank account every month to be transferred into an investment account rather than a traditional savings account, is not an easy psychological step to take especially for households that are

already financially established and set in their habits. For this reason, one suggestion I would make, particularly to families already inclined toward saving, is to consider setting up a saving plan for a child from birth, managed by the parents in the early years and eventually handed over once the child is mature enough to manage it independently. Approached this way, investing becomes not just an individual financial decision but a way of building financial wealthness across generations.

Investing starts with saving. Before any of this becomes relevant, however, there is a more basic point worth making explicit, one that I consider a precondition rather than a detail: investing can only begin once genuine savings exist, and genuine savings can only be built by first controlling costs distinguishing clearly between essential expenses, which cannot reasonably be avoided, and discretionary ones, which can. Only once this distinction is made does it become possible to create a consistent buffer of liquidity that can then be directed toward investing. A simple budgeting framework that I have found useful in this respect, and that I would suggest to anyone approaching this subject for the first time, is the 70/20/10 rule, applied on a monthly basis to one's net income: roughly 70% covers essential expenses (rent, food, utility bills, condominium fees and similar unavoidable costs); 20% is directed toward investing; and the remaining 10% is set aside for discretionary personal spending. It is worth stressing that these percentages are not a rigid, they are a starting point, to be adjusted according to one's own income and expenses. A different split, such as 60/30/10, may well be more appropriate depending on individual circumstances. What matters is not the specific numbers, but the underlying discipline: treating saving, and by extension investing, as a planned and structured allocation of income, rather than as whatever happens to be left over at the end of the month which, in practice, for many people, tends to be very little.

Starting from the goal, not from the instrument. Finally, one practical piece of advice that I believe is worth including in this thesis, precisely because it reflects what I wish I had understood from the very beginning, is this: rather than starting from the instruments themselves, it is far more effective to start from the objective. A range of tools and platforms now exist that allow an investor to project, based on realistic assumptions, what a given monthly contribution could realistically become over a horizon such as thirty years. One such tool that I have used personally is justETF, a website offering an ETF savings plan calculator: by entering a monthly contribution, an investment horizon, an expected market return and the relevant costs, it produces a simulation of how the invested capital could develop over time, net of fees very similar in spirit to the simulation presented earlier in this thesis. Once a clear objective has been defined in this way, the practical questions how much to invest, how frequently, and through

which instrument tend to follow naturally, rather than being decided somewhat arbitrarily.

Two Ways to Enter the Market: PIC and PAC

Before turning to the practical question of *what* to invest in, it is worth pausing on an equally important question that is too often overlooked: *how* to enter the market in the first place. In Italian financial vocabulary and, increasingly, in this thesis, this choice is usually described using two acronyms: **PIC** (Piano di Investimento del Capitale), a lump-sum investment made all at once, and **PAC** (Piano di Accumulo del Capitale), a systematic investment plan built through fixed, periodic contributions over time already introduced briefly in Chapter 1 under its English name, the "saving plan".

Both approaches rest on the same underlying principle discussed throughout this thesis, the power of compounding over a long time horizon, but they expose the investor to that principle in structurally different ways. A **PIC** puts the entire sum to work immediately: statistically, since markets rise over the large majority of multi-year periods, this tends to maximise the amount of time capital spends invested, and therefore, on average, the expected final return. The cost of this advantage is a concentrated timing risk: the investor's outcome depends disproportionately on the specific market conditions prevailing at the single moment the lump sum is committed a genuinely unlucky entry point, close to a market peak, can weigh on the result for years afterwards.

A **PAC** addresses this timing risk directly, by spreading purchases across many different price levels rather than committing to a single one an approach commonly known as dollar-cost averaging. This does not eliminate market risk, but it smooths it, and it carries an additional, practical advantage that is arguably more relevant to most ordinary investors than any statistical nuance: a PAC does not require the investor to already possess a large sum of capital. It can be built directly out of monthly income, which is precisely why, as the real portfolios presented in the next chapter will show, it tends to be the natural, and often the only realistic, starting point for a young investor without substantial savings already accumulated.

Neither instrument is inherently superior; as with several of the choices examined in the next chapter, the right answer depends on the investor's starting point, income stability and psychological relationship with risk. What follows, in Chapter 3, is a concrete, thirty-year quantitative illustration of how these two paths compound in practice.

CHAPTER 3

Modern Investment Tools, Delegation and Autonomy

Delegated and Self-Directed Investing

Having established, in the previous chapters, that investing has become considerably more accessible than it once was, it is worth asking: accessible *how*, exactly? Today's investor essentially faces two distinct paths, both enabled by the same underlying technological shift, but built on very different philosophies.

The first path is **delegated investing**, essentially robo-advisors: automated platforms that build and manage a portfolio on the investor's behalf, based on a risk profile derived from a questionnaire. The investor answers a series of questions, and the algorithm does the rest selecting instruments, allocating capital, and often rebalancing the portfolio over time.

The second path is **self-directed investing**: platforms such as Trade Republic, Revolut, Scalable Capital, or the trading sections increasingly offered by traditional banks, which give the investor direct control over what to buy, how much, and when, without an algorithm making the decisions for them.

Both paths are, in a sense, answers to the same underlying problem discussed in the previous chapter: financial markets used to be difficult to access, whether because of high minimum investments, steep transaction costs, or the need to go through a traditional intermediary. Technology has addressed this problem from two different angles one by automating the decision itself, the other by making direct access radically simpler and it is worth examining both in detail, because they carry different implications for the investor's need for financial education.

A robo-advisor is a digital platform that provides automated, algorithm-driven financial planning and portfolio management, typically with minimal human intervention. The process usually follows a well-defined sequence: the investor completes an online questionnaire, modelled on the MiFID requirements, covering aspects such as investment objectives, time horizon, income, existing wealth, and risk tolerance; the platform's algorithm processes these responses to assign the investor to a risk class; and a portfolio, typically composed of low-cost

diversified instruments such as ETFs, is then automatically constructed and, in many cases, periodically rebalanced to maintain the target allocation over time.

The appeal of this model lies in three main advantages:

1) Lower costs: without a human advisor to compensate for each individual client relationship, robo-advisors can typically charge substantially lower management fees than traditional advisory services.

2) Lower minimum investments: many robo-advisors allow investors to start with very small amounts, compared with the much higher thresholds often required for traditional discretionary portfolio management.

3) Simplicity and speed: the entire onboarding process, from account opening to a fully constructed portfolio, can often be completed in a matter of minutes, entirely online.

These advantages directly address several of the barriers to entry such as cost, complexity, and the perception that investing "is not for people like me". In this sense, robo-advisors can reasonably be described as one of the clearest practical expressions of financial democratization made possible by technology.

The Trade-Off Behind the Simplicity: A Case Study

However, this same simplicity conceals a structural trade-off, which is precisely the subject investigated by Giudici and Polinesi (2019) in the study at the centre of this chapter. Their research starts from a simple but important observation: the risk classification models used by many robo-advisors tend to be considerably simpler than those employed by traditional, human-led financial consultancy precisely, because speed and automation are part of what makes the service attractive in the first place. This creates a tension: the very features that make robo-advisors accessible are the same features that can make their risk assessment less rigorous, with the concrete danger that an investor ends up holding a portfolio that does not truly match their risk profile.

To investigate this problem empirically, the authors worked with real data supplied by a European robo-advisor: responses to a MiFID questionnaire from 139 investors, together with the daily returns of fifty-two ETFs observed over a period of nearly two years. Their approach unfolds in two complementary parts, corresponding to the two sides of the risk-matching problem the investor's side and the product's side.

On the investor side: the authors applied a statistical technique called correspondence analysis to the questionnaire responses. In simple terms, this method allowed them to identify which questions in the questionnaire were genuinely useful in distinguishing between more cautious and more risk-tolerant investors, and which questions, despite being asked, added little information because almost everyone answered them the same way. Out of fourteen questions in the original questionnaire, only three turned out to be truly discriminating: the investor's objective, their desired time horizon, and their stated risk tolerance, which are the very three concepts discussed in the first chapter of this work. Using only these three questions, the authors built a simple scoring system and grouped investors into four risk classes of increasing risk.

On the product side: the authors applied a hierarchical clustering technique to the historical returns of the fifty-two ETFs, using a measure of partial correlation between assets as a distance metric. This allowed them to group the ETFs into six homogeneous risk clusters, verified through a diversification index based on Value at Risk (VaR), essentially a way of checking how genuinely diversified (as opposed to merely numerous) the assets within each cluster actually were.

The key finding, and the one most relevant to the argument of this chapter, emerged when the authors compared their own investor risk classification with the one actually used by the robo-advisor providing the data: the two classifications matched in only 65% of cases. In other words, more than a third of investors were placed, by the platform's own methodology, in a risk class different from the one that a more rigorous statistical analysis of the same questionnaire data would have suggested. This is a striking result, precisely because it does not stem from any obvious error or malicious intent on the platform's part, but simply from the tendency described earlier in this chapter for automated systems to favour simplicity over depth in their risk classification methodology.

The authors then extended their analysis using feed-forward neural networks, showing that machine learning methods could, in principle, be used to automate both the investor classification and the asset classification in a way that could generalise to new investors or new financial products not included in the original dataset a theme we will return to, in a more applied form, later in this thesis. The relevance of this study to the argument of this thesis goes beyond its technical content. It offers concrete, empirical evidence for a point that might otherwise remain purely theoretical: automation and financial education are not substitutes for one another. A robo-advisor can lower the practical barriers to investing: cost, complexity, time, but it does not guarantee that the resulting portfolio genuinely matches the investor's needs. This burden of verification cannot be delegated

entirely to the algorithm; a certain degree of understanding on the part of the investor remains valuable even within a fully automated service. If only to be able to recognise whether the proposed portfolio actually makes sense given one's own circumstances. This is an important nuance for the argument of this work: financial technology and financial education, far from being alternative solutions to the same problem, function best as complements to one another.

Self-Directed Investing: The Rise of Direct Access

The second path, self-directed investing, has undergone an equally dramatic transformation over the past decade, driven largely by so-called neobrokers: digital-native platforms built around mobile apps, radically simplified interfaces, and cost structures that would have been unthinkable for a retail investor twenty years ago.

Several structural changes explain this shift.

1) Fractional shares allow investors to buy a portion of an expensive stock rather than a full share, removing what was once a significant barrier for smaller investors.


2) Commission structures have been driven down dramatically, in many cases to zero, removing the friction that historically discouraged frequent or small transactions.

3) Minimum investment amounts have shrunk to a few euros in many cases, making it possible to begin a systematic investment plan of the kind simulated in the chart discussed earlier in this work.

4) User-friendly interfaces have been redesigned around simplicity and immediacy, often drawing deliberately on the visual language of consumer apps and social media, rather than the more austere interfaces of traditional banking platforms.

A Practical Look at My Own Portfolio

To make the reflections in this chapter concrete, it is worth presenting an actual snapshot of my own portfolio, as held on Trade Republic. The figures below correspond to the screenshots included in this section.



Overall allocation. As the imagine 1 can show, my total portfolio was worth **€13,340.84**, split between three components: a securities account (*Conto titoli*) worth €11,769.21, representing 88% of the total, and a crypto wallet worth €1,571.63, representing the remaining 12% and the remain is the liquidity not invested yet. This split itself reflects a deliberate choice discussed earlier in this thesis: the large majority of the portfolio sits in the more predictable, long-horizon core (the securities account), while the crypto component is treated explicitly as a small, marginal experiment rather than a core holding.

The core ETF position. Within the securities account, the largest single holding and the one that corresponds directly to the systematic investment plan discussed throughout this thesis is the Vanguard FTSE All-World UCITS ETF (USD) Accumulating (ticker: VWCE; ISIN: IE00BK5BQT80). This is a physically replicated, accumulating ETF domiciled in Ireland, launched in July 2019, which tracks the FTSE All-World Index a broad benchmark covering large- and mid-cap

companies across both developed and emerging markets. Its total expense ratio (TER) is 0.19% per year. This position alone was worth €8,356.16, accounting for roughly 71% of the securities account. It is worth distinguishing two different figures shown in the app for this ETF: the fund's price is up 71.02% since its own inception in 2019, which reflects the ETF's overall market performance since launch and is not specific to me; my own personal return on this position, since I began investing in it in June 2024, is +21.69%, a figure that reflects the actual performance of my own systematic contributions over a much shorter period.

Individual stock positions. The remaining portion of the securities account (about 29%, worth €3,416.19) is allocated to a small number of individual stocks, held deliberately as a limited, as discussed earlier in this chapter:

- **NVIDIA** is a US semiconductor company, the dominant designer of graphics processing units (GPUs), which have become the core hardware underpinning modern artificial intelligence systems; it is one of the largest companies in the world by market capitalisation.
- **Alibaba Group (ADR)** is a Chinese multinational conglomerate, best known for its e-commerce platforms (Taobao, Tmall) but with significant operations in cloud computing and digital payments as well; the "ADR" (American Depositary Receipt) simply indicates that the shares are traded on a US exchange in the form of a depositary receipt rather than directly on a Chinese exchange.
- **BTQ Technologies** is a much smaller, more speculative holding: a Canadian quantum-technology company developing post-quantum cryptography solutions that is, security techniques designed to protect digital infrastructure (including blockchain networks) against the future threat posed by sufficiently powerful quantum computers.


These three positions illustrate, in practice, the point made earlier in this chapter about the greater volatility of individual stocks compared with a diversified ETF, a level of single-day fluctuation that a broad, diversified ETF such as VWCE essentially never experiences, precisely because gains and losses across hundreds of underlying companies tend to offset one another.

The crypto wallet. The remaining 12% of the portfolio is held in a crypto wallet, currently allocated entirely to Ethereum. This position was down approximately 8.3% over the period shown, and the "Max" (all-time) chart for this holding shows swings as large as -23% from its recent peak a level of volatility considerably more extreme than anything experienced by the ETF or even the individual stocks discussed above. This is consistent with the reflection offered earlier in this thesis:

without a genuine periodic investment strategy behind it and given this level of volatility, I have deliberately kept this position as a small, probably this position will increase overtime, but will be always around 10-ish % of the total portfolio's value.

Now we analyse some real cases

Case 1:




Portfolio Overview case 1

21 years old with no income and seasonal jobs only, invested since Genuary 2025. With a total value of €4,075.85, up 9.23% since date, the portfolio is split between ETFs (57.6%, €2,347) and individual stocks (42.4%, €1,728). The ETF component is built through two separate systematic investment plans (PAC): one on the MSCI EM ESG Enhanced (USD, Acc) ETF, which tracks emerging-market

companies with strong environmental, social and governance ratings; and one on the Core Stoxx Europe 600 (EUR, Acc) ETF, which tracks large-, mid- and small-cap companies across the European market. Together, these two positions give the ETF portion of the portfolio exposure to both emerging and developed markets, rather than concentrating it in a single geographic area a practical example of the diversification principle discussed in Chapter 1.

The MSCI EM ESG position, which receives a monthly contribution of €70, as well as the Euro Stock 600 and the single stock NVIDIA, currently stands at €1,354.20, and has generated a personal return of +28.61% since it was started.

case 2:



Portfolio Overview case 2


27 years old employee with a stable job and a mid source of income, invested since September 2024. The third portfolio presented offers a useful contrast to the two previous cases, illustrating a markedly more aggressive approach. With a total value of €43,744.65, the portfolio is split between stocks (68.5%, €29,980.84, up 12.62%) and a cryptocurrency wallet (31.5%, €13,763.81, currently down 19.25%).

Within the securities account, the large majority (98.5%) is held in three diversified ETFs: a global equity fund (FTSE All-World), a US technology-heavy index (NASDAQ100), and an emerging-markets fund (Core MSCI EM IMI). Alongside these core holdings, there is a small satellite allocation (1.5%) to speculative individual stocks in the quantum-computing sector, BTQ Technologies, D-Wave Quantum, and QMC Quantum Minerals, all three currently at a loss.

The cryptocurrency wallet, split between Bitcoin and Ethereum, represents nearly a third of the total portfolio and is experiencing a substantial drawdown at the time of observation.

Taken together, this portfolio illustrates a materially higher risk tolerance than the two previous cases: a larger allocation to volatile assets (crypto), a concentrated thematic bet within equities (quantum computing), and comparatively less diversification relative to the two other portfolios discussed in this thesis.

Case 3:



Portfolio Overview case 3

A fourth investor's portfolio: the entry-level case. The fourth investor is a 25 years old master's student with no stable income and part-time jobs only, invested since May 2025. With a total value of €552.59, up 10.33% since started, the entire portfolio consists of a single position: the Vanguard FTSE All-World (USD, Acc) ETF, accumulated through a modest systematic investment plan of €100 per month, which has generated a personal return of +10.52% to date.

Unlike the three portfolios discussed previously, all of which combine multiple instruments, and in two cases include individual stocks or cryptocurrency, this portfolio represents full diversification achieved through a single instrument, with no additional risk layered on top. It serves as a useful reference point at the conservative, low-complexity end of the spectrum illustrated by the four portfolios presented in this thesis, showing that a meaningful and genuinely diversified starting position requires neither a large monthly contribution nor multiple products.

The Other Side of Accessibility: Risks Without a Safety Net

Self-directed investing removes an important intermediary step: unlike with a robo-advisor, there is no algorithm standing between the investor and their decisions, checking whether a given trade is consistent with their stated risk profile or investment horizon. This has an important consequence for the argument of this thesis: in self-directed investing, financial education is not merely useful, it becomes the only safeguard the investor has.

Several behavioural risks are particularly relevant here, and are worth naming explicitly, since they connect directly to concepts introduced earlier in this work:

- **Home bias:** a well-documented tendency to over-concentrate a portfolio in domestic or otherwise familiar companies, at the expense of the diversification discussed in the first chapter, simply because they feel more familiar or trustworthy.
- **Excessive concentration:** buying a small number of individual stocks rather than diversified instruments such as ETFs, often driven by conviction or excitement about a specific company, which undermines the risk-reduction benefits of diversification explained earlier.
- **Trend-chasing ("FOMO"):** entering a position because its price has recently risen sharply, driven by the fear of missing out rather than by any underlying analysis a pattern that tends to be self-reinforcing.
- **Time horizon mismatch:** as discussed in the first chapter, investing funds that will actually be needed in the short term into instruments suited to long time

horizons. A robo-advisor's initial questionnaire is at least designed to catch this mistake; a self-directed investor, by contrast, makes this judgement entirely on their own.

- **Gamification and overtrading:** interface design elements push notifications, price alerts, a general aesthetic that can encourage frequent trading even though frequent trading has been consistently shown, in the empirical finance literature, to reduce net returns for the average retail investor, largely due to poor timing, the cumulative effect of transaction costs, and a lack of competence.

None of this is to suggest that self-directed investing is inherently reckless, or that delegated investing is inherently safer; both models have legitimate places depending on an investor's knowledge, interests and objectives. The point, rather, is that the type of risk shifts. With a robo-advisor, the main risk lies in the (often invisible) simplification of the underlying risk-matching methodology, as shown by the case study discussed earlier in this chapter. With self-directed investing, the main risk lies in the investor's own behaviour, unmediated by any structural safeguard.

A practical comparison

Dimension	Robo-advisor	Self-directed investing
Who decides	Algorithm, based on risk profile	The investor, directly
Typical costs	Low management fee (ongoing)	Low or zero commission per trade
Minimum investment	Often very low	Often very low
Diversification	Built in by design	Depends entirely on investor's choices
Risk-matching safeguard	Present, but potentially simplified (see case study)	Absent — relies entirely on investor knowledge

Dimension	Robo-advisor	Self-directed investing
Main risk	Misclassification of risk profile	Behavioural biases, overtrading, poor diversification
Financial education required	Useful, but not strictly essential	Essential

Toward a Balanced View

What emerges from this chapter is not a verdict in favour of one model over the other, but a more nuanced conclusion that will guide the rest of this work: technology has genuinely lowered the barriers to entry into investing, on both the delegated and the self-directed side, largely resolving the "practical accessibility" problem discussed in the previous chapters. What it has not resolved and arguably cannot fully resolve on its own is the need for a baseline level of financial understanding on the part of the investor. Whether to correctly interpret and question the output of an algorithm, as shown by the robo-advisor case study, or to make sound decisions independently, as required by self-directed investing. This is the thread that connects this chapter back to the central argument of this thesis, and that will be developed further in the empirical chapter that follows.

PIC vs. PAC: A Thirty-Year Projection

Chapter 2 introduced the distinction between PIC (a lump-sum investment made all at once) and PAC, a systematic investment plan built through fixed periodic contributions. It is worth making this distinction concrete with actual numbers, for two reasons. First, it anticipates a pattern that recurs throughout the real portfolios presented later in this chapter, all four of which were built, in whole or in part, through a PAC. Second, it offers a direct, quantitative illustration of the "time as a buffer" principle introduced in Chapter 1. Two realistic paths are compared below, corresponding to two different starting points an investor might reasonably find themselves in. Both assume an average annual return of 7% a conservative long-term assumption for a globally diversified equity ETF of the kind discussed throughout this thesis, compounded monthly over a thirty-year horizon, gross of fees and taxation, in the same illustrative spirit as the justETF simulation mentioned in the previous chapter.

1. **Path A is a mixed strategy:** an initial lump-sum investment (PIC) of €10,000, followed by a systematic monthly contribution (PAC) of €100 for thirty years. This is representative of an investor who begins with an existing sum savings already accumulated, or an inheritance, for instance, and then continues to build on it gradually out of income.
2. **Path B is a pure PAC strategy:** a systematic monthly contribution of €250 for thirty years, with no initial lump sum. This is representative of an investor without existing capital, but with a stronger and more consistent monthly saving capacity, typical of the four real portfolios presented later in this chapter, none of which began from a large lump sum.

PIC vs. PAC thirty-year growth projection, assuming a 7% annual return.





Table. PIC vs. PAC key figures after a thirty-year horizon

Metric	Path A	Path B
Total capital invested (30 yrs)	€46,000	€90,000
Projected portfolio value (yr 30)	€203,162	€304,993
Total investment growth	€157,162	€214,993
Value generated per €1 invested	€4.42	€3.39

Capital invested versus investment growth after thirty years, Path A and Path B



The comparison illustrates two related but distinct points. First, in absolute euro terms, Path B ends considerably ahead of Path A, €304,993 versus €203,162, but this gap is driven almost entirely by the difference in total capital committed by the investor (€90,000 versus €46,000), not by any inherent superiority of one method over the other. Second, and more subtly, each euro invested in Path A ultimately generated more value than each euro invested in Path B (€4.42 versus €3.39 per euro contributed), precisely because Path A's €10,000 lump sum was exposed to market growth for the full thirty years from day one, while Path B's capital entered gradually, euro by euro, with the average contribution spending considerably less time invested. This is the quantitative expression of the “time as a buffer” principle introduced in Chapter 1: it is not only how much capital an investor commits that matters, but how early that capital begins compounding.

This is also why, as the four real portfolios that follow illustrate, a PAC without a large initial lump sum remains a perfectly sound strategy for an investor who lacks capital but not consistency: a missing lump sum can, over a long enough horizon, be substantially compensated for by a higher and steadier monthly contribution, as Path B demonstrates. It is worth stressing that both projections rest on a simplified, constant-return assumption; real markets are considerably more volatile, and the actual sequence of returns can meaningfully affect the final outcome, particularly for the lump-sum component of Path A. The purpose of this illustration is not to predict an exact outcome, but to make the underlying mechanics of compounding and the practical trade-off between the two entry strategies concrete, in the same spirit as the real, PAC-built portfolios presented next.

CHAPTER 4

Empirical part/model construction

Data source and sample

This chapter uses microdata from the Survey on Household Income and Wealth (SHIW), conducted by the Bank of Italy (Indagine sui Bilanci delle Famiglie Italiane, edizione 2022), publicly released on 27 October 2025. The full dataset covers 9,641 households and 23,057 individuals. Following standard practice with this survey, the analysis is restricted to the household reference person (*Persona di Riferimento*, identified by PARENT = 1 in the CARCOM22 file), yielding one observation per household (n = 9,641).

Data were merged across four files on the household identifier NQUEST: CARCOM22 (socio-demographic characteristics), RFAM22 (household income), RICFAM22 (household wealth), and Q22C2 (financial asset composition and risk attitude, Section C of the questionnaire).

Variable	Source file	Description	Coding
ETA	CARCOM22	Age (years)	Continuous
STUDIO	CARCOM22	Educational attainment	1 = none ... 8 = post-graduate specialization
QUAL	CARCOM22	Occupation status	1 = blue-collar, 2 = white-collar, 3 = manager, 4 = entrepreneur/self-employed professional, 5 = other self-employed, 6 = retired, 7 = other non-employed
Y	RFAM22	Net disposable household income (€)	Continuous
W	RICFAM22	Net household wealth (€)	Continuous
RISFIN	Q22C2 (Q. C32)	Self-reported risk attitude	1 = high risk/high return ... 4 = no risk/low return

Outcome variable

Following the logic of Giudici and Polinesi (2019, discussed in Chapter 3), households were classified into three ordered risk categories based on the actual composition of their financial portfolio (Section C of the questionnaire):

- **Conservative:** holdings limited to bank/postal deposits, repos, or government bonds (BOT/BTP/other)
- **Moderate:** additionally holds asset management products, corporate/bank bonds, or mutual funds
- **Growth-oriented:** additionally holds ETFs, listed/unlisted shares, foreign securities, or crypto-assets

Resulting distribution: Conservative 76.5% (n=7,380), Moderate 14.1% (n=1,355), Growth-oriented 9.4% (n=906) a class imbalance consistent with the low equity-market participation documented for Italian households in Chapter 2

The ordinal logistic regression: results


This ordinal logistic regression was estimated on a training sample of 7,230 households (75% of the full sample of 9,641), with the remaining 2,411 households (25%) held out as a test set for the performance evaluation presented later in this section, predicting the three-level investor risk profile (Conservative < Moderate < Growth-oriented) from age, education, occupation, household income, household wealth, and self-reported risk tolerance.

Table 1. Ordinal logistic regression, coefficient estimates

Variable	Coef.	Std. err.	z	p-value	95% CI
Age	0.009	0.003	2.75	0.006	[0.003, 0.016]
Education level	0.297	0.021	14.37	<0.001	[0.257, 0.338]
Income (per €10,000)	0.063	0.007	8.90	<0.001	[0.049, 0.076]
Wealth (per €10,000)	0.0004	0.000	1.57	0.117	[-0.000, 0.001]
Risk tolerance	0.647	0.039	16.52	<0.001	[0.570, 0.724]
Blue-collar (vs. retired)	-0.774	0.171	-4.52	<0.001	[-1.110, -0.439]
Entrepreneur/self-employed prof. (vs. retired)	0.072	0.108	0.67	0.506	[-0.139, 0.282]
Manager (vs. retired)	0.408	0.136	3.01	0.003	[0.142, 0.674]
Other non-employed (vs. retired)	-1.126	0.178	-6.33	<0.001	[-1.475, -0.778]
Other self-employed (vs. retired)	-0.167	0.199	-0.84	0.401	[-0.557, 0.223]
White-collar (vs. retired)	0.068	0.117	0.58	0.564	[-0.162, 0.298]

Model fit: Log-likelihood = -4,191.9; AIC = 8,410; n = 7,230

Figure 1



presents the same estimates graphically, ordered by magnitude, making it easier to identify at a glance which predictors carry the strongest and most reliable effects.

Discussion

The results are, for the most part, consistent with the theoretical framework developed in Chapters 1 and 2 of this thesis. Self-reported risk tolerance is, by a clear margin, the strongest predictor in the model, echoing the central role that this same type of question plays in the MiFID-style questionnaire discussed in Chapter 3. Education is the second-strongest predictor and highly significant, lending empirical support to the link between financial literacy and investment behaviour argued throughout Chapter 2. Income is also positive and highly significant, though its effect size is considerably smaller than that of education or risk tolerance suggesting that, in this sample, *how* a household thinks about risk and *how much formal education* it has matter more for portfolio composition than income alone.

Occupation also plays a meaningful role, largely in the expected direction: households headed by a manager are significantly more likely to hold a higher-risk portfolio than the reference group (retirees), while blue-collar workers and the

"other non-employed" category are significantly less likely to do so the latter being, in fact, the single strongest negative effect in the model.

Two results are worth flagging as noteworthy rather than simply expected. First, household wealth is not statistically significant ($p = 0.117$) once income and the other predictors are held constant a somewhat counter-intuitive finding, since wealth is often assumed to expand an investor's risk *capacity*, in the terminology introduced in Chapter 1. One plausible explanation is that, in this sample, wealth is dominated by real estate (as shown in Chapter 2), which is largely uncorrelated with financial risk-taking, thereby diluting its explanatory power once income is already accounted for. Second, the coefficients for white-collar employees, and for self-employed professionals/entrepreneurs, are not statistically distinguishable from the retiree baseline

suggesting that, once risk tolerance and education are controlled for, occupation alone adds comparatively little additional information for these particular categories.

Table 2. Confusion matrix (test set, $n = 2,411$)

Actual \ Predicted	Conservative	Moderate	Growth-oriented
Conservative	1,813	0	33
Moderate	309	0	30
Growth-oriented	171	0	55

Shaded cells indicate correct classifications (diagonal). Overall accuracy = 77.5%; balanced accuracy = 40.8%.

Overall accuracy on the test set is 77.5%. Taken at face value, this figure appears reasonably strong; however, it conceals an important limitation, visible clearly in Table Y: the model never predicts the "Moderate" category all 339 households in this class are misclassified, the overwhelming majority (309) as "Conservative". Precision and recall for the Moderate class are therefore both 0, despite the class

representing 14% of the sample. This is a known challenge with linear ordinal models applied to imbalanced, three-category outcomes: the "Moderate" class, sitting between two more clearly separated categories and containing comparatively little independent signal in the available predictors, tends to be absorbed by its neighbours rather than identified as a distinct category in its own right an outcome that plain accuracy, dominated by the large Conservative class, does not reveal on its own.

This finding is itself informative for the broader aims of this thesis: it illustrates that even a well-specified, theoretically grounded model can struggle with the practical task of *fully* separating investor profiles when one category is both a numerical minority and conceptually a "middle ground". This motivates the comparison, in the following sections, with tree-based models that may handle this kind of intermediate, imbalanced category differently.

The simple neural network: results


Following the same training/test split used for the ordinal logistic regression (7,230 households for training, 2,411 held out for testing), a feed-forward neural network with a single hidden layer of 9 nodes was estimated, directly mirroring the architecture used by Giudici and Polinesi (2019) for the "expected risk" classification discussed in Chapter 3 (their Figure 5). Continuous predictors were standardized prior to training, and the training set was rebalanced across the three risk classes (via oversampling) to address the class-imbalance problem identified in the previous section. Training converged after 80 iterations via early stopping (validation-based), reaching a best validation balanced accuracy of 0.603.

Table 3. Simple neural network, classification performance (test set, n = 2,411)

	Precision	Recall	F1-score	Support
Conservative	0.92	0.70	0.80	1,846

	Precision	Recall	F1-score	Support
Moderate	0.24	0.37	0.29	339
Growth-oriented	0.30	0.63	0.41	226
Accuracy			0.649	2,411
Balanced accuracy			0.566	

Figure 2. Simple neural network architecture (1 hidden layer, 9 nodes)



Note: connection thickness and opacity reflect the magnitude of each trained weight; slate blue indicates a positive weight, burgundy a negative weight

Discussion


Compared with the ordinal logistic regression, the simple neural network trades some overall accuracy (65.0% vs. 77.5%) for a markedly better balance across classes (56.6% vs. 40.8% balanced accuracy). Concretely, it recovers a meaningful share of the Moderate class 37% recall, against 0% for the regression while also improving recall for the Growth-oriented class (63% vs. 24%). This confirms, with a second and quite different modelling approach, that linear models anchored to the dominant class struggle with intermediate, minority categories, while more flexible models are better able to identify them, at some cost to overall accuracy.

Table 4. Permutation importance

Variable	Mean importance	Std. dev.
Risk tolerance	0.0551	0.0095
Income	0.0507	0.0075
Wealth	0.0394	0.0070
White-collar occupation	0.0271	0.0052
Retired occupation	0.0246	0.0064
Age	0.0155	0.0099
Entrepreneur/self-employed occ.	0.0105	0.0060
Education	0.0089	0.0064
Other non-employed occupation	0.0088	0.0057
Manager occupation	0.0048	0.0028
Blue-collar occupation	0.0048	0.0044

Variable	Mean importance	Std. dev.
Other self-employed occupation	0.0021	0.0022

Figure 3. Permutation importance, ranked



Because a neural network does not offer individually interpretable coefficients, permutation importance a standard model-agnostic technique that measures how much predictive performance degrades when a single variable is randomly shuffled is used here as the network's counterpart to the regression's coefficient plot. The ranking that emerges is more evenly distributed than either the regression (dominated by risk tolerance and education) or the tree-based exploration earlier in this chapter (dominated almost entirely by wealth): risk tolerance, income and wealth all contribute meaningfully and in relatively similar measure, with occupation category particularly the white-collar and retired 45roup salso playing a non-trivial role. Education, by contrast, contributes comparatively little to the network's predictive performance, in contrast to its central role in the

regression a genuine point of disagreement between models worth flagging rather than glossing over.

Concrete illustrations from real households

To connect the model's output back to the kind of concrete, human illustration this thesis has aimed for throughout, three households from the test set real, anonymized observations from the SHIW 2022 sample are presented below, alongside the probabilities the network actually assigned to each risk class.

Table 5. Real households examples

Actual class	Household profile	Model prediction
Conservative	59 y/o, primary education, no recorded income, ~€1,000 wealth, lowest risk tolerance, "other non-employed"	97.5% Conservative
Growth-oriented	62 y/o entrepreneur/self-employed professional, postgraduate degree, ~€414,000 income, ~€6.58M wealth, moderate risk tolerance	82.5% Growth
Moderate	59 y/o, professional diploma, ~€113,600 income, ~€3.67M wealth, low-moderate risk tolerance, "other non-employed"	63.6% Moderate

The conservative case closely mirrors the "Michela" example introduced earlier in this thesis: with no financial cushion and no appetite for risk, a conservative allocation is predicted with near-certainty. The growth-oriented case is intuitive given the household's substantial financial capacity, even though its self-reported risk tolerance is only moderate suggesting the model weighs financial capacity alongside stated risk preference rather than relying on either alone. The moderate case is the most informative of the three: despite substantial wealth (€3.67 million), the network correctly resists classifying this household as growth-oriented, instead weighing its comparatively cautious stated risk tolerance to arrive at an intermediate classification precisely the kind of nuanced judgement a simple rule of thumb ("wealthy households take more risk") would mis

Table 6. Confusion matrix (test set, n = 2,411)

Actual \ Predicted	Conservative	Moderate	Growth-oriented
Conservative	1,299	345	202
Moderate	87	124	128
Growth-oriented	31	53	142

Shaded cells indicate correct classifications (diagonal).

The confusion matrix confirms the pattern described above: while the regression's errors were almost entirely concentrated in the Moderate class (misclassified 100% of the time), the simple neural network distributes its errors more evenly, at the cost of misclassifying a larger share of the (numerically dominant) Conservative households as Moderate or Growth-oriented (30% of Conservative households are misclassified, against just 2% for the regression). This is the practical face of the accuracy/balance trade-off discussed above: neither model is unambiguously “better” each optimises for a different notion of what a good classifier should do.


The deep neural network: results

The same training/test split and rebalancing procedure used for the simple network were applied here, but with a deeper architecture: three hidden layers of 5, 3 and 2 nodes respectively, directly mirroring the architecture used by Giudici and Polinesi (2019) for the “actual risk” classification of financial products discussed in Chapter 3 (their Figure 6). Compared with the simple network, this architecture introduces two additional hidden layers and a narrowing structure (5 → 3 → 2 nodes) before the output layer, allowing the model, in principle, to learn more complex, hierarchical combinations of the input variables at the cost of a harder optimisation problem and a greater risk of overfitting on a dataset of this size. Training converged after 54 iterations via early stopping, reaching a best validation balanced accuracy of 0.569 lower than the 0.603 achieved by the simple network.

Table 7. Deep neural network, classification performance (test set, n = 2,411)

	Precision	Recall	F1-score	Support
Conservative	0.92	0.67	0.78	1,846
Moderate	0.22	0.31	0.26	339
Growth-oriented	0.26	0.68	0.38	226
Accuracy			0.621	2,411
Balanced accuracy			0.555	

Figure 4. Deep neural network architecture (3 hidden layers: 5-3-2 nodes)



What changes relative to the simple network

The additional depth does not translate into better performance on this dataset. Overall accuracy falls slightly, from 65.0% (simple network) to 62.1% (deep network), and balanced accuracy falls from 56.6% to 55.5%. Class-by-class, the deep network trades some Moderate-class recall (31% vs. 37%) for a modest gain in Growth-oriented recall (68% vs. 63%), but the overall picture is one of comparable if anything, slightly weaker performance despite the additional architectural complexity.


This is a substantively important finding for this thesis, not merely a technical footnote: it confirms, with real data and a real modelling exercise, a pattern well documented in the machine learning literature that additional depth does not automatically improve performance on small-to-medium, tabular datasets with a limited number of features, and can in fact make optimisation harder without a corresponding gain in predictive power. With only twelve input variables and roughly seven thousand training observations, there is arguably too little complexity in the underlying relationships, and too little data, for a deeper architecture to exploit its theoretical advantage over a shallower one.

Table 8. Permutation importance

Variable	Mean importance	Std. dev.
Retired occupation	0.0779	0.0080
Other non-employed occupation	0.0457	0.0048
Income	0.0424	0.0110
White-collar occupation	0.0388	0.0052
Blue-collar occupation	0.0377	0.0047
Wealth	0.0368	0.0113
Risk tolerance	0.0363	0.0084

Variable	Mean importance	Std. dev.
Entrepreneur/self-employed occ.	0.0203	0.0046
Education	0.0106	0.0061
Manager occupation	0.0093	0.0035
Other self-employed occupation	0.0054	0.0025
Age	0.0011	0.0049

Figure 5. Permutation importance, ranked



The permutation importance ranking changes considerably relative to the simple network itself a further illustration of how sensitive neural networks can be to architectural choices. Where the simple network relied most heavily on risk tolerance, income and wealth, the deep network relies most heavily on

occupation category, particularly “retired” and “other non-employed”, with the continuous financial variables (income, wealth, risk tolerance) all contributing a roughly similar, more moderate amount. Age, which already contributed little in the simple network, becomes almost entirely irrelevant here. This shift is a useful, concrete demonstration of a broader point made in this thesis: unlike the regression's coefficients, which are stable, additive and directly interpretable, a neural network's internal “logic” can shift substantially with architecture alone, even when trained on the very same data a further reason why interpretability, and not only predictive accuracy, matters when choosing a model for this kind of application.

The same three households, through both networks

To make this comparison concrete, the same three households presented for the simple network are shown again below, this time alongside the predictions of both networks.

Actual class	Household profile	Simple NN	Deep NN
Conservative	59 y/o, primary education, no income, ~€1,000 wealth, lowest risk tolerance, "other non-employed"	97.5% Conservative	95.6% Conservative
Growth-oriented	62 y/o entrepreneur, postgraduate degree, ~€414,000 income, ~€6.58M wealth, moderate risk tolerance	82.5% Growth	61.9% Growth
Moderate	59 y/o, professional diploma, ~€113,600 income, ~€3.67M wealth, low-moderate risk	63.6% Moderate (correct)	54.4% Growth (WRONG)

Actual class	Household profile	Simple NN	Deep NN
	tolerance, "other non-employed"		

The first two cases are classified consistently by both networks, though with somewhat lower confidence from the deep network in the growth-oriented case (61.9% vs. 82.5%). The third case is the most instructive: the moderate household that the simple network correctly identified (63.6% Moderate) is misclassified by the deep network, which instead assigns it a 54.4% probability of being Growth-oriented seemingly over-weighting the household's substantial wealth (€3.67 million) at the expense of its comparatively cautious stated risk tolerance. This single example mirrors, at the level of one household, exactly the aggregate pattern discussed above: added architectural complexity does not guarantee a better-reasoned prediction, and can in some cases produce a less nuanced one.

Table 9. Confusion matrix (test set, n = 2,411)

Actual \ Predicted	Conservative	Moderate	Growth-oriented
Conservative	1,237	338	271
Moderate	75	106	158
Growth-oriented	26	46	154

Shaded cells indicate correct classifications (diagonal).

Taken together, the results of this section support a clear conclusion for this thesis: when comparing the two networks, the simple architecture is, if anything, the stronger choice on this dataset achieving comparable or better accuracy with a simpler, more stable structure. This is a useful complement to the discussion in Chapter 3, where Giudici and Polinesi (2019) apply a deeper architecture specifically to a higher-dimensional problem (predicting ETF risk class from thirty days of return data); the present chapter's findings suggest that architectural

depth should be matched to the complexity of the underlying problem and the volume of available data, rather than assumed to be beneficial by default.

Conclusions

Summary of the Path Taken

This thesis set out to explore a simple but consequential question: is investing genuinely as inaccessible as most people believe, or has it become, in practice, a realistic option for anyone willing to build a minimal foundation of financial understanding?

The work unfolded in four stages, each building on the one before it. Chapters 1 and 2 laid the theoretical groundwork: the distinction between saving and investing, the relationship between risk and return, the role of diversification and time horizon, and closely tied to these concepts the state of financial education in Italy and the cultural and behavioural reasons why so many households continue to leave their savings idle. Chapter 3 moved from theory to practice, examining the two dominant paths through which a modern investor can actually access financial markets: delegated investing through robo-advisors and self-directed investing through low-cost digital platforms. These were compared not to declare one superior to the other, but to weigh their respective costs, benefits, and the level of financial knowledge each realistically demands. This chapter also grounded the discussion in real evidence, presenting four actual investment portfolios, each reflecting a different risk profile and a different starting point from a single-ETF plan of €100 per month to a considerably more aggressive, diversified position including individual stocks and cryptocurrency. Chapter 4 then turned to formal empirical analysis, taking Giudici and Polinesi's (2019) study of robo-advisory risk profiling as a methodological starting point and building on it. An ordinal logistic regression and two feed-forward neural networks of different architectures were estimated on real microdata from the Bank of Italy's Survey on Household Income and Wealth (SHIW 2022), in order to test with actual data rather than only theoretical argument which household characteristics are genuinely associated with a more or less risk-tolerant investment profile.

Answering the Central Question

The central thesis of this work that investing is more accessible than commonly believed, provided a minimal degree of awareness is in place is, on the basis of everything presented in the preceding chapters, confirmed without reservation. Investing is no longer a domain reserved for the wealthy or for those with access to private banking and exclusive instruments such as private equity. It is, today, something an ordinary person can begin the moment they decide to, provided they have grasped the relatively simple foundational concepts set out in this thesis.

This confirmation, however, comes with an important qualification: accessibility does not mean speed. Nobody becomes wealthy overnight through investing, and this thesis has, if anything, argued the opposite: building meaningful wealth through financial markets requires time, and time is the single most powerful lever available to an investor. This is not merely a theoretical claim it is the central, recurring finding of Chapter 1's discussion of time horizon, and it is confirmed empirically by the simulations presented earlier in this work, which show, in concrete euro terms, how the same modest monthly contribution compounds into a substantially larger outcome the earlier it begins. A young investor who starts small has a structural advantage that no later, larger contribution can fully replicate: an investor who begins with modest sums at twenty-five will, in almost every realistic scenario, end up considerably better positioned than one who begins with far larger sums at thirty-five. The single most important piece of practical advice this thesis can offer, if it can be reduced to one sentence, is simply this: start early, start small if necessary, but start.

Limitations of This Study

Several limitations of this work should be acknowledged openly, both in the interest of academic honesty and as a guide to the future research directions discussed in the next section.

First, the empirical models presented in Chapter 4, while grounded in real and representative data, achieved only moderate predictive performance. All three models the regression and both neural networks struggled to correctly identify the intermediate "Moderate" risk category, which proved structurally harder to separate from its neighbouring classes than either the Conservative or Growth-oriented groups. This reflects a genuine methodological challenge rather than an error in the analysis, but it does mean the resulting models should be read as

illustrative of underlying patterns rather than as production-ready classification tools.

Second, the SHIW dataset does not include a direct, standardised test of financial literacy comparable to the OECD/INFE instruments discussed in Chapter 2. This thesis relied on a self-reported risk-attitude question as the closest available proxy, which, while conceptually similar to the MiFID-style questions used in the case study of Chapter 3, is not a perfect substitute for a genuine knowledge assessment.

Third, the four real-world portfolios presented in Chapter 3 were included for illustrative purposes, to make abstract concepts concrete. They are not, and were never intended to be, a statistically representative sample of Italian investors.

Finally, this thesis intentionally focused on a specific and relatively narrow set of instruments principally ETFs, individual equities, and, briefly, cryptocurrency as the practical vehicles best suited to illustrating the core concepts developed throughout the work. Other asset classes relevant to a complete financial education, such as real estate, insurance-based investment products, and private pension schemes, were mentioned only in passing and would merit dedicated treatment of their own.

Directions for Future Research

Several natural extensions follow from the limitations above and from the themes developed in Chapter 3 in particular. Future work could expand the empirical analysis in Chapter 4 using multiple waves of the SHIW panel, allowing risk profiles and portfolio composition to be studied longitudinally rather than in a single cross-section, and potentially improving the identification of the harder-to-classify intermediate risk categories. A further natural extension, directly connected to the self-directed investing discussion in Chapter 3, would be a more thorough treatment of the FIRE (Financial Independence, Retire Early) movement, which builds directly on the same principles of disciplined, long-horizon investing discussed throughout this work, but pushes them toward the explicit goal of early financial independence. Finally, future research could broaden the scope of instruments considered beyond ETFs, equities and cryptocurrency, to include a more complete map of the financial and real-asset instruments available to a retail investor, from pension products to real estate investment vehicles.

A Personal Reflection

Beyond its theoretical and empirical contributions, this thesis reflects a conviction that has taken shape gradually, over the course of both the research presented here and my own direct experience with investing, discussed in Chapters 2 and 3: investing is something that everyone can do, provided two conditions are met, neither of which is out of reach for the ordinary person. The first is a minimal financial foundation not a deep or sophisticated one, but a basic, working understanding of the concepts set out at the start of this thesis. The second, perhaps less obvious but equally important, is the courage to change one's own relationship with money, and with the family's accumulated wealth, however modest it may be.

Both of these conditions can be met starting today, not in some distant future contingent on higher income or greater expertise. And the potential consequences of meeting them are significant: not only over the course of one person's lifetime, but across generations, laying a foundation that can be passed on rather than rebuilt from zero each time. This is, to me, one of the most important ideas this thesis has tried to convey, and one that deserves to be far more widely understood than it currently is.

Developing a genuine curiosity for this subject, over the course of writing this thesis, made clear just how much there is to learn for anyone who wishes to go deeper the range of instruments, strategies and ideas available to a motivated investor is genuinely vast. But what I have come to see as the more important, and more encouraging, insight is the opposite one: building meaningful wealth over the course of a lifetime does not require mastering all of that complexity. The fourth portfolio presented in Chapter 3 a single ETF, accumulated at €100 per month makes this point better than any argument could. It is, quite plausibly, the simplest, most accessible and most affordable starting point available to anyone; and yet, sustained over thirty years and eventually passed on to the next generation without ever needing to be liquidated, it is entirely capable of changing a family's financial trajectory.

What is required, above all, is a change in mindset. An investment account is not a current account, and should not be judged, or feared, by the same standards. The path is not a straight line upward. It is, in shape, closer to exponential growth over the long run, but one made up of constant short-term fluctuation along the way: days of gains of six percent, days of losses of ten percent on the back of a single piece of bad news. What matters, in the end, is not what happens along the way, but where the path ultimately leads. This single point the ability to tolerate the middle of the journey without losing sight of its destination is, in my view, one of the greatest psychological obstacles standing between ordinary

families, and ordinary young people in particular, and the decision to begin investing. But it is an obstacle that, once overcome, tends to be followed by considerable and lasting reward.

Bibliography

- Giudici, P., & Polinesi, G. (2019). Scoring models for roboadvisory platforms: a network approach. *Journal of Network Theory in Finance*

Sitografia

Bank of Italy data and documentation

- Banca d'Italia. *Indagine sui bilanci delle famiglie italiane distribuzione dei microdati.* <https://www.bancaditalia.it/statistiche/tematiche/indagini-famiglie-imprese/bilanci-famiglie/distribuzione-microdati/index.html>
- Banca d'Italia. *Indagine sul 2022 Descrizione degli archivi.* <https://www.bancaditalia.it/statistiche/tematiche/indagini-famiglie-imprese/bilanci-famiglie/documentazione/documenti/2022/Legen22.pdf>
- Banca d'Italia. *Indagine sul 2022 Questionario.* https://www.bancaditalia.it/statistiche/tematiche/indagini-famiglie-imprese/bilanci-famiglie/documentazione/documenti/2022/Questionario_2022.pdf

Financial literacy

- Banca d'Italia. *Indagini sull'alfabetizzazione finanziaria e le competenze di finanza digitale in Italia: adulti 2023.* <https://www.bancaditalia.it/pubblicazioni/indagini-alfabetizzazione/2023-indagini-alfabetizzazione/index.html>
- Banca d'Italia. *Alfabetizzazione finanziaria pagina generale con accesso ai dati.* <https://www.bancaditalia.it/statistiche/tematiche/indagini-famiglie-imprese/alfabetizzazione/index.html>
- Il Sole 24 Ore. *L'educazione finanziaria degli italiani migliora, ma troppo poco per evitare rischi.* <https://www.ilsole24ore.com/art/l-educazione-finanziaria-italiani-migliora-ma-troppo-poco-evitare-rischi-AFPbeEI>

Wealth of European families

- EFAMA. *Households continue to keep a disproportionate amount of money in bank deposits in most European countries.* <https://www.efama.org/newsroom/news/households-continue-keep-disproportionate-amount-money-bank-deposits-most-european>

Financial tools

- justETF. *ETF Savings Plan Calculator*. <https://www.justetf.com/en/etf-savings-plan/etf-savings-plan-calculator.html>
- justETF. *Vanguard FTSE All-World UCITS ETF (USD) Accumulating Scheda ETF*, ISIN IE00BK5BQT80. <https://www.justetf.com/en/etf-profile.html?isin=IE00BK5BQT80>

Investing platforms

- EQS News. *Trade Republic receives full banking license from ECB*. https://www.eqs-news.com/news/corporate/trade-republic-receives-full-banking-license-from-ecb/714e21cb-b294-48ac-b1d7-a2e3d4d51734_en