UNIVERSITY OF SOUTHAMPTON

FACULTY OF SOCIAL AND HUMAN SCIENCES

Division of Economics

Learning and Strategic Asset Allocation

by

Michael Kearns

A thesis produced for the degree of Doctor of Philosophy

September 2015

This thesis is dedicated to Jesus Christ, my saviour. He has been a very close source of help in times of need. “Christ Jesus came into the world to save sinners - of whom I am the worst. But for that very reason I was shown mercy so that in me Christ Jesus might display his immense patience for those who would believe in him and receive eternal life.”

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL AND HUMAN SCIENCES

DIVISION OF ECONOMICS

Doctor of Philosophy

Essays on Portfolio Choice and Information Acquisition

by Michael Kearns

This thesis investigates whether or not models that portray the relationship between what an investor learns and how he allocates his portfolio can explain phenomena related to household behaviour in the stock market. Endogenous modelling of household learning is utilised, which builds on a growing literature called bounded rationality with increasing explanatory power, offering an alternative to the classical rational expectations theory. Such phenomena include firstly why households often hold portfolios that are little diversified, secondly why household beliefs about the stock market exhibit widespread heterogeneity despite past data being publicly available and lastly whether or not they employ strategic motives in the stock market and whether they complement others' actions or substitute them for their own.

In particular, Paper 1 addresses the observation that a significant number of investors hold concentrated portfolios, apparently forgoing the benefits of diversification. In a static portfolio choice model with limited capacity constraints, Van Nieuwerburgh & Veldkamp (2010) show that the observed lack of diversification is rational amongst investors with a strong preference for an early resolution of uncertainty. This paper studies whether or not endogenous information acquisition can also rationalise the observed peculiarity within a dynamic portfolio choice model. It is found that in the steady state, a hedging demand component appears in the optimal portfolio, which attempts to spread risk across the investor's investment horizon: The investor chooses additional information precision the smaller is the compensation for bearing risk. A numerical approximation to the agent's decision rule suggests that the research question can be answered in the affirmative. This is for an investor who is indifferent to the time of

uncertainty resolution and is risk averse over wealth.

Paper 2 tackles the question “Is demand for information positively correlated with returns?” Results from this study indicate “no” in any permanent sense. Using a dynamic model of endogenous information acquisition and portfolio choice, simulations reveal that once an agent learns the underlying process behind returns, his demand for information is constant and hence acyclical. This is surprising given survey results found by Coibion et al (2015), which, along with the PAT€R survey 2014 wave, document widespread heterogeneity in agent beliefs. Hence it appears that economic agents in reality do not treat returns as though they are driven by an underlying, learnable process. This paper contributes to literature on bounded rationality and limited processing of information. It also finds a rationale for belief heterogeneity: ignoring realised returns and observing signals that support prior beliefs allows degenerate distributions to emerge. This result can be produced under the assumption that the investor does not access public information.

Lastly, the third paper addresses whether or not households behave strategically regarding each other when making stock market participation decisions. The study examines the contribution of strategic considerations in stock market return expectations on the demand for risky assets empirically, exploiting novel data from the 2014 PAT€R survey wave, representative of the population by age and wealth. The strategy is to identify whether individual stockholding decisions are consistent with strategic substitutes or complements prevailing in the stock market, under the null hypothesis of efficiency. The study finds evidence for strategic complementarity and additional information variables can explain this effect. Given the substantial heterogeneity in expectations and perceptions of returns, and the relatively low degree of sophistication of the median investor identified in the empirical literature, the project concludes that as strategic substitutes prevail even amongst them, a portion of the excess volatility observed in stock markets may be driven by expectational motives in coordination.

Contents

List of Tables

List of Figures

Declaration of Authorship

I, Michael Kearns, declare that the thesis entitled *Essays on Portfolio Choice and Information Acquisition* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself (chapter 3);
7. Either none of this work has been published before submission.

   –

Signed: ....

Date: ....

Acknowledgements

"My heart took delight in all my labour, and this was the reward for all my toil." Ecclesiastes 2:10, The Holy Bible.

Firstly, I would like to acknowledge and thank my supervisory team, which consists of Dr Hector Calvo-Pardo and Dr Thomas Gall. Hector has been a reliable source of advice and has provided me with many helpful references. Thomas' office door has always been open to me and has been a wonderful source of support. I have also been blessed by the comments and feedback from various members and former members of the Southampton University Economics department. In particular, I would like to mention Dr Roman Sustek, Professor John Knowles and Dr Arnau Valladares.

Crucially, I want to acknowledge and sincerely thank the Economic and Social Research Council, which has provided me with a studentship. Without it I would not have been able to finance this venture and it has allowed me to access many superb resources including training courses, conferences and a summer school that would otherwise have been beyond my reach.

Truly, I want to mention my fellow PhD students, together we have had a tremendous source of community and affection whilst spending long hours studying together. I am so thankful for the friendship and wisdom given to me from, to mention some, Carlos Pineda, Marcos Gomez Mella, Hao Xu, Tong Xue, Panos Giannarakis, Andres Luengo, Jana Sadeh (nee Farrugia), Nicholas Lazarou and Veridiana Nogueira.

Both Bitterne Park CC and Southampton Lighthouse International Church have given me excellent opportunities to unwind and relax outside of my studies. Your love and friendship have really boosted me.

Lastly, and very importantly, I want to acknowledge my family. I am sorry that the amount of time we have spent together has suffered, you have all been patient with me as I have devoted long periods of time to these studies. Thank you for only being a message or phone call away.

1

Introduction

The aim of this introduction to place the thesis within a general context of relevant literature and to give a brief overview of the three papers, of which it is comprised, that aim to contribute to that literature. Although each constituent paper is unique and aims to answer a separate research question, all three overall fit into the topic of Macroeconomic Learning and in particular, agent-based learning, as well as asset allocation, which is part of Financial Economics. The ensuing literature covers broad research areas that could be considered to be "neighbouring" as those in which this thesis can directly be placed.

Comprising this thesis are three studies in the fields of Macroeconomic learning and asset allocation, the first of which tackles the puzzle of why investors decide to hold portfolios concentrated in few risky assets when the theoretical benefits of diversification are well understood by economists. The second project asks when, given capacity constraints, information should be acquired and documents the fact that widespread heterogeneous beliefs about the stock market exist among households despite returns being publicly accessible and proposes a theory that can replicate the stylised fact based on the assumption that an investor does not observe public information. Lastly, cross-section data on the perceptions and expectations of households about the stock market and related decisions are empirically analysed for evidence of strategic motives between them in information processing and actions.

1.1 Portfolio selection

The Markowitz Paradigm is the well-known textbook foundation of portfolio theory and derives from Markowitz (1952). It states two key principles of investing in risky assets. Firstly, if risky assets are imperfectly correlated, then an investor can achieve superior risk and return characteristics by diversifying his portfolio choice. Secondly, after diversifying as such, he can only achieve a higher return by taking on more risk. What this means is that standard theory expects investors to hold diversified portfolios.

Markowitz (1952) is also an early paper that alludes to the combined importance of learning and portfolio selection, firstly beliefs must be set through observation and secondly wealth must be allocated based on these beliefs, though it exclusively treats the second. The main theoretical enterprise of this thesis is to express that formulation in a dynamic model with discrete time. Markowitz strongly rejects the maxim of investing to solely maximise expected future payoffs from assets as it indicates that the investor would invest only in the asset with the highest expected payoffs and it would render equally valuable assets equally desirable. Maximising expected future payoffs only gives no regard to time preferences or risk. So his work was influential in establishing the use of mean-variance preferences over discounted future payoffs. In the paper the exclusion of short-selling is critical and can rule out efficient portfolios by making them unattainable, this can even include the minimum variance portfolio. The development of another key theory in financial economics, the Capital Asset Pricing Model, followed the work of Markowitz and applies mean and variance preferences to treat the issue of a diversified portfolio. Seminal work on it took place in the 1960s and Black, Jensen & Scholes (1972) construct a version with the no risk free asset assumption that held up better to empirical studies and saw the theory become more popularly used. Empirical studies had found that regressions of the expected excess return on estimated systemic risk of an asset saw a positive constant (indicating a positive excess return for no anticipated risk) and a coefficient on risk different to the expected market return (as the theory stated). Black et al introduce a beta factor that is the return on a portfolio with zero covariance with the market portfolio. Evidence they analyse shows that the beta factor grew by 1-1.3% over 35 years, which was similar to the average market return (1.2-1.3%) rather than being zero as the previously held wisdom suggested. Further evidence is the high t-statistics found for the beta factor. In earnest, CAPM is for treating general equilibrium pricing of assets.

Present versions of Markowitz's theory allow the investor to borrow at the risk free rate to leverage his chosen portfolio. Introducing short-selling removes the upper bound on the efficient frontier and Merton (1972) shows that the efficient set is the upper part of a hyperbola. Markowitz's theory applies to static portfolio choice and a graphical representation of his efficient frontier is below, taken from Brandt (2010).

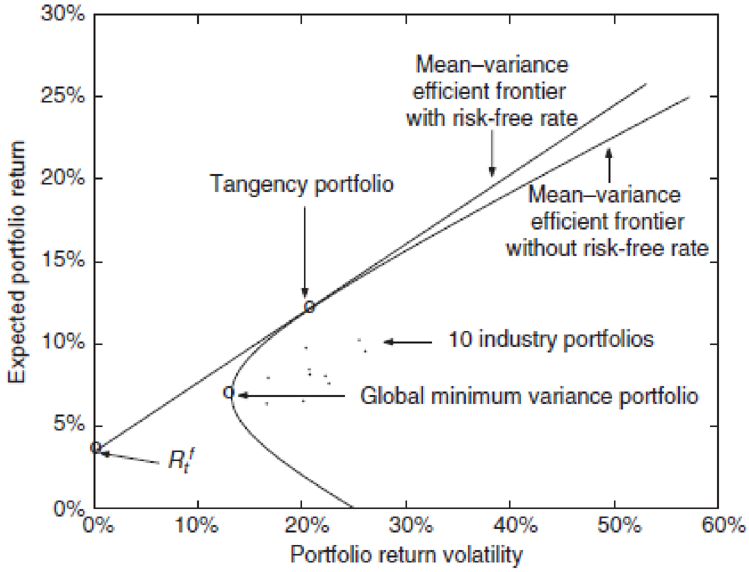

Each individual dot is a portfolio and these can be combined. Optimal combinations form the upper part of the frontier where the variance is minimised for a given expected return. The combination of portfolios that varies the least is the left-most point of the frontier, the global minimum variance portfolio. Achieving a greater expected return than this portfolio requires the investor to accept higher risk. With the possibility of lending and borrowing at the risk free rate (buying and selling government bonds), the efficient frontier expands, with R_f the risk free rate. A single portfolio becomes preferred and this is the tangency portfolio between the line extending from R_f and the hyperbolic frontier. Any leftward or rightward portfolio gives an inferior expected return for the same variance on the new linear frontier. The frontier itself is a linear combination of the risk free asset and the tangency portfolio. Before the tangency portfolio the agent lends (buys government bonds) and beyond it he borrows (sells) at the risk free rate. Minimising the variance of a portfolio constructed from N portfolios by selecting portfolio weights subject to a target return and a limit on the weights traces out the Markowitz frontier.

$$min_x x'\Sigma x$$

subject to

$$\bar{R} = xR, \qquad \sum_{i}^{N} x_i = 1.$$

x is the vector of portfolio weights, Σ is the variance-covariance matrix and R is the vector of expected returns. The given return is given by $\bar{R}$ and varying it traces out the frontier and restricts how much the variance is minimised. This mathematical representation is based on Brandt (2010). For portfolios of more than three assets a graphical representation was typically used but Merton (1972) derives optimal and unique portfolio weights required to minimise portfolio variance for a target expected return.

Survey data of American households indicates that, contrary to the Markowitz Paradigm, diversification is oft-eschewed. This is hard to square with the portfolio choice literature. The observation is documented in Cucuru et al (2010) and that text identifies two subgroups within the one that diversifies less than expected. The first group is one that consists of individuals that decide to invest in the companies that employ them. Normally, such selections are part of retirement plans that allocate stock holdings to workers. This is a particularly crazy-looking strategy because it means that there is a strongly positive correlation between the labour income and portfolio income of such individuals. Hence the risk faced by the household is increased and not decreased in the portfolio choice. For example, if the employer were to perform extremely badly, then its stock market performance would deteriorate and so the household's portfolio income would receive a negative shock. Moreover, the company may force the employee to take a pay cut or even fire him. Hence household labour income would also receive a negative shock. So the asset allocation would enhance negative outcomes rather than soften them. Viceira (2001) investigates this relationship between stock market and labour market behaviour. He finds that the solution to his theoretical models shows that agents prefer to hold safer assets in the stock market when risk to their labour income increases. So the observation that many households invest in the stocks of their employers is very peculiar. Guiso et al (1996) find empirical evidence to support the theory that labour (amongst others) income risks reduce household willingness to take risk in the stock market. However, they treat labour income risk as being independent of risky asset returns

The second subgroup is one that appears to deliberately hold its wealth in a small number of stocks. The behaviour of this group also contravenes the Markowitz Paradigm because it seems like these individuals consciously forgo the benefits of diversification. It is into this context that the first paper in this thesis fits. A paper that attempts to tackle this puzzle is Van Nieuwerbugh & Veldkamp (2010). These authors build a static model of portfolio choice that is preceded by an endogenous information acquisition decision. The model is static because the investor makes each decision only once. The endogenous information acquisition is modelled as an allocation of costly informative signals between available risky assets. In particular, the investor chooses the number of signals to receive about each asset subject to a limit to the total number.

Many portfolio models were static but some early dynamic models are studied in Merton (1969), (1971) & (1973) using continuous time and Samuelson (1969) and Phelps (1962) do so in discrete time. These were the first stochastic models of saving behaviour. Merton applies mean-variance analysis in each. This thesis is purely concerned with portfolio choice problems though these papers find the optimal consumption rule when risky storable assets are available. Merton (1969) finds the optimal decision rules under CRRA utility to determine constant portfolio weights over time (independent of wealth and time) and consumption is a constant fraction of stochastic wealth. Also, the consumption decision is independent of the parameters of optimal portfolio weights through the CRRA utility (constant proportion of wealth invested) and stochastic price generator used. The less risk averse typically save more and consume less when higher returns are available as the substitution effect dominates but the more risk averse will do the opposite as the income effect is stronger for them. Merton (1971) is able to prove the Mutual Fund Theorem in continuous time and eliminate reliance on normally distributed returns and quadratic utility. The former allow unlimited liability and the latter produces increasing absolute risk aversion. The Mutual Fund Theorem states that any optimal portfolio for an investor can be constructed from a linear combination of two unique mutual funds. It is useful because it suggests that an investor can avoid costly financial transactions by diversifying across just two funds.

Two advantages of using continuous time over discrete time models are that only two types of stochastic process are required: Poisson and functions of Brownian motion. These are both well understood so there is much literature to draw on and the second advantage is that the number of parameters required in the model is reduced.

Interest in dynamic portfolio choice models started to an extent because the popular CAPM model is strictly a one period model and was used akin to a dynamic model. With constant preferences and investment opportunities, there was some justification for this but the second assumption in particular is quite restrictive. Merton (1973) argues for the use of dynamic models that retain the tractability and simplicity of the CAPM but are able to model the effects of stochastic investment opportunities. Shown by Merton (1971) to be different to myopic choices, the optimal portfolio of an investor in a dynamic model is determined by a changing investment opportunity set. Facing a multi-period problem and stochastic investment opportunities the investor adds a term to the optimal myopic portfolio, which is called a hedging demand. This is an attempt to spread risk over the time dimension. Brandt (2010) is one study that took the discovery of hedging demands and outlines three conditions under which myopic investing remains optimal. First is that utility is logarithmic meaning that periods are additively separable, second is that investment opportunities are independent so that state variables today indicate nothing about returns tomorrow and third is that there is no way to hedge stochastic investment opportunities. This thesis uses Brandt's condition of independent returns (IID returns are used) to show that hedging demands can arise in such situations if a learning technology is available to the investor. Merton (1973) illustrates hedging demands in a continuous time model that reflect the investor's attempt to hedge against unfavourable future shifts in the investment opportunity set. An unfavourable shift would be an increase in the variables x that affects asset returns when the marginal effect of x on consumption or wealth is negative. Demand for the asset through hedging demands rises when the correlation between x and returns becomes more positive. When investment opportunities are static then the demand for assets in continuous time also reduces to that of the myopic investor as shown by Merton (1971).

The discovery that empirical evidence contradicts two of the sufficient conditions for myopic portfolio choice has further driven research into long-horizon portfolio choice models. Firstly, expected asset returns vary through time, ruling out the assumption that investment opportunity sets are constant. Secondly, research on the equity premium puzzle has found excess returns on stocks to be too high to be reasonably explained by logarithmic utility. Texts to refer to on the equity premium puzzle are Campbell (1996), Cochrane & Hansen (1992), Hansen & Jagannathan (1991), Kocherlakota (1996) and Mehra & Prescott (1985).

Some studies have investigated hedging demands further. Campbell & Viceira (1999) compare a myopic investment choice to an optimal portfolio when returns can be predicted by a particular factor and use this technique to quantify hedging demands. They find that hedging demands are negative for risk lovers. As the risk aversion coefficient rises, so do hedging demands but as the agent becomes very risk averse, he minimises his exposure to risk through both the myopic and hedging demands components of asset demand. The risk lover wants assets that payoff when assets are productive, when current returns are high and increasing payoffs through negative hedging demands but the risk averse agent has positive hedging demands because he wants assets that payoff when investment opportunities are poor. Moreover, hedging demands have a massive impact, they account for at least 20% of stock demand but they often double the risky portfolio allocation.

Returns are predictable in Michaelides & Zhang (2015) through a single factor. Hedging demands are higher when the correlation between innovations to that factor and the stock market are negative compared to when they are zero. However, changes to the correlation between the predictive factor and permanent income shocks has little impact whilst raising the correlation between permanent income and stock market shocks reduces hedging demands. Bertaut & Haliassos (1997) and Davis & Willen (2000) study finite-horizon models whilst Dammon, Spatt & Zhang (2001) include taxes. Cocco, Gomes & Maenhout also include labour income that is not fully insurable and find that labour income is used as a substitute for a risk free asset under a positive correlation. The life cycle effect induces the household to move more into bonds, consistent with standard portfolio advice, as labour income becomes less important. Also using incomplete markets than cannot fully insure labour income, Gomes & Michaelides (2008) can produce a realistic risk premium through imposing a fixed entry cost for stock market participation on households. However, there proves to be a trade off between matching market-level facts (risk premium) and individual portfolio holding patterns.

1.2 Information choice

Endogenous information acquisition is a form of agent-based learning, along with others such as rational inattention, sticky information models and models of strategic information choice. Models of these types aim to offer an alternative to the traditional theory of rational expectations that posits that agents know the true parameters behind the data-generating processes of the economy. Seminal works on rational inattention are Sims (1998) and Sims (2003) and the literature appears to be a growing one. It models the information choice, of agents, as being subject to a limited capacity to process information, which agents rationally elect to spread between different objects about which to learn. An agent can chooses which information to carefully attend to, which to partially ignore and which to fully ignore. Sims proposed modelling attention as an information flow using the measure of uncertainty. Allocating attention increases the precision of the signal received about an unknown variable that is necessary for making a particular decision. This could be optimally adjusting a price as in Wiederholt (2010), Mackowiak & Wiederholt (2009), or Matejka (2010) or optimally selecting a consumption-savings combination as in Sims (2006).

Matejka finds that the observed behaviour by which prices remain stable and then suddenly jump can be explained as well by rationally inattentive price setters as the standard costly price adjustment theory. Sims (2006) argues that unlike traditional optimisation, rational inattention can produce discrete and externally unpredictable responses that are able to fit the slow responses of aggregate variables to shocks over time. Rational inattention applied to portfolio theory studies, such as Van Nieuwerburgh & Veldkamp (2009) and Mondria (2010), have found that learning exacerbates information advantages, agents prefer to learn more about what others do not know rather than about what they do know. This can help explain the Home Bias in Assets puzzle. Also, when investors can observe a private signal that is a combination of asset payoffs, comovement between asset prices can increase.

An alternative is to use sticky information and a paper that embodies it is Mankiw & Reis (2007). They model a constant fraction of workers, consumers and firms as receiving new information each period. So information is sticky in that it is only transmitted to a limit portion of the population over time. Other research has investigated the relationship between coordination and information choice, including Hellwig & Veldkamp (2009) and Morris & Shin (2002). These papers emphasise that an agent may wish to alter his demand for information depending on what others do. That is agents do not make information choices in isolation.

Van Nieuwerbugh & Veldkamp, in their study, are able to replicate the concentrated portfolio choices of those investors who appear to deliberately forgo the opportunities of diversification. They are able to do this when the agent they model has a impatient preference over resolving uncertainty. That is, when he wishes to learn early about the unknown payoffs tomorrow of the risky assets. In that case, he allocates his available costly signals to only one risky asset, meaning that he specialises in it. His uncertainty about the return falls and the related risk decreases. Therefore, he wishes to hold more of it in his portfolio. This in turn means that every signal he allocates to it provides more information about the portfolio he expects to hold. Thus there is a feedback effect that supports the holding of a concentrated portfolio through increasing returns to information.

However, the model of Van Nieuwerbugh & Veldkamp is a static model so my first paper contributes to the understanding of household portfolio choice by reconstructing their model to a dynamic one. This means their theory can be tested when investors face an investment horizon of more than a one period. It also means that my model can test for the presence of hedging demands. These are present when an investor chooses to spread the risk he faces over time. He would over- or under-investing compared to the portfolio he would choose if he only invested for one period. Merton (1973) was the paper that introduced this concept and a number of papers have worked further on the idea since. A good reference is Brandt (2010) but Viceira (2001) and Davis & Willen (2000) are two others papers that investigate how agents' hedging demands are affected by the presence of income, particularly labour income, shocks. The reason that the development of hedging demands is important stems from the fact that they cause the portfolios of the dynamic and myopic investors to differ. That is, the dynamic investor either over- or under-invests compared to the one who participates in the market for only one period (or considers only one period at a time). Thus, those with hedging demands hold what appear to be extra risky portfolios from the perspective of myopic investors.

1.3 Thesis overview

The first paper discovers that it is indeed possible to rationalise a concentrated portfolio in a dynamic model of portfolio and learning allocation. However, a further discovery that is very useful is that the result does not depend on investor preferences. Instead, it can be obtained with a simple ratio between the expected return and risk, from the agent's perspective, of the assets in question. As expected returns and risks of assets are pieces of data that can be collected and preferences cannot, it is now possible to test the theory empirically. What is more, the research finds that not only do hedging demands appear in the steady state but that there is also a "learning factor" present. It indicates that agents who endogenously acquire information and invest dynamically take more risk than myopic ones (hedging demands) and than those who invest dynamically but do not acquire signals (learning factor).

Whereas the first paper investigates the link between learning and portfolio choices, the second follows it by asking when an investor should acquire information given that he faces capacity constraints to information. In particular it examines the relationship between risky asset returns and the demand for information. So the angle of the paper is whether an agent wants to allocate more, less or the same amount of learning to an asset when its return rises from one period to the next. Given the relevance of information choices for investment behaviour and given the interest in Economics literature in agent-based learning, the effect of risky asset returns on the demand for information appears to be an interesting topic. Furthermore, it appears to be an issue that is little-addressed so far.

The research finds that the fluctuations of assets' returns actually have little effect on the demand for information about those assets. It is the nature of the representative agent's learning that causes this, he has an adaptive style of learning that places an increasing weight on prior beliefs and a decreasing weight on new information. Hence, as the agent learns the true distribution of returns, he is less prone to "trend following" behaviour. Trend following occurs when agents base their decisions on information from the recent past under the assumption that future outcomes will be in the same vein. Additional results indicate that, unlike in Van Nieuwerburgh & Veldkamp (2010), learning strategies can be time-dependent and that investors need deep pockets. The latter means that being uninformed can lead to an investor defaulting due to making portfolio choices that lead to large losses.

Finally, the third paper works within a strand of literature mentioned earlier; the strategic element of information choice. Hellwig & Veldkamp (2009) are able to show that agents who wish to coordinate their actions choose to learn similar information to each other whereas agents who wish to differentiate themselves from others choose to learn information that nobody else has. This paper uses Hellwig & Veldkamp's separation theorem to empirically test the strategic element of information and portfolio choice. Data from the 2014 wave of the PAT€R survey of French households is used to empirically assess whether or not they behave strategically given knowledge of others' information and actions in the stock market. Unique information on household perceptions of current and future stock market returns and beliefs about the percentage of the population that is informed and participates in the stock market are all available. Households are found to behave strategically not only in choosing whether or not to participate in the stock market but also in forming their own information sets and future expectations.

Following this introduction are three sections that contain the contributions to economic knowledge of the papers in this thesis. Various results are derived and presented in graphs and tables, which can all be found in the appendixes at the end. Also, this is where the references used can be found.

2

Paper 1

2.1 Introduction

A cursory look at standard portfolio theory shows that it is generally expected for investors to hold diversified portfolios. By trading imperfectly correlated risky assets, an investor can, under complete markets, eliminate all idiosyncratic (market-specific) risk (Markowitz 1952). However, it is observed in data that a significant number of investors, 13.7% according to the Consumer Finance Survey in the USA (Cucuru et al, 2010), hold concentrated portfolios, that is ones which are diversified very little, a surprising phenomenon indeed! Van Nieuwerburgh & Veldkamp (2010) (VN&V) explain why investors, with a preference for resolving uncertainty early, might hold concentrated portfolios by studying a static model of endogenous information acquisition[1] and

[1]Endogenous information acquisition is the technology by which an agent purchases costly signals that are discrete and noisy. This contrasts with Rational Inattention in which signals are costless and their

portfolio choice via a feedback effect between the two choices. Instead, this paper builds a dynamic model of portfolio and signal allocation with a single agent that spans an infinite horizon. Its research question is "Why do investors hold concentrated portfolios?". The first aim of this paper is to test if a specialised learning strategy (all signal allocation devoted to one asset) can explain portfolio concentration or not. The second aim is to test whether or not the dynamic results found differ from VN&V's static ones (used as a benchmark) and, if so, under what conditions the static ones are recovered.

There are two key motivations for choosing a multi-period approach. Firstly, mathematically, the solution to an infinite-horizon model is a fixed point argument. So it could be that increasing the number of periods means that the optimal portfolio and learning allocations may be different functions of the state variables. Secondly, a dynamic framework allows a dynamic investment strategy to arise and differ from a static one through the existence of hedging demands[2]. Additionally, the multi-period approach will be useful for future research as the model can be calibrated and simulated.

This paper contributes to the bounded rationality[3] literature in Macroeconomics and portfolio theory. Though portfolio concentration may appear inconsistent with Rational Expectations, an agent choosing a portfolio conditional on his own information set may act rationally despite not knowing the entire data generating process of the market. These behaviours are called 'internally rational' (maximising discounted expected utility given dynamically consistent subjective beliefs) and 'externally rational' (knowing fully and incorporating the true stochastic process in portfolio decisions) by Adam & Marcet (2011).

This study's main contribution is to develop, find steady states for and approximately solve, via a numerical method, an infinite horizon model of endogenous information acquisition and portfolio choice. The latter has not been done before, in discrete time,

noisiness depend on the attention allocate to them.

[2]These are attempts by the investor to spread risk over the dimension of time, distinct from diversifying across assets.

[3]" 'Bounded rationality' is used to designate rational choice that takes into account the cognitive limitations of the decision-maker - limitations of both knowledge and computational capacity" (Simon 1990).

according to the author's knowledge. In doing so the model further contributes by allowing for the possibility that the agent employs a dynamic investment strategy[4], that is that he desires to spread risk over time (known in relevant literature, such as Merton (1973) and Brandt (2010), as hedging demands and these are further explained in section 2.1.2). Section 2.3 shows that hedging demands remain in the steady state with exogenous learning whilst section 2.4 shows that a multiplicative factor also exists in the transition dynamics with endogenous learning. Section 2.5 uses indifferent time preferences with endogenous learning to test a static result of VN&V. A final contribution is to provide both a model and theoretical predictions about the dynamic interactions between endogenous information acquisition and portfolio choice that can be tested using both simulations and data.

This paper discovers that hedging demands appear as a result of using a dynamic model. Section 2.2 shows that the steady state portfolio choice of an agent with a preference for a late resolution of uncertainty is similar to but, significantly, different from both the solution found by VN&V and the diversified portfolio of standard portfolio theory. The first reason is that the agent has hedging demands. The second reason is something this paper identifies as a learning factor, extra risk the investor takes due to increased certainty from learning. So the horizon length really does matter for the investor because he invests more aggressively than both the myopic one of standard portfolio theory and the myopic one who acquires costly signals as in VN&V. Alternatively, the agent holds a riskier portfolio than both the Markowitz and the Merton (1973) agents.[5] In fact, even if a risky asset has a negative expected excess return (and hence should be sold short) hedging demands cause him to buy more units of it because he can spread the risk of unexpectedly bad returns across today and tomorrow. Also, the investor who acquires signals endogenously will buy even more units of that asset despite having a finer information set, because he knows his uncertainty about returns

[4]A myopic investment strategy occurs when an investor maximises all investment periods separately, whilst a dynamic investment strategy opposes this; he accounts for future periods when making his current period portfolio optimisation decision.

[5]The portfolio is riskier from the perspectives of the other agents. It is actually not so for the agent who acquires signals because his beliefs have increased precision compared to theirs.

will be lower both today and tomorrow.

This study can explain portfolio concentration without relying on a strict early resolution of uncertainty preference unlike VN&V. Section 2.5 describes this finding, whose purpose is to directly place VN&V's generalising investor in a dynamic model to test the robustness of his myopic behaviour. Intuitively, the agent learns about the asset he is strongly more optimistic and certain about as he has consistent beliefs (he expects information acquired to strengthen his beliefs). Learning increases his certainty and makes that asset even more attractive so he concentrates his wealth in it. Thus dynamic learning, allocating signals over multiple periods, perpetuates prior beliefs, given signal realisations that the investor expects. Precisely, the investor holds a concentrated portfolio when his expected Sharpe ratio (the excess return per unit of variance the agent expects) is significantly higher about one risky asset compared to the other. This is important because, whilst preferences are not measurable, expected returns and variances (confidence and uncertainty) of risky assets are and such data exists. Hence, an empirically testable theory of portfolio concentration is found in this study and future research could quantitatively assess the theory of how learning affects investment behaviour.

A further key result, is that the agent's decision rule indicates that he switches between a specialised and generalised (allocating signals to equalise posterior uncertainty across assets) learning strategy, according to the numerical solutions (see section 2.5). Intuitively, given consistent beliefs, the agent expects to keep the same strategy always but actually switches every time he learns something that he finds surprising. VN&V find that under the same preferences the static agent solely chooses a generalised strategy. This result is very significant because it implies the investment horizon affects the learning strategy that an agent chooses. Section 2.5 also finds that the agent's learning strategy depends on the expected excess return. These are new findings compared to VN&V's static benchmark.

As dynamic learning strategies are time-dependent, the model predicts that an agent with accurate initial beliefs will have a learning strategy that varies little over time (new observations rarely surprise him) whereas one who does not will have a strategy that varies significantly (new observations surprise him often). A further prediction is that the investor who employs a dynamic strategy will be less likely to default than a myopic investor because he will respond more to the signals he receives. The former will correct falsely optimistic or pessimistic beliefs quicker by changing his learning strategy. Thus, this paper finds dynamic results that differ from the static ones of VN&V and it is further able to produce multiple theoretical predictions that are testable with both simulation techniques and data.

The remainder of this paper is organised as follows. Section 2.2 elucidates the full model. Section 2.3 describes an exogenous learning benchmark version of the infinite horizon model solved for the steady state portfolio allocation, in which the agent has CARA preferences. Section 2.4 describes the version with endogenous learning wherein the agent prefers a late resolution of uncertainty, the decision rules for the optimal portfolio and learning allocations are also found, which reveal the presence of a learning factor in the transition dynamics to the steady state. Section 2.5 describes the approximate numerical decision rules for the version where the agent is indifferent to the time of uncertainty resolution, section 2.6 summarises the paper and an appendix (section A.1) that derives the results found appears after the references.

2.1.1 Portfolio theory and concentration

This subsection gives a brief outline of the standard portfolio that the observation of observed underdiversification appears to violate. It is relevant to the reader wishing to know more about diversification and portfolio choice. It also describes the model of VN&V.

Standard portfolio theory generally expects investors to spread their wealth over many assets because idiosyncratic risk can be reduced to zero, if the idiosyncratic risk between assets differs, through diversification and under the assumption of complete markets. Economy-wide risk is not affected. So, by combining (by buying and selling) assets that are not perfectly correlated, an investor can achieve a lower variance for the same expected return or a higher expected return for the same variance. An efficient portfolio frontier can be drawn by finding the lowest possible variance for each possible expected return (or the highest possible expected return for each possible variance). He can utilise this possible trade off between imperfectly correlated assets, by choosing between risk and return combinations, to find the portfolio with his preferred risk and return characteristics among all efficient portfolios. This is the first insight of the Markowitz paradigm[6]. The second insight is that to achieve a higher expected return, among efficient portfolios, the investor must bear more risk. That is, there is a trade off.

Thus the observation that some investors have "holdings concentrated in a few individual stocks", Curcuru et al (2010), who provide data on this puzzling fact, is very surprising given standard portfolio theory. They note that although underdiversification has fallen since the 1990s, it was still significant and observed in 13.7 % of all households holding equity, at the beginning of the 2000s[7]. It appears that they wastefully discard the opportunity to achieve an efficient risk-return portfolio. Theory suggests that a better chosen portfolio would mean that they could relieve themselves of some risk borne or achieve a higher expected return.

VN&V develop a static model in which a strong preference for an early resolution of uncertainty can rationalise the underdiversification peculiarity. They incorporate endogenous information acquisition into their model which gives the investor the ability to reduce the conditional variances of asset payoffs by purchasing costly signals subject to a constraint on the total change in precision. An investor may like to acquire information firstly because he may prefer an early resolution of uncertainty to a late one. If so he will purchase costly signals that reduce the uncertainty of his beliefs, which is something that he dislikes. The second reason is that he may gain useful information for rebalancing his portfolio.

A key driver in their model is the feedback effect between the choice of which assets

[6]Brandt (2010) is one possible reference amongst many for further detail

[7]The data used is from the Survey of Consumer Finances, which is based in the USA. They define households with under diversified positions as those with more than 50% of their equity holdings in less than 10 different stocks.

to learn about and which assets to invest in. The mechanism works as follows, if the investor learns about an asset, the expected amount held increases. Precision of the asset increases through learning and the investor prefers assets that he knows more about. Then, it is more valuable to learn about assets which he expects to hold more of because a single signal purchased about that asset will be informative about more of the portfolio he expects to hold. So specialised learning can motivate an investor to hold a concentrated portfolio through this feedback mechanism described. This mechanism arises when the maximisation problem is convex in the information choice. Since the second derivative of the Lagrangian is positive, subsequent signals have an increasing effect on utility and the investor will specialise his learning. What is more, an agent operating in a model wherein the choice over information precedes that of the portfolio allocation, does not find the diversified portfolio of standard theory optimal. In such a model, the nature of the portfolio optimal to the investor depends on both his preferences and the learning mechanism.

2.1.2 Hedging Demands

This subsection describes hedging demands in some detail. It is relevant to the reader who is unfamiliar with them and wishes to have some background knowledge.

Circumstances in which a myopic portfolio choice is best are well-covered in portfolio choice literature. Brandt (2010) lists three conditions under which investing myopically is preferable. Firstly, it is true when utility is logarithmic as then terminal wealth can be separated into a sum of per period utility functions to be maximised. The second condition is when investment opportunities are constant (e.g. identically and independently distributed, or IID, returns) because they are unpredictable so state variables become uninformative about future returns.[8] No inference about future returns can be made from the state variables as there is no correlation between them. Lastly, investing myopically is optimal when investment opportunities are stochastic but unhedgable. In

[8]This is unlike a VAR process in which tomorrow's return is determined by lags that are observed today and in previous days. Hence, in such a case, state variables would be informative.

the first two cases any link between the returns and the state variables is severed and in the last, periods may be linked but the investor is unable to act upon any information he has about future trends in investment opportunities. Such a case could be when markets are incomplete.

Consider then that the investor knows something (informative) about tomorrow, then he could incorporate that knowledge into his investment decision today, as today and tomorrow are linked by information. Endogenous information acquisition provides a mechanism through which state variables (beliefs about the mean and variance of returns) can become informative about future returns - signals indicate something useful about the IID realisation tomorrow. Exogenous learning is also sufficient because observed returns contain information about returns tomorrow given that the underlying distribution is stationary. Learning (of both types) effectively makes the conditional distribution of returns tomorrow on the state variables dependent on returns today and so breaks the independence of returns over time. Signals are informative in this paper because they have the same underlying mean as the returns process and zero-expectation noise.[9] If investment opportunities are also hedgable then the investor can utilise his signals to hedge risk in returns over time. Attempts to smooth investment opportunities over time are called hedging demands by Brandt (2010) and refer to the amount by which a dynamic investor "over invests" in risky assets from the myopic agent's perspective. This term derives from Merton (1973) in which the author finds optimal portfolio demands in an intertemporal framework consists of two terms. The first is the demand for risky assets that solves a single period problem whereas the second hedges against "unfavourable" shifts in the investment opportunity set. For instance, a lower expected value of future investment opportunities may persuade him to invest more today, than he would myopically, as a higher return today will somewhat offset the lower future return and smooth returns over time.

Consequently, using an infinite horizon lets this paper test for the presence of hedging demands and finds that they arise. To separate the effects on portfolio allocation of

[9] It would also be possible for returns to follow a persistent process in which the underlying parameters change over time. Signals would then be the realised return next period with some noise.

endogenous information acquisition and state variables simply becoming informative, a benchmark model is created in which the investor learns exogenously. This switches off the information acquisition facility so the agent updates his beliefs using only observed asset returns after they occur.

2.2 The model

The section presents the dynamic model of simultaneous endogenous information acquisition and portfolio allocation by a single agent. Portfolio choice depends on the information set, updated through endogenously acquired and exogenous (observed returns) information. Endogenous information acquisition can be switched off meaning only realised returns are observed, the Exogenous learning version.

This model contains two sub periods within each standard time period. A sub period contains exactly one information update, either exogenous or endogenous. Exactly the same three actions are taken each period and these are spread across the sub periods.

2.2.1 Investor preferences

The investor maximises utility over wealth expected tomorrow and there is no intermediate consumption. This is more like modelling a pension fund manager's utility rather than the client's. The first expectation is taken over prior (outer) beliefs, denoted by a negative sign and the second expectation is taken over posterior (inner) beliefs, denoted by a positive sign.

$$U_t = E_t^-[u_t^-(E_t^+[u_t^+(W_{t+1})])] \tag{2.2.1}$$

Both u_t^- and u_t^+ are continuous and twice differentiable (consistent with Solow's convention). The investor's (outer) preferences are over the time of resolution of the unknown wealth tomorrow and can be linear, concave or convex (see 2.2.1 for more). His (inner) preferences over wealth are described by a CARA utility function. So that the result is comparable to VN&V's, CARA preferences must be used despite the common use of CRRA preferences in macro finance literature, as authors can normalise using the growth component. CRRA preferences could alternatively be used but the model would be harder to solve.

$$u(W) = -exp(-\rho W_{t+1})$$

ρ is the absolute risk aversion coefficient and W_{t+1} is the wealth generated next period once the risky asset return has been realised. The wealth process, taken from Admati (1985), is below:

$$W_{t+1} = W_t r + q_t'(f_{t+1} - p_t r). \tag{2.2.2}$$

The excess return earned on each risky asset is $(f_{t+1} - p_t r)$. Multiplied by the units of risky assets held, q_t, it gives the total excess return on all risky assets. This figure added to the wealth that would be generated by dolloping all wealth into the risk free asset gives the wealth precipitated for the next period.

r is the risk free return which is known and constant and p_t is the vector of risky asset prices in t. It is exogenous and known. q_t is the quantity of shares that the investor chooses of each risky asset in t, a transposed vector. $p_t q_t$ is expended on risky assets and the remaining wealth is tossed into the risk free asset. f_{t+1}, unknown until $t+1$, is the risky assets return vector, about which the agent forms beliefs. The data generating process for returns is:

$$f_{t+1} = \theta + \epsilon_{t+1} \tag{2.2.3}$$

where $\epsilon \sim N(0, \Sigma_\epsilon)$.

Every period, the risky return is a constant (underlying parameter) plus a random, zero-mean shock. The investor learns about the constant term θ by drawing a vector of signals η_t. The signal is noisy but unbiased, $\eta_t = f_t + e_\eta$ with $e_\eta \sim N(0, \Sigma_\eta)$. As his information set grows, his beliefs about the underlying parameter become more accurate, his uncertainty (variance of the mean belief) tends to zero and his belief about the variance converges to the true variance, that of the shocks. The formulas for updating the expectations of the mean, uncertainty and variance[10], respectively, based on the Kalman Filter and no information acquisition are:

$$\mu_{t+1}^{-} = \mu_t^{-} + K_t(f_t - \mu_t^{-}),$$

$$\sigma_{t+1}^{-} = \frac{\sigma_t^{-}\Sigma_\eta}{\sigma_t^{-} + \Sigma_\eta}$$

$$\Sigma_{t+1}^{-} = \Sigma_\epsilon + \sigma_t^{-}$$

where $K_t \equiv \frac{\sigma_t^{-}}{\sigma_t^{-}+\Sigma_\eta}$, the Kalman gain. The same formulas when endogenous information acquisition is switched on are below.

$$\mu_t^{+} = \mu_t^{-} + K_t(\eta_t - \mu_t^{-}) \tag{2.2.4}$$

$$\sigma_t^{+,i} = \frac{\sigma_t^{-}\Sigma_\eta}{d_t^i\sigma_t^{-} + \Sigma_\eta}. \tag{2.2.5}$$

$$\Sigma_{t+1}^{+} = \Sigma_\eta + \sigma_t^{+}. \tag{2.2.6}$$

These formulas are derived in the appendix, section A.1.4. The investor's belief about risky returns is the current expectation plus some coefficient, K_t, times the error in the previous expectation. It is like an adaptive expectations formula. A subscript t indicates the time period of the variable in question and a negative superscript (-) indicates the variable is a prior not a posterior (+). In equation (2.2.5), d_t^i refers to the number of signals acquired about risky asset i. It is subject to the budget constraint

[10] These formulas are for updating information about an individual asset. Matrix forms could be used to describe all assets together.

$\sum_i d_t^i \leqslant D$, meaning signals are costly not in terms of wealth rather the number of signals that can be drawn is limited. The constraint could intuitively be considered the limited time available in the day for gathering pertinent information.

Sub periods

A sub period contains exactly one information update, either exogenous or endogenous. There are two sub periods in each time t. The first is denoted by the superscript $^-$, to denote prior values and the second by $^+$, for posteriors (prior and posterior to endogenous information acquisition). The same actions are available in each $^-$ & $^+$ across all t.

In sub period one the investor learns about the vector of risky asset payoffs, f. The agent is endowed with a prior belief that $f_t \sim N(\mu_t^-, \Sigma_t^-)$. He knows the true variance, Σ_η, so he learns only about the underlying mean, θ. The agent allocates his learning capacity, L, between the risky assets by picking the signals drawn about each, a vector d_t^i, based on his priors.

In sub period two the agent receives a vector of signals, η_t, that corresponds to his learning allocation, d_t^i. Subsequently, he updates his beliefs and chooses his portfolio based on his posteriors by picking q_t, the units of each risky asset. The first sub period of the next period will reveal the wealth earned by his portfolio, W_{t+1}.

The agent must observe signals in order to update prior to posterior beliefs within a period. He must observe risky asset returns in order to update a posterior in the previous period to a prior in the current period. Ipso facto the value function of a Bellman equation in which choice variables are picked in the first sub period of the current period is conditional on the current period prior mean and variance.

To reiterate the notation, a positive sign signifies that a mean or variance is a posterior whereas a negative sign signifies a prior. The mean or variance subscript indicates the period to which it belongs. For example μ_{t+1}^+ and σ_{t+1}^+ would be the posterior mean and uncertainty for the period one after the current. The posterior values in period t become the prior values in the following period, $t+1$, once the agent has updated again using the observed risky asset returns, f_{t+1}.

Expected utility

The form of u_t^- determines the preferences over the time when uncertainty is resolved. An agent with a linear u_t^- is indifferent about when the return is revealed. When learning the agent, with u_t^- either convex or concave, sees $E_2[u_2(W)]$ as random (the signal received is uncertain too). As u_t^- becomes more convex, the agent prefers an earlier resolution of uncertainty.

To see this, Jensen's inequality indicates that $E_t^-[u_t^-(E_t^+[u_t^+(W_{t+1})])]$ increases, where $E_t^-[u_t^-(W)]$ is treated as a random variable, as u_t^- becomes more convex. Also, $E_t^-[u_t^-(E_t^+[u_t^+(W_{t+1})])]$ decreases as u_t^- becomes more concave. So an increase in concavity means an investor prefers a later resolution of uncertainty. This concept is described further in chapter 2 of Chew & Epstein (1989).

Learning choice

The learning choice of the investor is modelled similarly to VN&V. The main difference is that they constrain the increase in the precision (variance inverse) of the agent's beliefs whereas this paper restricts the number of signals acquired.

Information choice is influenced by the amount of information acquired across all assets and the way that information is allocated between them. With enough learning capacity the investor could immediately eliminate all uncertainty and the information allocation would not matter. This possibility is negated because the total signals acquired across all assets must not exceed the finite learning constraint, L, effectively limiting the number of signals drawn. It is not a financial or wealth cost but more like a time restriction; an investor has limited hours each day to gather information.

$$\sum_i d_i \leqslant L$$

Learning reduces the uncertainty about the return on a specific risky asset. Prior variances are mapped into posteriors using the Kalman filter described earlier. When the investor makes his learning choice it is an irreversible decision. To illustrate, imagine that each signal drawn is a piece of paper with a return written on it contained in a specific pot for the asset learned about. Once he draws he cannot change his choice or redraw in the same period.

Two assumptions from VN&V are used. Firstly, signals are independent, reducing uncertainty only about the chosen asset. So the agent does not choose some correlation structure between the assets. Secondly, the investor is forbidden from forgetting information. He cannot increase the prior variance of some assets in order to reduce it on others. Explicitly, the posterior uncertainty must not exceed the prior uncertainty. Mathematically, $0 < \sigma_i^+ \leqslant \sigma_i^-$.

2.2.2 The Bellman equation

As the model spans an infinite horizon, the investor's maximisation problem can be described by a Bellman equation. This is seen below. Note that, as the maximisation is taken in the current period (t) in the first sub period ($^+$), the agent has not observed signals nor risky asset returns. Therefore the value function is conditional on current period prior variables.

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t, \Sigma_t^+\}} || E_t^- \{u^-(E_t^+\{u^+(W_{t+1})|\mu_t^+, \Sigma_t^+\})|\mu_t^-, \Sigma_t^-\} \\ +\beta E_t^- \{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\} \end{aligned} \tag{2.2.7}$$

Subject to:

- $W_{t+1} = W_t r + q_t'(f_{t+1} - p_t r)$
- $f_{t+1} = \theta + \epsilon_{t+1}$ where $\epsilon \sim N(0, \Sigma_\epsilon)$.
- $\sum_i d_t^i \leqslant L$
- $\sigma_t^+ = \frac{\sigma_t^- \sigma_\eta}{d_t^i \sigma_t^- + \sigma_\eta}$
- $\Sigma_t^+ = \sigma_t^+ + \Sigma_\epsilon$
- $\sigma_{t+1}^- = \frac{\sigma_t^+ \sigma_\eta}{\sigma_t^+ + \sigma_\eta}$
- $\Sigma_{t+1}^- = \sigma_{t+1}^- + \Sigma_\epsilon$
- $\mu_t^+ = (1 - K_t)\mu_t^- + K_t \eta_t$
- $\mu_{t+1}^- = (1 - K_t)\mu_t^+ + K_t f_{t+1}$
- $0 < \Sigma_{t,i}^+ \leqslant \Sigma_{t,i}^-$
- $\eta_t = f_{t+1} + e_\eta$ with $e_\eta \sim N(0, \Sigma_\eta)$

The investor selects the choice (control) variables to maximise his utility over time subject to the state variables and the constraints he faces. The optimal choice over each control variable will be a single policy function specific to each which is a functional equation. That is a function of functions. The policy function for a specific choice variable maximises the investor's objective function in terms of that choice variable. The steady state for each choice variable indicates the amount of it that the model will converge to, diverging from which would not be optimal.

In this model, the state variables, in the current period, are:

- W, previous period wealth (and the f that determines it) and
- μ^- and Σ^-, the investor's prior expectations.

The choice variables in this model for the current period are:

- q, the risky asset allocation and
- d^i, the number of signals acquired about an asset.

The endogenous variables are:

- μ_{ii}^{+}, the posterior mean and
- W_{t+1}, the currently unknown wealth next period which depends on q, d^{i}, μ^{-} and Σ^{-}.

2.2.3 What comprises a solution?

A solution to this model consists of two things, a learning allocation, d^{i}, and a risky asset allocation, q, both in the forms of time-invariant policy functions that maximise the investor's infinite horizon problem subject to the constraints. The learning allocation for each asset that determines the posterior variance, $(\Sigma_i^{+})^{-1}$, is subject to the given capacity constraint, L, and the prior variance for each asset $(\Sigma_i^{-})^{-1}$. The quantity of risky assets purchased, q^{*}, depends on the posterior mean and variance and is constrained by wealth inherited from the previous period. The risk free asset wealth allocation is the remaining wealth.

2.3 Exogenous learning

This section presents the dynamic model with the agent's ability to endogenously acquire information switched off rendering observing realised returns as the sole way to learn. This is the Exogenous learning version with only portfolio allocation chosen by the agent. It is useful to compare the result to the static benchmark to identify the effects of using a dynamic model on portfolio choice. To make this comparison with VN&V's result, CARA preferences must be used despite the common use of CRRA preferences in macro finance literature. CRRA preferences could alternatively be used but the model would be harder to solve.

2.3.1 The Bellman Equation: exogenous learning

The Bellman equation describing the exogenous version of the agent's problem is below. The maximisation is over only portfolio allocation, q, and subject only to prior beliefs.

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t\}} || E_t^- \{ -exp(-\rho W_{t+1}) | \mu_t^-, \Sigma_t^- \} \\ + \beta E_t^- \{ V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^- \}. \end{aligned} \tag{2.3.1}$$

Subject to:

- $W_{t+1} = rW_t + q_t'(f_{t+1} - p_t r)$
- $f_{t+1} = \theta + \epsilon_{t+1}$ where $\epsilon \sim N(0, \Sigma_\epsilon)$.
- $\sigma_{t+1}^- = \frac{\sigma_t^- \Sigma_\eta}{\sigma_t^- + \Sigma_\eta}$
- $\Sigma_t^- = \sigma_t^- + \Sigma_\eta$
- $\mu_{t+1}^- = (1 - K_t)\mu_t^- + K_t f_{t+1}$.

The optimal diversified ('Markowitz') static portfolio is below, adapting VN&V's notation. Solving the maximisation problem above with one period renders the result. Comparing the static and dynamic exogenous learning portfolios later helps identify the impact of the dynamic model.

$$q_t^{div} = \frac{\mu_t^- - p_t r}{\rho \Sigma_t^-} \tag{2.3.2}$$

where div signifies a diversified portfolio. The (Markowitz) investor chooses this portfolio when he diversifies but does no learning. This benchmark is uninformed relative to rational expectations. He holds more units when the prior expected return, $\mu_t^- - p_t r$, increases or risk aversion or uncertainty decreases.

This result is contained in the steady state portfolio when the agent with CARA preferences learns exogenously so a comparison reveals the effect of the dynamic model on portfolio choice. Equation (2.3.3) containing the steady state definition and portfolio are below.

DEFINITION: Learning converges when

$$\sigma_t^- = 0.$$

Learning has converged when the agent is completely confident in his belief about the underlying mean of the distribution of returns. Hence he acquires no more information. This takes place in the limit when $T \to \infty$ and the Kalman gain places no more weight on new information. It implies that in the limit:

- $\Sigma_t^- = \Sigma_\eta$
- $\mu_t^- \to \bar{\mu} = \theta$
- $W_t \to \bar{W}$.

Proposition 1. *The steady state portfolio of risky assets in the exogenous learning with CARA preferences case is the following.*

$$\bar{q} = \underbrace{\frac{\bar{\mu} - \bar{p}r}{\rho\Sigma_\eta}}_{S/s\ myopic\ portfolio} + \underbrace{\frac{\beta r}{\rho\Sigma_\eta}}_{S/s\ hedging\ demand} \tag{2.3.3}$$

Equation 2.3.3 is found by substituting the process for wealth into 2.3.1, taking the first order condition in terms of q_t, using the Benveniste-Scheinkmann condition to find the Euler equation and solving for the steady state (derived in appendix, section A.1.1).

The result differs from the Markowitz portfolio, equation (2.3.2), in two ways, firstly through the fact that the excess return and variance beliefs are now in their steady state forms, showing that the investor is now informed. In the steady state, beliefs have converged, meaning $\sigma_t^- = 0$, $\Sigma_t^- = \Sigma_\eta$, $W_t = \bar{W}$, $p_t = \bar{p}$ and the agent has learned the true mean of the asset return distribution, $\bar{\mu} = \theta$. Thus he uses the true parameters of the asset distribution to allocate his portfolio. Relative to an uncertain investor, $\Sigma_0^- > \Sigma_\eta$, he will be more certain and hold more risky assets. Relative to a pessimistic investor, $\mu_0^- < \theta$, he will be more optimistic and hold more risky assets. However, one might ask if the agent's beliefs will become certain ($\sigma_t^- = 0$) and he stops learning before discovering the true mean ($\bar{\mu} = \theta$). That is, will the investor learn the underlying mean before he stops processing new information? In the face of that question, see section A.1.1 for a proof that mean beliefs converge.

Equation (2.3.3) secondly differs from (2.3.2) by the addition of a hedging demand (in its steady state form) scaled by the true rather than the prior variance, Σ_t^-. This is a key finding: Brandt (2010) shows IID to be a sufficient condition for dispelling hedging demands; state variables are uninformative about future returns, but the result shows that a learning mechanism (here exogenous) is sufficient for them to exist in a multi-period model. The learning mechanism stores information about returns in the state variables and with an ability to hedge, the investor uses that information to spread risk over time. That defines a hedging demand. Even though learning has converged and the conditional distribution of returns no longer depends on returns today, the distributions of returns today and tomorrow are the same conditional on the investor's converged information set. Notably, hedging demands remain in the steady state. Hence they are not a learning effect, they are purely the investor's attempt to diversify against underlying risk using the information learned. Hedging demands may exist in the transition dynamics too and represent an attempt to spread uncertainty (about returns) as well as underlying risk over time.

Through hedging demands the agent no longer invests myopically but accounts today for possible fluctuations in investment opportunities tomorrow. He "over invests" (in Brandt's, 2010, words) relative to the diversified (Markowitz) portfolio to hedge against the (time dimension of) risk that returns will deteriorate between today and tomorrow. The alteration in the investor's behaviour illustrates the significance of opening the infinite horizon channel with an exogenous learning mechanism.

Consider the positive sign of the hedging demand: The investor takes more risk than the myopic investor, given the same investment opportunities today, by allocating a higher portion of his portfolio to risky assets from the risk free. The hedging demand increases with the risk free rate, which, in conjunction with the positive sign, is interpreted like an income effect. A rise in the risk free rate guarantees the agent to earn more tomorrow; the minimum expected return tomorrow is r because he can forgo risky assets if he believes their returns will be inferior. The minimum return between today and tomorrow is higher so he moves some of those resources[11] to today attempting to exploit positive excess returns today. Further, even if the investor believes that risky returns will diminish from today to tomorrow then the higher risk free rate ensures his expected wealth will be higher by the increase in r. A rise in r also has a countervailing substitution effect (see the numerator of the myopic component of (2.3.3)) that excess returns today diminish. Lower expected excess returns today mean that the agent substi-

[11] These resources are not wealth, he does not borrow against the risk free return. The resources are a willingness to bear risk allowing him to hold more risky than risk free units today.

tutes away from risky assets and towards the risk free asset in the current period through the myopic component. Hence the dynamic investor is motivated by an income effect (investing today becomes cheaper) that increases the units of risky assets he holds today and counteracts the myopic motive to substitute towards the risk free asset (risky assets become relatively more expensive). In all, the hedging demands smooth wealth between today and tomorrow. Puzzlingly, even when today's expected excess return is negative and he should sell, the hedging demand causes the agent to buy more units. This is the effect of being able to smooth wealth, that an informed investor will buy more bad assets than seems sensible because hedging opportunities and his knowledge allow him to spread wealth over time. Additionally, the hedging demand will increase when the agent becomes less risk averse (ρ falls) and more patient (β increases). In fact, it does not depend on the Markowitz portfolio at all. [12]

Equation (2.3.3) shows the importance of the infinite horizon that opens the channel through which hedging demands appear and crucially it shows that the myopic and diversified portfolio of standard theory is no longer optimal under an exogenous learning technology. Instead a dynamic strategy becomes superior and the portfolio contains more risky assets. The evidence for this is the hedging demand that is appended to the myopic solution but also, of lower importance, the myopic solution is morphed into its steady state form. Combined, they mean that the investor appears to take a riskier position than the Markowitz investor from the perspective of one who does not learn and invests myopically. In fact, it is only riskier from the perspective of the non-learning investor but not the exogenous learner. This exogenous learning result is reused in section 2.4 to interpret the endogenous learning solution because the latter builds on the former.

[12]Both sides of all of these results could be divided by wealth today, W_t, and multiplied by the price of the risky assets today, p_t, to give $\alpha^* \equiv \frac{p_t q^*}{W_t}$ as the subject, the optimal share of wealth devoted to the risky asset.

2.4 Endogenous learning: Preference for late resolution of uncertainty

This section looks at the case when the investor prefers for uncertainty to be resolved late (logarithmic "outer" preferences), which, along with CARA "inner" preferences, makes a closed-form decision rule obtainable, this is not so for other time resolution preferences, so the decision rules for optimal portfolio and signal allocations are found. A further advantage for using CARA preferences is that the result is comparable to the static one found by VN&V and this is a motivation for not using CRRA preferences, which are common in Macro finance literature. The optimal portfolio is presented and compared to the exogenous solution (section 2.3), showing how endogenous information acquisition affects portfolio choice in the transition dynamics, when the distribution of returns tomorrow remains non-independent of today's. Indeed there is a difference between the two and this is labelled as the learning factor.

2.4.1 When u^- is concave

In this case, the outer utility is concave which means that the investor prefers a late resolution of uncertainty. The inner utility function is CARA and hence exponential. ρ is the risk aversion coefficient.
$u_t^-(y) = ln(-x)$ and $u_t^+(W_{t+1}) = -exp(-\rho W_{t+1})$

The per period utility function, using (2.2.1) is:
$U_t = E_t^-\{ln(E_t^+\{exp(-\rho W_{t+1})\})\}$.

Also, the Bellman equation, derived in section A.1.2, is:

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t, \Sigma_t^+\}} || - \rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^- q_t \\ + \beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}.$$

The Diversified Portfolio

The diversified portfolio of standard theory with no learning and CARA preferences is the expression below, akin to (16) in VN&V.

$$q_t^{div} = \frac{\mu_t^- - p_t r}{\rho \Sigma_t^-}$$

This is what could be called the 'Markowitz portfolio'. A rise in uncertainty reduces the allocation of wealth to the risky assets by the investor whereas a rise in the prior expected return increases his position in the risky assets. Also, a rise in risk aversion quite plainly reduces the amount the agent exposes his portfolio to risky assets.

The Optimal Portfolio

The Bellman equation, given above, for the concave-CARA combination of preferences can be used to uncover the optimal portfolio of risky assets (the decision rule), q^*. To find it, substitute the process for wealth (2.2.2) into the Bellman equation, take the first order conditions in terms of q^* and apply the Benveniste-Scheinkmann condition before solving for the optimal portfolio. This produces the optimal portfolio below and the derivation is in the appendix, section A.1.2.

Proposition 2. *When the investor's preferences are described by concave outer utility and CARA inner utility and investment opportunities are both stochastic & hedgable, then the optimal portfolio is similar to the diversified portfolio with a positive hedging demand.*

$$q_t^* = \frac{\mu_t^- - p_t r}{\rho \Sigma_t^+} + \frac{\beta r}{\rho \Sigma_t^+} \tag{2.4.1}$$

The optimal portfolio with a preference for late resolution is almost exactly q^{div} but plus a term which is the discounted value of investing one unit in the risk free asset (the hedging demand) and scaled by the posterior variance rather than the prior. It also looks the same as the exogenous learning result, 2.3.3, but notice two differences. Firstly, it is not in steady state values and secondly, both terms are scaled by the posterior variance rather than the prior variance. Regarding 2.4.1, an increase in the prior expected return still increases the wealth allocation to risky assets whilst increased risk aversion and imprecision both still decrease it. The second term is the hedging demand, which shows, as in the exogenous learning case, that the optimal decisions in the static and dynamic models differ. The first term incorporates the prior mean and the posterior variance. Without engaging in any learning, q^{div}, as VN&V label it, is actually the portfolio that an investor would choose if he were forming a diversified portfolio according to standard portfolio theory, the Markowitz portfolio. A necessary condition for recovering q^{div} is that the hedging demand is zero. Also, a second condition would be that the agent does no learning and hence acquires no signals. This entails that $d_i = 0$ and ensures that $\Sigma_t^- = \Sigma_t^+$ (also $\sigma_t^- = \sigma_t^+$). To see the effect of the learning capacity on the portfolio choice, it is helpful to multiply by $1 = \frac{\Sigma_t^-}{\Sigma_t^-}$. That is the ratio between the prior variance and the posterior variance. The optimal portfolio would then be the next equation[13].

$$q^* = \frac{\Sigma^-}{\Sigma^+}(q^{div} + H.D.) \tag{2.4.2}$$

This result is essentially the Markowitz portfolio plus a hedging demand (Merton, 1973 & Brandt, 2010) scaled by some factor. It is also like the Merton (1973) portfolio scaled by the same factor. This factor is the ratio of the prior to the posterior variance, re-

[13] The result relies on endogenous learning. For the intuition it could be more plausible to imagine the investor employing both endogenous and exogenous learning techniques but exogenous learning has no impact on the derivation of the decision rule. By itself, exogenous learning is sufficient for convergence of learning. With only endogenous learning, the possibility of many kinds of degenerate distributions of beliefs arises. Kearns (2016) studies this combination of learning technologies.

ferred to in this paper as the learning factor - the extra risk the investor is willing to bear through the uncertainty reduction he gains via endogenous acquisition of information. Essentially when the investor expects to learn something he allocates more wealth to that asset (holds more units) but he does not when he expects to learn nothing. Note that, under endogenous learning, the posterior variance is smaller than the prior variance and never less than one on account of the no-forgetting constraint. The learning factor does not exist in the steady state as equation (2.4.3) shows (the investor has learned the true distribution in the steady state and is then in effect a Rational Expectations investor). This is a key point. Whilst he is learning, the distribution of returns tomorrow depends on signals acquired today. It is not independent. So the agent is willing to invest more aggressively in assets for which he gains a reduction in uncertainty. Further, the uncertainty reduction entails a more aggressive portfolio than either the Merton or Markowitz investors, shifting his investments more towards assets that appear risky to other investors. Merton's agent holds a riskier portfolio than Markowitz's agent because he uses a dynamic strategy rather than a myopic one. The agent who also acquires signals endogenously takes even more risk than that (from the perspective of other investors). Hence learning tilts the portfolio towards risky assets from the prior perspective.

As (2.4.2) is (2.3.3) scaled by the learning factor, $\frac{\Sigma_t^-}{\Sigma_t^+}$, the implication is that not only are the exogenous and endogenous results different from Markowitz's portfolio, but they are also different to each other. So whilst exogenous learning provides a mechanism for investing dynamically, endogenous learning is a wholly new channel through which uncertainty reduction fuels the investor's risky asset purchases. In fact, it is one through which the the uninformed agent invests more aggressively than the informed. Through acquiring signals the uninformed investor knows he will be less uncertain about returns tomorrow as well as those today and so bears more risk in his portfolio allocation, trading that uncertainty reduction to today, whereas the informed agent there is never any uncertainty change.

Conditions for recovering the diversified portfolio

An infinite horizon combined with the exogenous learning mechanism was shown earlier to be a sufficient channel for hedging demands to arise. The investor's portfolio is further complicated beyond the Merton solution by the endogenous information acquisition technology. Here the two conditions for recovering Markowitz's solution are stated and explained by juxtaposing the Markowitz portfolio and equation (2.4.2). Both are necessary and together they are sufficient. They identify why an infinite horizon and endogenous information acquisition cause an investor to divert from the standard diversified portfolio. The optimal portfolio with endogenous information acquisition diverts through a hedging demand and a scaling factor, the ratio of the posterior to the prior precision. The conditions follow.

CONDITION 1: $\frac{\Sigma_t^-}{\Sigma_t^+} = 1$

So $d_i = 0$ *must be true, meaning that* $\sum_i d_i = 0 < L$ *and the agent learns nothing within the period. This condition will turn the endogenous learning portfolio into the Merton one.*

CONDITION 2: $H.D. = 0$

Thus either $r = 0$ *or* $\beta = 0$ *as in the exogenous learning case. This condition will turn the Merton solution into the Markowitz one.*

The no-forgetting constraint ensures that the learning factor multiplies the myopic portfolio and the hedging demand by a factor of at least one. So the first condition to recover q^{div} means that the variance of beliefs does not change and the investor learns nothing within periods. The ratio of posterior precision to prior precision must be one, meaning uncertainty about returns ($\sigma_t^+ = \sigma_t^-$) does not change within periods. Thus, the investor must acquire no signals endogenously. Consequently, the portfolio would revert to the

Merton one (and also the exogenous one, (2.3.3), if learning had converged).

The interpretation of the scaling factor, $\frac{\Sigma_t^-}{\Sigma_t^+}$, is that the agent is aware that he will know more about risky asset returns (both today and tomorrow) through the signals acquired, and hence have higher utility seeing as he likes precision (he dislikes uncertainty), so he is willing to take more risk today. Thus he purchases a higher quantity of risky assets than he would in either the Merton or myopic cases.

The hedging demand increases with the risk free return, which seems unusual. Yet the myopic component decreases with the risk free return. However, the investor knows that no matter what his prior beliefs are tomorrow, the minimum he can earn, the risk free rate, is higher. Hence he is guaranteed greater wealth tomorrow as r rises. The hedging demand means he is willing to trade that guaranteed higher return, like an income effect, against taking more risk today. An increase in the investor's patience raises the hedging demand too meaning he puts more weight on the future. The hedging demand is scaled by the learning factor so acquiring more signals makes the hedging demand more potent. Intuitively: the agent steers his portfolio more towards risky assets than the myopic investor since he is less uncertain about them and they appear less risky to him. Hedging demands fall as precision decreases and risk aversion increases.

The second condition requires either $\beta = 0$ or $r = 0$ to eliminate the hedging demand and recover the diversified portfolio of standard theory. Both of these are extremely restrictive conditions. $\beta = 0$ would entail absolutely no patience on the part of the agent meaning that he would not value future wealth at all and hence hedging demands would be unimportant to him. $r = 0$ would mean that the wealth invested into the risk free asset attenuates fully. These conditions are found by setting the hedging demand equal to zero and solving for each variable. Under either condition, q^* is undefined because Σ_t^+, the posterior variance, is undefined (section 2.4.2 covers this). The investor either wants to buy an infinite amount or sell an infinite amount of the risky asset in question depending on whether the prior expected return is positive or negative.

Late resolution preference with RE beliefs

In Proposition (3), the limit in learning of (2.4.2) is taken to discover to how a Rational Expectations agent would behave. With a late preference for resolution of uncertainty and Rational Expectations then the investor's portfolio choice is given by (2.4.3).

Learning has a twofold impact: it reduces the uncertainty that the investor has about the returns on the asset about which he acquires signals and it ambiguously affects his mean belief. Through the first effect, learning persuades the investor to purchase more of and bear more risk from the asset to which he has allocated signals; he likes certainty. Through the second effect, if the signals acquired give a favourable impression about returns (high absolute excess return), then the investor unambiguously holds more units and portfolio concentration can result from endogenous learning as both effects agree. This mechanism provides a rationale for concentrated portfolios observed in the CFS in the USA. If the mechanism can induce concentration even under a late time resolution preference then the effect could exist and be stronger under preferences for earlier resolution.

Proposition 3. *When the investor has Rational Expectations and his preferences are described by concave outer utility and CARA inner utility and investment opportunities are both stochastic & hedgable, then the optimal portfolio has the same components as that of an uninformed investor.*

$$q_t^{*,RE} = \frac{\bar{\mu} - \bar{p}r}{\rho\Sigma_\eta} + \frac{\beta r}{\rho\Sigma_\eta} \tag{2.4.3}$$

The optimal portfolio under Rational Expectations has exactly the same structure as the optimal portfolio for an uninformed investor (the endogenously-learning investor whose beliefs have not converged): a diversified portfolio with a hedging demand. The

Rational Expectations investor believes that the variance is lower on each risky asset than the uninformed investor does because the uncertainty of his mean belief is zero ($\sigma_t = 0$). Thus he will purchase more risky assets if his mean belief is more optimistic (in absolute terms) than the uninformed investor's belief. So discounting extraordinarily optimistic initial prior beliefs, the uninformed investor will continue to purchase more risky assets as he learns until he achieves Rational Expectations. Interestingly, the steady state portfolio choice for the investor who is indifferent to the time of uncertainty resolution, (2.5.1), is identical to the RE agent's portfolio choice with a late preference for time resolution, (2.4.3), and the steady state portfolio under exogenous learning, (2.3.3). These two results are logical because all learning (with only one source of noise) will converge to the truth and hence the RE set of beliefs.

2.4.2 The Optimal Posterior Variance

The optimal learning choice determines the posterior variance which, as already seen, determines the choice of risky assets. By substituting the latter into the Bellman equation and taking the first derivative in terms of the posterior variance, the first order conditions for the learning choice can be found. Then the Euler equation can be obtained and from it, the optimal posterior variance (which, inverted, gives the optimal posterior precision) can be found.

Proposition 4. *When the investor's preferences are described by concave outer utility and exponential inner utility and investment opportunities are both stochastic & hedgable, then the optimal posterior variance is a combination of the prior expected excess return, discounted risk free return and the risk aversion coefficient.*

$$\Sigma_t^+ = \left(\frac{(\mu_t^- - p_t r + \beta r)(\mu_t^- - p_t r) - \frac{1}{2}(\mu_t^- - p_t r + \beta r)^2}{\beta \rho r}\right)^{\frac{1}{2}} \tag{2.4.4}$$

A graphical example representation can be found in figure A.1 with the parameters set to: $\beta = 0.99$, $\rho = 4$ & $r = 1$. The red line shows the optimal posterior precision, the investor desires a higher precision the smaller the excess return as he wants more certainty to compensate him for a perceived a lower payoff. It tends towards infinity as the excess return tends to zero. The blue line represents the optimal posterior variance. The optimal posterior variance is not zero when there is no excess return because the investor would require a negative variance to compensate him for such a low return to taking on risk, if that were possible.

For investigating the optimal posterior variance, four particular cases can be considered which are described below.

The Markowitz portfolio is greater than the hedging demand

$$(\mu_t^- - p_t r > \beta r)$$

When the prior expected excess return is greater than the discounted risk free return, the posterior variance is positive, $\Sigma_t^+ > 0$. So the investor potentially learns in this case. The lower bound on $(\Sigma_t^+)^{-1}$, the posterior precision, is 0 which is achieved when the discounted risk free return is 0 or when the investor tolerates no risk ($\rho = 0$), setting the denominator of (2.4.4) to zero. So in this case the investor wants no precision at all.

This could be because the investor has no patience ($\beta = 0$) and hence would not learn about the future or it could be because $r = 0$. Both cases would eliminate the hedging demand and lead to a myopic portfolio choice. None risk aversion ($\rho = 0$) has the same effect. In each situation the investor does not care about his payoff uncertainty either because the risk free asset is wholly unattractive (it will attenuate all his wealth) and therefore risky assets dominate or because he is insensitive to risk. In all these cases, as posterior precision would be zero, q^* would be infinite, a corner solution.

The more interesting feature is how the relationship between the prior expected excess return and the discounted risk free return affects the value of the posterior precision. The prior expected return must be greater otherwise the posterior precision will be undefined. Yet posterior precision grows the closer these two factors become. Intuitively the investor cares more about precision when the difference between taking a risk today and being safe tomorrow is tight. Moreover, posterior precision increases with both $\mu_t^- - p_t r$ and βr meaning as the size of potential returns to choose between increases, the investor takes more care over the decision made by desiring more information.

Buying or selling, the investor prefers the risk free asset

$$(|\mu_t^- - p_t r| < \beta r)$$

When the prior expected excess return is less than the discounted risk free return, the posterior variance (and the precision) is undefined (the root of a negative number is undefined) and there is no solution to equation (2.4.4). If the investor expects the risky asset to return less than the risk free asset then he does not consider it worthwhile to learn anything about it. In fact, he would prefer to forget what he knows, if that were possible, to learn more about assets that he expects to return more than the discounted risk free rate. Hence this is a corner solution bounded by the no-forgetting constraint.

The investor prefers short selling

$$(\mu_t^- - p_t r < -\beta r)$$

Once the prior expected return falls to less than negative one times the discounted risk free return the investor becomes interested in the risky asset again (the numerator and denominator are both negative so the precision is defined). Now he believes its return is so negative that it will be worth short selling it (his asset allocation would be below zero, $q^* < 0$). As a result, the expected posterior variance of the risky asset and learning become interesting to the investor again. The closer the prior expected excess return is to the discounted risk free return, the lower he would like the posterior variance to be and the inverse is also true. As $(\mu_t^- - p_t r)$ approaches βr a lower variance is more important in order to check the prior belief so the investor is sure that the risky investment will offer a superior return to the risk free one. In the opposite case, the widening gap between the two returns means that he is willing to accept greater uncertainty in receiving a larger excess over the discounted risk free return.

There is no expected reward for risk-taking

$$(|\mu_t^- - p_t r| = \beta r)$$

When the prior expected excess return is exactly equal to the discounted risk free rate, the investor desires to have a posterior variance of zero. This is quite plain and it is a knife-edge case; he urgently wants to confirm whether or not his prior expectation of the mean is correct due to the absolute proximity of the two returns in his current view and be sure that there is no compelling reason to take on risk. The inverse, the precision, is undefined because the agent would always like to have more information in this case and needs infinite compensation in precision for no excess over the discounted risk free return.

2.5 Indifference to the time of resolution

This section considers the agent who is indifferent to the time at which the uncertain risky asset payoffs are revealed. The neutral preference is driven by the linear outer utility function. The following subsections present the Bellman equation under linear time resolution preferences and concave (CARA) preferences over wealth along with its steady state solution and a rough[14] approximate numerical solution for the decision rule. The results are then compared to the benchmark static solutions of VN&V to test if their agent's myopic behaviours survive in a dynamic context. This comparison necessitates linear time and concave risk aversion preferences, which is why time preferences change from section 2.4.

2.5.1 The Bellman Equation

The agent has linear preferences over the resolution of uncertainty, $u_t^-(x) = x$, which mean that he is indifferent regarding timing, and CARA preferences over wealth, $u_t^+(W_{t+1}) = -exp(-\rho W_{t+1})$, making him risk averse over wealth.

The per period utility function, using (2.2.1) is:
$U(W_{t+1}) = E_t^-[(E_t^+[-exp(-\rho W_{t+1})])]$.

Also, the Bellman equation is, the derivation of which is in section A.1.3:

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t, \Sigma_t^+\}} || - exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t \\ + \frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\ + \beta E_t^- \{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^-\}. \end{aligned}$$

[14] The state & choice variable grids are sparse due to the high computational costs associated with the multiple dimensions of the problem. Nonetheless, the grids are sufficiently big for ruling out a "budget constraint effect" as the mechanism for the results. That is, the numerical environment is not so restrictive as to a priori eliminate strategies that the agent could employ.

Here, the value of the problem to the agent depends on his initial wealth in the current period (W_t) and his prior beliefs about the mean (μ_t^-) and his uncertainty about it (σ_t^-), using which he forms his posterior variance (Σ_t^+). K_t is the Kalman Gain, which is the weight between the prior and error components of the formula for updating the belief about the mean. The agent maximises the value of the problem in terms of the choice variables q_t, the quantity of risky assets purchased (he can short sell), and Σ_t^+, the posterior variance which he manipulates by drawing signals about the risky assets available.

2.5.2 A generalised learning portfolio

In a static model with linear time and concave risk preferences, the agent wishes to have a generalised learning strategy, which means that he gathers signals in such a way as to equate posterior variances. This is the result that Van Nieuwerburgh $\&$ Veldkamp label as Corollary 1 on page 12 of their paper. The portfolio of risky assets that the agent chooses in the static model under this learning strategy and the preferences described above (linear-CARA), using the notation style of this paper, is the following.

$$q_t^{gen} = \frac{\mu_t^+ - p_t r}{\rho \Sigma_t^+}$$

So if the signals that the agent gathers make him more optimistic about the asset in question (the posterior excess return rises) then he will buy more of it, yet if the signal is adverse, he will buy less of it. Moreover, the more signals he gathers about a particular asset, the more of it he wishes to buy because he is more certain (the posterior variance is lower) about its return. Lastly, the more risk averse he is (the greater the risk aversion coefficient) the less risky assets he wants to buy.

The reason that the static-strategy investor chooses a generalised learning strategy and related generalised portfolio is that the Lagrangian problem that can be set up to describe the investor's portfolio maximisation decision is concave in the signal allocation and the constraint is linear. Therefore there is an interior solution from a concave objective and a linear constraint and an interior solution spells out a generalised learning strategy in this scenario. The Lagrangian set up can be seen in equation (9) of VN&V. This means that the investor wants to equalise the posterior variances of assets. The concavity of the Lagrangian depends on the expected covariance between marginal utility of wealth and the second derivative of wealth in terms of posterior variance. If they are more negatively related then the problem is more concave and indicates that the benefits to specialised learning (second derivative of wealth in terms of the posterior variance) are high when the marginal utility of wealth is low. Thus specialisation is not worthwhile for such an investor who consequently generalises his learning.

2.5.3 The steady state portfolio

In the dynamic model of this paper with linear time and concave risk preferences, the agent holds a portfolio is the steady state which is the one below. This result differs from that of Van Nieuwerburgh & Veldkamp generalised learning portfolio firstly through the addition of a hedging demand term. It signifies the importance of the dynamic framework in that an agent wants to optimally invest more today than a purely myopic strategy would suggest so long as he cares about the future ($\beta \neq 0$). It further means that although a rise in the risk free rate today would hinder his ability to earn excess returns today and so reduce his myopic holdings of risky assets, it would increase his holdings under a dynamic strategy through a higher minimum guaranteed return tomorrow.

$$\bar{q} = \underbrace{\frac{\bar{\mu} - \bar{p}r}{\rho \Sigma_\eta}}_{Myopic\ component} + \underbrace{\frac{\beta r}{\Sigma_\eta}}_{Dynamic\ component} \tag{2.5.1}$$

The investor selects his steady state portfolio by allocating wealth to each individual asset depending on the excess of his mean belief ($\bar{\mu}$) once his beliefs have converged

($\sigma_t^- = 0$, meaning that his uncertainty reaches zero) over the expected price times the risk free rate scaled by the true variance and his risk aversion in the myopic component. His hedging demand component depends on the product of the risk free rate and his discount rate divided by the risky asset's true variance (the minimum guaranteed return tomorrow scaled by his impatience and the risk he ultimately bears on the asset).

Recalling equation (2.4.3), it is possible to compare the steady state portfolio choice in each asset to that under RE when the agent prefers a late resolution of uncertainty. One observes that the choice is identical in the hedging component and the same in the myopic component except that the mean belief is a steady state rather than a true value. Yet section A.1.1 in the appendix indicates that mean beliefs will converge as learning does and hence the two would be the same, the RE and the steady state beliefs about the true mean. Thus the two decisions ultimately are congruent and show that decisions over risky assets are unaffected by indifferent or late preferences over time in the limit of learning.

Additionally, equation (2.3.3), the steady state portfolio choice over each risky asset, is absolutely the same as under indifferent time preferences. Pertinently and logically this shows that the steady state is the same for the exogenously- and endogenously-learning agents. That is, once all benefits from learning have been extracted, both types of agent make the same decision[15].

2.5.4 The approximate decision rules

The standard ways to find the decision rules (policy functions) for a Bellman equation that does not have logarithmic utility and Cobb-Douglas constraints or quadratic utility and linear constraints are either value function iteration or policy function iteration. This paper uses the former to solve for the agent's optimal choices over the quantity of risky assets to purchase and the signal allocation. As there are two assets in the numerical solutions there are four decision rules for which the solutions are found. The results

[15]This relies on the fact that signals are true on average.

can be found in A.1.3 and this subsection both highlights some important features of and explains those results. Table A.1 lists the states of the world and the values of the state variables in each. These are determined by all of the possible combinations of wealth (1) at the start of each period and of the prior variances on the two assets (9) which gives 81 states. A selection of the optimal selections of risky asset quantities and signals are contained in tables A.2-A.4. They proceed downwards as the state increases. For each state there are four matrices, one for each of the choice variables. Moving rightwards across a matrix, along the columns, the prior mean of asset 1 increases by one value in the grid, ceteris paribus. Moving down a matrix, the prior mean of asset 2 increases by one value in the grid, ceteris paribus. Please see the three small tables at the beginning of section A.1.3 for the state variable grids, the values of the parameters and the transition matrix used.

Generalisation is not uniformly utilised

The first result stems from noting that VN&V's static model indicates, due to the concavity of the Lagrangian, that an investor with linear-CARA preferences wishes, always, to equalise posterior variances between assets. Refer to Corollary 1 in VN&V, page 12 for comparison purposes. So when the agent initially knows more about asset 1 than asset 2, he will want to learn more about asset 2 and with sufficient learning capacity he would equate them. Additionally, in the static model, VN&V's Corollary 1 shows that this result does not depend on the prior mean. However, the numerical results in this paper do not conform to these static model results. Firstly, he does not always prefer to equalise posterior variances and secondly the agent does take into account the prior mean when making his signal allocation choice. These results are summarised in Facts 1 & 2.

Fact 1. *Within the numerical results, the agent sometimes prefers a specialised learning strategy over the generalised one predicted by the benchmark static model.*

Consider the tables of results A.2-A.4, which describe states 1-9 (states 10-81 follow similar patterns in groups of 9). From states 1-9, the prior variance on asset 1 rises whilst than on asset 2 remains constant, meaning the agent is more confident about asset 2's returns in all but state 1. An agent with linear-CARA preferences in VN&V's static model would prefer a generalised learning strategy whereby the signals are spread evenly[16] making the posterior variances as close to each other as possible, so the matrices representing the signal choices for the two assets would always be full of the number 1, yet this is not so. Hence the dynamic model results deviate from the static model ones.

Instead of the agent choosing a generalised learning strategy throughout the agent sometimes prefers to specialise in learning about one asset. There are three ways in which the investor specialises that are detailed next. However, the reason for specialisation is essentially that the agent has time-consistent beliefs. When he expects a strongly superior Sharpe ratio[17] on one asset he allocates signals to it because he anticipates to learn nothing new about its underlying mean. That is, he believes his prior mean to be true and expects his posterior mean to be identical to it. Yet by learning, he knows that the posterior uncertainty about the asset will fall thus improving his post-learning expected Sharpe ratio. Thus by specialising in the more attractive asset (higher expected Sharpe ratio) he believes he will improve his future investment opportunities.

Firstly, specialisation can be seen when the prior variances of assets i and j change but their prior means do not. Observe the bottom right quadrant of tables A.2-A.4, states 1-9, for signals acquired about asset 1, where prior variances rise on asset 1 from states 1-9. In state 1, the prior variances are identical for the two assets and the lead diagonal contains the value 1, indicating generalisation. Above the diagonal asset 1's prior mean is higher and the investor specialises in the first asset whereas below the diagonal, the narrative is the inverse. In subsequent states, the prior variance rises and no longer

[16] The calibration of the dynamic model in this paper means that, with two available signals, the agent's best allocation of signals for making posterior variances as close as possible is one for each asset.

[17] The expected Sharpe ratio of an asset can be thought of as the excess return per unit of risk or the bundle of prior mean and variance together of that asset.

does the agent specialise in the asset that has the higher prior mean. As asset 1's prior variance rises, its higher prior mean is no longer so attractive so the agent reallocates signals to asset 2. This means that specialisation (2) in asset 1 & generalisation (1) are increasingly replaced by generalisation (1) & specialisation in asset 2 (0). In state 9 reallocation occurs to such an extent that the agent specialises in asset 1 only when the prior expected excess return on asset 2 is 0. Generalisation would mean that when prior means are identical (the lead diagonal) the agent would allocate one signal to each asset and in a model with more learning capacity he would devote all signals to the asset with the higher prior variance (lead diagonal full of 2 in states 2-9). Therefore the investor does not spread his signals to increase his belief precision about asset 1 when he initially knows less about it, he chooses to gain further knowledge about asset 2. So clearly, in the dynamic model, there is scope for the agent to specialise under linear-CARA preferences, namely when he is optimistic about an asset about which he is more certain of the payoffs.

Secondly, specialisation can be seen when the prior mean of assets i and j change but their prior variances do not. Consider table A.2, state 1, and particularly the signal choices for asset 1. Moving rightwards across each row sees the prior mean rise on asset 1. More signals are devoted to asset 1 at the extremities of each row than in the middle. Realising that the excess return falls on asset 1 in the middle of a row explains this pattern. What is more, moving down a column sees the prior mean on asset 2 rise. This means that the excess return on asset 2 falls in the middle of a column and rises at the end. Hence more signals are drawn for asset 1 in the middle of each column. The reverse of these two patterns is true for the allocation of signals to asset 2.

Thirdly, specialisation can also be seen in row 5 of the matrix for signals about asset 1 and in column 5 of that for asset 2, again in state 1. To understand these results, it is important to realise that the fifth grid point of the prior mean state space has a value of 1. This means that when the agent holds this belief about an asset he expects its excess return to be zero and hence it is not worth buying or selling any of it. He is uninterested in trading it. Thus, this case of specialisation is trivial but important to clarify. Therefore, no matter the state, when the agent expects either asset to give no excess return he sets $Q_t^i = 0$, where $i = 1, 2$, so column 2 of each matrix representing the quantity of asset 1 purchased and row 2 of those representing that of asset 2 purchased are always filled with zeros. When $Q_t^i = 0$, μ_t^i & $\Sigma_t^{+,i}$ drop out of the Bellman equation. Consequently, an asset for which the expected excess return is zero is uninteresting to learn about. That is why the second column of all matrices representing the signal allocation to asset 1 and the second row of those relating to that of asset 2 are always 0. A special case occurs when the agent expects both assets to give zero-expected returns because the quantity purchased of both is zero and he gathers no signals at all. Hence he appears to waste his learning capacity (the corresponding signal choices for assets 1 and 2 do not sum to the constraint of 2) and the reason is that all terms involving the vector of risky asset quantities drop out of the Bellman equation, equation (2.5.1). Hence signals no longer have any effect on the value of the agent's problem.

Decidedly, the investor has a motivation to specialise and then subsequently concentrate his portfolio in the asset learned about. This occurs not always but only when the expected Sharpe ratio[18] of one asset is sufficiently superior to that of the alternative. Suppose it is asset 1 that is substantially more attractive. Its comparatively superior Sharpe ratio persuades the investor that not only is worth specialising his learning in asset 1 but he is so confident in the better excess returns it offers so as to dedicate his whole portfolio to it and so leave no room for asset 2. The next subsection deals with concentration further.

Another thing to note is that the signal allocation choices vary depending on the prior mean, the matrices are not homogenous. In the static benchmark there is no dependence at all, see again Corollary 1 in VN&V. Hence this is a second significant way in which the dynamic learning decision of this paper differs from the optimal static choices.[19] This outlines Fact 2.

[18]The expected Sharpe ratio of an asset can be thought of as the excess return per unit of risk or the bundle of prior mean and variance together of that asset.

[19]In fact, signal allocation depends on the excess return that the investor expects, which is why fewer signals are allocated to an asset whose prior mean is close to 1 than one whose prior mean is far from 1.

Fact 2. *According to the estimation, the agent's learning allocation choice depends on the prior mean he holds, unlike in the benchmark static model where the choice is independent.*

It is not simply the prior mean that the agent takes into account but rather the expected excess return: $\mu_t^i - p_t^i r$, where $i = 1, 2$. Consider state 1 again. When the agent expects to earn no excess return on asset 1, the fifth column, signals have no impact on the value of his problem so he acquires no signals. If the prior mean moves upwards or downwards by one grid point then the expected excess return increases in absolute value. The agent now considers it worthwhile to both learn about and trade it (the quantity of asset 1 is always 0 in column 5). The absolute value of the expected excess return is important for the quantity of signals purchased and the magnitude of wealth devoted to the particular risky asset whereas the sign of the expected excess return is important for whether the agent buys or short sells the asset. The sign does not impact the learning allocation, which is why rows 1-4 $\&$ 6-9 and columns 1-4 $\&$ 6-9 are always mirror images of each other.

Portfolio concentration in a dynamic model

The numerical results display a choice by the investor to concentrate (select only one of the two risky assets) his portfolio choice. View table A.2. In state 1, concentration occurs in asset 2 whenever the quantity of asset 1 held is 0.00. This is evident as the excess return on asset 1 falls towards 0 (towards the middle of each row). Specialisation occurs below the lead diagonal of the bottom right quadrant and concentration is a subset of this; it occurs below the main diagonal but less frequently. Thus arises the intuition that concentration arises from specialisation motives. That is, when the prior expected Sharpe ratio is sufficiently attractive on one asset compared to the other, the investor will specialise his learning in it and this can induce him to concentrate his portfolio in the same asset.

In the static model of VN&V, underdiversification occurs when the investor selects a specialised learning strategy and this requires a strict preference for an early resolution of uncertainty, as described in Proposition 2 on page 15 of their paper. In fact their Corollary 2 (page 16) explains that under such a preference for early resolution of uncertainty and an increase in the value of the Lagrangian in the learning capacity, the investor will allocate all signals to the asset with the highest Sharpe ratio. In the dynamic model of this paper, the value function is increasing in the learning capacity but the same outcome, specialised learning and concentrated portfolio choice is attained with indifference to time preferences.

What is the rationale for underdiversification (concentrated portfolio choice)? As the agent has time-consistent beliefs, his expectation of the posterior mean is the prior mean, he anticipates no change in mean expectations from learning. Yet, he knows that each signal drawn will increase the attractiveness of an asset through the diminution of the posterior variance. So this is the rationale for specialised learning, which can then induce concentration through the resulting superior expected relative Sharpe ratio of one asset. The mechanism is also one of simultaneity akin to VN&V; as the investor draws more signals about one asset, its posterior variance falls increasing its expected Sharpe ratio and encouraging him to hold more of it. Then as he desires to hold more of it, he desires to gather more signals about that asset because each signal is informative about more of the portfolio that he wishes to hold. Yet strikingly, this mechanism that arises in this dynamic model under indifference to the time of uncertainty resolution induces concentration only under a strong preference for early resolution of uncertainty in a static model. So this is an interesting development that the preference combination in their paper that leads to a generalised learning strategy ends up with a specialised one in the dynamic model of this paper. This leads to one of the most important results found.

Proposition 5. *In a dynamic model, concentrated portfolio choice can be explained without reliance on a strong preference for an early resolution of uncertainty.*

The static model results of VN&V depended on a strong preference for an early resolution of uncertainty to explain the concentration in portfolio choice. However, the dynamic model used in the present study does not depend on this preference. Instead, it relies on the prior mean and variance that the investor holds about the risky assets. This is very useful because whilst it is not possible to observe the time preferences of agents, it is possible to glean from them their prior beliefs about the risky assets available to them. So this theoretical finding, that concentrated portfolio choice under dynamic investment models depends on the mean and variance of prior beliefs, indicates that it is possible to investigate investor behaviour empirically. Indeed, such data on investor expectations already exists.

The research question this paper attempts to address is why do investors hold concentrated portfolios? In this section, the numerical approximations to the investors decision rules over his portfolio and learning allocations when he is risk averse over wealth and indifferent to the time of uncertainty resolution have been described and analysed. The implications are significant: even in a dynamic setting with the possibility of hedging both wealth and learning an investor may choose to specialise in both his portfolio and signal allocation choices. The model asserts that this will occur when he has a significantly higher ratio of confidence to uncertainty about one asset than another.

2.6 Summary

It is, on the surface, surprising that investors hold portfolios that exhibit a high degree of concentration thus contravening standard portfolio theory. Van Nieuwerburgh & Veldkamp (2010) find a rationale for such behaviour wherein, by a feedback effect, endogenously acquired information about an asset reduces uncertainty about its payoffs, persuading the investor to hold more of it and conveying benefits to specialisation as further information will inform about more of the expected portfolio thus concentrating his portfolio.

This paper employs a dynamic endogenous information acquisition model to test the robustness of their static model results and find the conditions under which the diversified portfolio of standard theory (Markowitz's portfolio) is recovered. The solution to a relevant case, that when the investor prefers a late resolution of uncertainty, is found and benchmarks for dynamic and static models with exogenous learning are found too. That the results diverge from the static model solutions illustrates the relevance of the extension to an infinite horizon. Under an exogenous learning technology (gathering public information only) hedging demands (spreading risk over time) arise showing that myopic and dynamic investment strategies differ through learning despite time independent returns. Under an endogenous learning technology (acquiring private information), not only is a hedging demand present in the transition dynamics but so is a learning factor (a multiplicative ratio showing the factor by which uncertainty is reduced). The latter disappears in the steady state meaning that whilst learning, the agent invests more aggressively assured of greater certainty about returns both today and tomorrow. So not only do dynamic investment strategies matter but so does the ability to acquire signals endogenously.

The second general contribution of this paper is to identify the conditions under which the new results differ from Markowitz's portfolio. Portfolio allocation with exogenous learning reverts to the static choice under two restrictive conditions: Either the risk free return is zero or the investor has no patience whatsoever. Further to these conditions, endogenous signal acquisition must be switched off to recover Markowitz's portfolio under an endogenous learning technology.

Van Nieuwerburgh & Veldkamp find that a strict preference for early resolution of uncertainty is necessary in a static model for specialised learning and concentrated portfolios. This paper uses the same preference structure in a dynamic model and the numerical solution suggests that strict unobservable preferences are not necessary to explain underdiversification. Instead the investor specialises his learning strategy and concentrates his portfolio when he expects a strongly superior Sharpe ratio on one asset than another. Thus the model indicates that underdiversification can be addressed empirically through the reported stock market beliefs of households without relying on unobserved preferences. The result also implies that dynamic decision-making can account for concentrated portfolio choice through the learning strategy employed. Interestingly, in a dynamic setting this result fortifies Van Nieuwerburgh & Veldkamp's feedback effect between the gathering of noisy information and expected portfolio holdings that produces a concentrated static portfolio. Surprisingly, dynamically, concentration can occur under indifferent time preferences that cause the myopic investor to generalise his learning and form a diversified portfolio.

3

Paper 2

3.1 Introduction

The first paper in this thesis investigates the effect of constraints to information processing capacity on the information and portfolio that an investor chooses. A natural extension is to consider when, given such constraints, an investor should choose to acquire such information. Given a rise in the returns of an available risky asset, should the investor respond by acquiring more information, less or even not respond at all? So this section of the thesis concentrates on the responsiveness of demand for information to fluctuations in risky asset returns. The investor responsiveness of the investor is assessed through the change in the number of signals endogenously acquired about a particular risky asset given that there was a positive change in its observed return. Risky asset returns are artificially generated in the model in order to assess their impact on demand for information. In order to relate this study to other work and ground it in a stylised fact, a further motivation follows.

Widespread heterogeneity in beliefs about past stock market returns is documented by the 2014 wave of the PAT€R survey[1]. This is surprising given the availability of historical CAC40 returns. Households' average belief about returns in the past three years is +3.61% with a standard deviation of 12.04 and concerning future returns they expect 1.62% growth over the following five years with a standard deviation of 8.94. The true return was 34.3% growth so agents are also badly informed and many even thought the stock market had fallen significantly.

The aim of this paper is twofold: to study whether or not demand for information is positively correlated with risky asset returns in a model of information acquisition[2] and portfolio acquisition and to investigate the widespread heterogeneity in beliefs observed. See sections 3.1.1 & 3.1.2 for further detail. This study contributes to literature concerned with bounded rationality. It attempts to answer the following research questions: "Is demand for information positively correlated with returns?" and "Can a dynamic information acquisition model with limited learning capacity explain heterogeneity in stock market beliefs?". It finds a theory that can replicate heterogeneity of beliefs when there is no access to public information.

Modelling portfolio allocation with information acquisition has previously been studied by Van Nieuwerburgh & Veldkamp (2010)[3] in a static setting. By extending this framework to a dynamic horizon, this paper can address the relationship between return fluctuations and demand for information and finds the latter to be unresponsive to the former, counter to the findings of Coibion, Gorodnichenko & Kumar (2015). It also allows a theory for the widespread heterogeneity in beliefs despite the public availability of information to be found. A dynamic investor may forgo learning about one asset to acquire signals only about another that he believes has a superior distribution. Specialisation increases his expected utility of future wealth because he expects to hold more of that asset so learning raises certainty about portfolio returns. Observing no returns

[1]This is a survey of French households that is representative of the population by age & wealth.
[2]The agent purchases costly signals that are noisy but informative.
[3]VN&V hereafter

and signals that confirm his priors encourages specialisation and the agent develops an information set different to another agent with opposing initial priors. So with these two assumptions the theory can create heterogeneity in beliefs like that documented.

A model is built in which an investor allocates his wealth can learn exogenously (observing returns) and endogenously (endogenously acquiring signals). Learning exogenously, he observes risky asset returns after allocating his portfolio and updates his beliefs about their underlying distributions. Endogenous learning lets him acquire costly signals prior to allocating his portfolio. This provides more timely information that improves accuracy and reduces uncertainty of beliefs. Signals are not financially costly but the agent has a limited capacity to acquire them. He forms his portfolio by combining two risky assets and one risk less asset. The model is simulated for different initial prior beliefs about the risky assets and three interesting insights are found.

Firstly, demand for information tends to be unresponsive to risky return fluctuations. Stationarity of the underlying distribution means, once he has learned it, he picks a stable learning strategy independent of the state of the world. The static investor does not vary his learning strategy regardless of how much he has learned. This is also true for the dynamic investor who observes no returns and receives signals that support his prior beliefs, which perpetuate his initial decision to specialise. However, the dynamic investor who observes returns utilises a stable learning strategy only once he has learned the underlying distribution. His Kalman filtering learning which weights new information less as time increases drives this. This result cannot add insight to Coibion et al's survey finding that demand for information is (anti-cyclical) negatively related to the business cycle but it does offer a theoretical explanation for why they observe significant dispersion in beliefs, which is the prevalence of permanent specialised learning strategies. That is, two different agents could hold different initial prior beliefs, specialise in learning about different assets and subsequently hold different information sets and beliefs about their past performances.

Secondly, a rational investor (processes information logically) may default due to inaccurate initial beliefs about the underlying distribution of the risky assets. Thus an uninformed, rational agent needs deep pockets. Employing a dynamic investment strategy[4] over a myopic one reduces the probability of default. An endogenous learner is less likely to default than one who does not - he holds a larger information set (he acquires signals as well as observing returns) and receives more timely information (before allocating his portfolio). Also, a dynamic investor has positive hedging demands[5] meaning he will always buy more (sell fewer) units of an asset than an agent maximising a myopic portfolio problem, as paper 1 explains. So, for example, an agent may invest more today than is myopically optimal in order to balance the risk of lower returns tomorrow. Brandt (2010) writes further about hedging demands. This result confirms the prediction of paper 1 that dynamic investors will be less likely to default than myopic ones. paper 1 implies that, building on VN&V, the standard Markowitz portfolio is not an investor's optimal choice in the presence of a learning technology in a dynamic model. In VN&V and paper 1 learning makes the agent's portfolio choice appear riskier than the Markowitz portfolio (he buys more risky asset units) as the agent is more certain about the assets he learns about whereas the dynamic horizon makes a dynamic investment strategy superior to a myopic one by a hedging demand term.

Thirdly, when investing dynamically, the agent may settle on a specialised learning strategy. In a static model, VN&V's myopic investor maintains a constant information acquisition choice, which is to generalise (spread his learning evenly). Their investor is risk averse and indifferent to the time of uncertainty resolution. Generalisation means he wants the same posterior precision about each asset and allocates signals to achieve this. Consequently, he invests more in assets that he initially knew less about and so spreads his wealth across more assets than one who does not learn.

So, the third result offers evidence that information acquisition is different in a dynamic model to a static one. The simulations sometimes see the agent specialise and change his learning whereas VN&V's agent always generalises, with the same preference combinations. In this paper the agent's learning strategy depends on his informa-

[4] A dynamic investor considers tomorrow's as well as today's investment opportunities when allocating his portfolio.

[5] An attempt to spread risk over time.

tion set. paper 1 builds on VN&V's result by making the theoretical prediction that the learning strategy of the agent who invests dynamically is time-dependent and will vary less the more accurate the initial beliefs. However, this turns out to be incomplete because the learning strategy depends on the agent's information set. That is, the choice of learning strategy is time-dependent only before the investor is convinced that his beliefs have converged[6] to the truth. After that, his choice of learning strategy is permanent due to his Kalman filtering.

Widespread heterogeneous beliefs are also studied by Coibion et al (2015) who find that information acquisition is state dependent, 60% of New Zealand firms are less willing to buy new information when they perceive the economy to be in a good state whereas 75% are more willing to when they perceive the economy to be in a bad state. Their results indicate that agents prefer to acquire information when the economy contracts, anti-cyclical information demand. They also find that over half of firms (51%) surveyed are uninformed[7] about recent past (the last 12 months) inflation and that their opinions about past and future inflation exhibit wide dispersion despite the fact that inflation has been stable for the past 25 years.

This paper can also contribute to explaining another stylised fact. As Curcuru et al (2010) document, many US households, 13.7%, hold undiversified portfolios. Both VN&V and paper 1 attempt to explain this peculiar fact. The power of VN&V's result is that it can explain, in a static context, how a increasing returns to learning lead an agent to hold a diversified portfolio. Paper 1 finds something similar in a dynamic context. Simulations in this paper can replicate concentrated portfolio choice when the investor acquires signals that support his prior beliefs and does not observe returns.

So the dynamic model built in this paper is capable of testing VN&V's generalisation result by extending the investment horizon to multiple periods. Thus a major contribution is providing evidence that their static result changes in a dynamic horizon. The simulation results indicate that the investor who learns exclusively endogenously

[6] Convergence in learning means that he no longer weights new information when updating his beliefs because the Kalman gain reaches zero.

[7] They define informed as having beliefs about past inflation within 2% of the true value.

(does not observe actual returns), under certain initial beliefs (necessary conditions), specialises. These entail that the agent is initially optimistic about one asset and pessimistic about the other one, with low precision about both. This requires the fairly strong assumption that the investor does not observe actual risky asset returns. Observing returns stops the investor reaching a degenerate distribution of beliefs and he does not specialise permanently, instead, he sometimes uses a generalised learning allocation and sometimes a specialised one before settling on a generalised strategy when his beliefs converge.

Although the assumption that an agent does not observe reality may seem strong Coibion et al (2015) find that, through surveying businesses operating in various industries in New Zealand, there is a significant number of firms that are uninformed about past economic conditions (49%), as previously mentioned. They also find that there is wide dispersion in beliefs about past inflation rates. Hence many decision makers have imperfect information about previous realisations, so the assumption that an agent does not observe actual returns is not as strong as it may first appear. They attribute the fact that so many firms are uninformed to rational inattention - that agents have limited capacity to process signals and allocate that capacity optimally meaning agents do not include some observations in their information sets.

3.1.1 Stock market returns & investor beliefs

This subsection provides data on recent & historical stock market returns as well as investor perspectives on these statistics. The 2014 wave of the PAT€R survey of French households, representative of the population by age & wealth, provides household perspectives on the stock market and recent & historical returns are taken from the French CAC 40 index. Statistics quoted are based on the opening values of the market between 1 December 2011 and 1 December 2014 for recent returns (to coincide with the relevant survey wave) and between 1 March 1990 and 1 March 2016 for historical returns.

Table (A.5) describes the recent (3 years) and historical (past 26 years) of CAC40 returns. Plainly there is significant risk as shown by the standard deviation of the index's value but returns were sizeable for both time horizons. So for those willing to bear risk, significant returns were obtainable.

Evidence on how informed households are about the positive returns that were available in the recent past is below. Households responded to a question[8] asking them to allocate statistical weights to a range of discrete possible outcomes for recent stock market returns in the set [more than -25%, -10-25%, less than -10%, 0%, less than +10%, +10-25%, more than +25%]. As the two extreme categories are open ended, the midpoints used are +37.5% and -37.5% respectively. A distribution for each household's belief about the stock market can then be plotted and the "Perceived Returns" data (Appendix section A.2.1) shows the mean past belief of each household. "Expected Returns" are data from a similar question about households' expectations of returns over the next 5 years under the same categories. Perceived and expected returns data are in tables A.6 & A.7.

These statistics show that households are significantly uninformed about recent returns. They also seem uninformed about historical returns, given that these are higher per year than recent returns, it would be understandable for some households to have higher perceptions than the actual figure should their recent information updates have been arbitrarily noisy. Uninformed beliefs are surprising given that financial returns are publicly available. This project interprets this fact by appealing to the theory of endogenous information acquisition within the field of bounded rationality. As such, a household has limited capacity with which to learn about the stock market and must allocate that capacity between different stocks.

The whole sample of households perceived a return of just 3.6% in the recent past, meaning 1.2% per year and less than a thirtieth of the actual figure (38.17%). Certain subgroups are better informed than others including high earners (incomes exceeding €40,000 per year), males and the highly educated. Of these, high earners know the most (7.10%) but even this is only about a fifth of the return actually realised. Married

[8] Copies of the two questions can be found in the appendix section A.2.2

individuals actually are worse informed about recent returns than the average (3.29% perceived return).

Expected returns for the next 5 years are roughly half perceived returns. The annual expected return is even less than half the perceived as expectations cover 5 rather than 3 years. So households expect a decrease in returns even as recent returns are lower than historical returns as presented in this study. Nevertheless, the overall sample expects a yearly return of just 0.32% over the next 5 years and the most optimistic subgroup, high earners, only expects a yearly return of 0.69%.

3.1.2 Heterogeneity in investor beliefs

Households have beliefs about the stock market that are significantly heterogeneous. Graphs A.2, A.3, A.4 & A.5 (Appendix section A.2.1) are histograms of household perceptions of recent returns (red line indicates the actual return) and future expectations by subsample.

Perceptions are approximately normally distributed (A.2) meaning that there is a high standard deviation describing how informed the population is. Responses are slightly more skewed towards positive returns and somewhat more concentrated for better informed subgroups (males and high earners) (A.3) but beliefs are consistently and strongly heterogeneous. This is somewhat surprising given that, again, returns are publicly available information and hence for perceptions to be more accurate and more similar would be natural. Though accurate and homogenous beliefs sound intuitively right, largely inaccurate and heterogeneous beliefs persist.

Histograms of perceived returns also show how uninformed the general population is. Less than 5% of the sample correctly identified this return (over 37.5% which was the highest possible response of survey participants) with a small population density next to the red line in each graph. Nevertheless, a few individuals are well-informed meaning that figures A.2 & A.3 have a bump at the extreme right end, disturbing the approximately normal distribution. More extreme is the fact that only just over half of participants correctly identified that the stock market had grown in the 3 year period and some individuals thought it had lost over 37.5% of its value, the small left-hand bump. The histogram for high earners is more to the right of other subsamples (figure A.3) more individuals correctly identifying both a positive return and the true return as well as fewer incorrectly identifying a extreme losses. Having a high income could be simultaneously linked to being well informed as the better informed can earn more and those who earn more can access more information through higher participation.

Expected returns are strikingly pessimistic; nearly 50% of all participants forecast a return of just over zero and actually 31% expect no growth whereas 25% expect a stock market fall (figure A.4). However, expectations are significantly more similar than perceptions. They are also quite symmetric in distribution, around zero, indicating that the average household believes that the stock market will not change. Although expected returns are more strongly grouped, they are more extreme than perceptions, the maximum and minimum expected returns are absolutely greater, though this is an effect of the longer horizon posed to participants. Heterogeneity is again very prevalent across all subgroups. The male, highly educated and top earning subgroups (figure A.5) all have less symmetric distributions that are shifted to the right (negative skew). This is particularly pronounced in the case of top earners of whom 60% expect there to be growth in the stock market. However, married folk tend to have similar expectations to the average member of the population, with a relatively symmetric distribution, and a minority (44%) expecting growth.

The rest of this paper is organised as follows. Chapter 2 describes the dynamic model in which the investor learns and allocates his portfolio. Chapter 3 firstly describes the simulations that are carried out and then details the three main results found, as mentioned earlier. Lastly, Chapter 4 concludes and includes a few suggestions for further work. Some figures from the simulations and some derivations for the model can be found in the appendix following.

3.2 The model

The chapter presents the model, which has two versions based on different learning technologies. Both model the investment decision of a single investor who chooses between one risk less and two risky assets. The constant return on the risk less asset, r, is known. Both risky assets follow normal distributions with the same mean θ and variance Σ_ϵ and realisations per asset are determined by $f^i_{t+1} = \theta + \epsilon^i_{t+1}$, where $\epsilon_{t+1} \sim N(0, \Sigma_\epsilon)$ and are randomly drawn IID (Independently and Identically Distributed) shocks. So they are identical. However, the investor does not know this, he is endowed with a prior belief about the mean and the variance of each asset, $(\mu_{i,0}, \Sigma_{i,0})$, where $i = 1, 2$ represents the asset considered and 0 indicates the initial period.

The agent also is endowed with initial wealth, w_0. He allocates it between the three assets to maximise his CARA expected utility over next period wealth

$$- E_t[exp(-\rho W_{t+1})]. \tag{3.2.1}$$

W_{t+1} is wealth generated tomorrow and ρ is the risk aversion coefficient. The process for wealth (and budget constraint) is given by

$$W_{t+1} = rW_t + q'_t(f_{t+1} - p_t r). \tag{3.2.2}$$

This process was developed by VN&V from Admati (1985). Vectors q_t and p_t describe the quantity of units and price of the risky assets. f_{t+1} is the returns vector on all risky assets. Wealth tomorrow can be broken down into two components. The first is the return on investing wealth inherited from yesterday in the risk free asset. The second is the excess return earned on all risky assets, $(f_{t+1} - p_t r)$. If all wealth is allocated to risky assets then, multiplying out the brackets, the first and third terms cancel giving wealth generated as the quantity of each risky asset by its return. The investor's belief about the variance of each risky asset is

$$\Sigma_t = \sigma_t + \Sigma_\epsilon. \tag{3.2.3}$$

Σ is the diagonal variance-covariance matrix - the risky assets are independent and the diagonal elements are the two true variances. σ_t is a diagonal matrix describing the investor's uncertainty about his belief about each risky asset's true mean ($\sigma_{ii,t} \in \mathbb{R}^+$). It decreases in absolute value when he learns. Σ_ϵ is the true variance-covariance matrix. Σ_t is of course bounded below at zero. The initial prior variance is given by (3.2.3) when $t = 0$ and σ_0 is randomly assigned. The initial beliefs about the risky assets means, $\mu_{i,0}$, are also randomly assigned ($\mu_{i,t} \in \mathbb{R}$). The mean and variance posterior beliefs are found by updating the priors using the Kalman Filter formulas below, derived in the appendix, subsection A.1.4.

$$\mu_{t+1} = \mu_t + K_t(f_{t+1} - \mu_t) \tag{3.2.4}$$

updates his belief about the mean like adaptive expectations and

$$\sigma_{t+1} = \sigma_t - \sigma_t(\sigma_t + \Sigma_\epsilon)^{-1}\sigma_t \tag{3.2.5}$$

updates his belief about the variance. K_t, a matrix, is the Kalman gain, a term that describes the weighting between the agent's prior belief and what he learns. This weight decreases over time so the agent weights his prior more as learning increases (more updates). There is nothing random in (3.2.5), it is completely deterministic and the agent knows the path his variance beliefs will take.

This version of the model is when the agent learns exogenously, he updates his beliefs only with observed returns. The next subsection describes the other version when the investor also endogenously acquires signals. The investor's problem is given by

$$max_{q_t}||E_t[\sum_{t=1}^{T} -exp(-\rho(W_{t+1}))|\mu_t, \Sigma_t] \tag{3.2.6}$$

subject to (3.2.2), (3.2.3), (3.2.4) and (3.2.5). The agent is forced to treat the problem in a myopic fashion in the current version of the model, each time period is additively separable. This rules out the possibility of dynamic portfolio choice through hedging demands[9] as in Merton (1973). There are three sufficient conditions under which it is optimal to invest myopically rather than dynamically. They are below and see Brandt (2010) for more detail.

1. When preferences are logarithmic form, then the problem can be separated into a sum of individual utility maximisation problems.

2. When investment opportunities are constant over time, for example when returns are independently and identically distributed, as state variables are uninformative about future returns.

3. When investment opportunities are stochastic but unhedgable.

The model violates the first and third conditions and negates the second condition as the investor's beliefs about the mean and variance today are sufficient statistics for the beliefs yesterday. So a dynamic strategy is superior to a myopic one. However, in order to contrast the effects of these two strategies both are employed in the simulation analysis, with myopic investment first. Paper 1 predicts that the dynamic investor will adapt his portfolio choice faster to new information received and so be less likely to default than the myopic investor. VN&V find that an agent who is indifferent to the time that uncertainty is resolved and who has CARA preferences over wealth chooses a generalised learning strategy. Using both investment strategies allows this prediction and this result to be tested.

Substituting (3.2.2) into (3.2.6) and taking the expectation means that the problem becomes

$$max_{q_t}|| \sum_{t=1}^{T} -exp(-\rho(rW_t + q_t'(\mu_t - p_t r) + \frac{1}{2}\rho^2 q_t'\Sigma_t q_t)) \tag{3.2.7}$$

[9] These are the desire to spread risk across time by bearing more or less risk today in order to compensate for lower or higher expected risk tomorrow. They arise when the investor finds it optimal to spread risk not only across different assets to address idiosyncratic risk but also across time to hedge against fluctuations in investment opportunities.

subject to (3.2.3), (3.2.4) and (3.2.5). Both the mean (μ) and variance (Σ) beliefs are priors in the sense that they will be updated tomorrow once actual risky returns are observed and the result, today's posterior, will be tomorrow's prior.

The sequence of actions that the investor takes is as follows.

1. Maximise the current period problem to find the optimal vector of risky assets, q_t.
2. Deposit any remaining wealth into the risk less asset to complete the period t portfolio.
3. Observe the risky asset realisation, f_{t+1}, next period and discover the wealth earned, W_{t+1}, the budget constraint in $t+1$.
4. Update beliefs, exogenously, about the mean and variance, μ_t and Σ_t.

3.2.1 Endogenous learning

This is the second version of the model and the investor can acquire information endogenously. His time preference is indifferent. The signals are costly to acquire as there is a constraint, L, on the number drawn each period. It is not a financial cost but can be thought of as a constraint on the time the investor gathers financial information each day. So the endogenous learner faces an extra constraint to his maximisation problem.

$$\sum_i d_{i,t} \leqslant L \tag{3.2.8}$$

L limits the number of times that the agent can update his beliefs about a particular asset within a period, before he observes the actual return. $d_{i,t}$ is the number of signals drawn about asset i in time t. So the endogenous learner can influence his own information set. Each signal drawn:

$$\eta_{i,t} = \theta + e_{i,t},$$

is true in expectation because $e_{i,t} \sim N(0, \Sigma_\eta)$. Hence signals provide useful information to the investor about the underlying mean, θ.

Endogenous learning provides two more advantages to the investor. Firstly, he can learn quicker than the exogenous learner as, in a given period, he draws up to L informative signals as well as observing the return. So his information set can grow quicker. Secondly, he learns prior to his portfolio allocation, and conditions upon signals drawn unlike the exogenous learner. Drawing (accurate) signals reduces the probability of conditioning portfolio choice on inaccurate priors, though this is not guaranteed (signals may be extreme with low probability). The equation for updating mean beliefs with signals is below.

$$\mu_{t+1} = \mu_t + K_t(\eta_t - \mu_t) \tag{3.2.9}$$

The agent combines the difference between his prior and the signal received with the prior. The endogenous learner uses both equations 3.2.4 and 3.2.9 to update his mean beliefs each period. Again, the investor's variance belief is guaranteed to become more accurate with each signal drawn as (3.2.5) is deterministic. Uncertainty (3.2.5) tends to zero and (3.2.3) tends to $\Sigma_{i,\eta}$ regardless of each signal's value.

The endogenous learner's problem is to maximise (3.2.7) subject to (3.2.3), (3.2.4), (3.2.5), (3.2.8) and (3.2.9). The problem is the same as the exogenous learner's except for the addition of the last two constraints.

The sequence of actions that the investor takes is similar to that taken by the one who learns exogenously and is as follows.

1. Allocate the learning capacity, L, between the two risky assets.

2. Observe the signals and update mean and variance beliefs, μ_t and Σ_t.

3. Maximise the current period problem, conditional on updated beliefs, to find the optimal portfolio choice, q_t.

4. Deposit any remaining wealth into the risk less asset to complete the period t portfolio.

5. Observe the risky asset realisation, f_{t+1}, next period and discover the wealth earned, W_{t+1}, the budget constraint in $t+1$.

6. Update beliefs exogenously.

3.2.2 Dynamic investment strategies

So far exogenous and endogenous learners have been modelled who solely use a myopic approach to investing. This subsection describes how the investor can instead employ a dynamic investment strategy combined with either of the technologies previously explained.

Modelling both myopic & dynamic investment strategies means the simulated results can be compared. Then paper 1's prediction (lower probability of default for dynamic than myopic investors) and VN&V's result (a myopic investor with CARA wealth and indifferent time preferences chooses a generalised learning strategy) can be investigated.

The dynamic investor considers future investment opportunities when he invests today whereas the myopic investor simply maximises over what is available to him today. The dynamic investor could spread risk over time; he may expect worse (better) investment prospects tomorrow compared to those available today. Thus, he could invest more (less) aggressively in risky assets today than the myopic investor would to hedge lower expected returns tomorrow by greater returns today. Therefore, dynamic and myopic portfolios can differ and the difference is called a hedging demand, the extra demand for risky assets to spread risk over time. Merton (1973) first used the term. Brandt (2003) and (2010) provide further explanations. A necessary condition is contemporary correlation between investment opportunities and the risky assets returns to inform the investor about how investment opportunities may change in the future. With such knowledge he can effectively choose optimal hedging demands. The returns in the present model are IID and hence violate this condition. However, paper 1 indicates that the endogenous and exogenous learning technologies in the model render investor beliefs as sufficient statistics for past returns so hedging demands can arise. Hence the Markov property is satisfied and the agent's problem can be written recursively.

First, is the Bellman equation for the dynamic investor who learns exogenously.

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t\}} || E_t^- \{-exp(-\rho W_{t+1}) | \mu_t^-, \Sigma_t^-\} \\ + \beta E_t^- \{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^-\}. \end{aligned} \tag{3.2.10}$$

The value function represents that the agent considers future investment opportunities when he allocates his portfolio today. The negative superscripts indicate that expectations and state variables are for exogenous learning - expectations are conditional only on observed risky assets returns and not endogenously acquired signals. The exogenously learning dynamic investor maximises (3.2.10) subject to (3.2.2), (3.2.3), (3.2.4), (3.2.5). This requires substituting (3.2.2) into (3.2.10) and then taking expectations over the unknown future returns, conditioning on prior beliefs.

Second is the Bellman equation for the dynamic investor who learns endogenously.

$$\begin{aligned} V(W_t, \mu_t^+, \Sigma_t^+) = max_{\{q_t\}} || E_t^+ \{-exp(-\rho W_{t+1}) | \mu_t^+, \Sigma_t^+\} \\ + \beta E_t^+ \{V(W_{t+1}, \mu_{t+1}^+, \Sigma_{t+1}^+) | \mu_t^+, \Sigma_t^+\}. \end{aligned} \tag{3.2.11}$$

Again, the value function represents that the agent considers future investment opportunities today. The positive superscripts indicate that expectations and state variables are for endogenous learning - expectations are conditional on both observed returns and endogenously acquired signals. The endogenously learning dynamic investor maximises (3.2.11) subject to (3.2.2), (3.2.3), (3.2.4), (3.2.5), (3.2.8) and (3.2.9). This requires substituting (3.2.2) into (3.2.11) and then taking expectations over the unknown future returns, conditioning on posterior beliefs.

The next chapter describes the simulations of the two versions of the model and then states and explains some interesting results found.

3.3 Simulations

The two learning technologies and the investment strategies (section 3.2) are combined and simulated. The investor's behaviour is observed and differs when the learning technology and investment strategy combination changes and the connection between his beliefs, learning choices and portfolio allocation is revealed. Due to computational constraints the dynamic investment strategies are approximated, the investor accounts for one additional period ahead instead of every future period.[10] Interesting results emerge about the relationship between demand for information and the fluctuations of returns, how inaccurate beliefs can lead the investor to default and how he can specialise his learning choice (described in subsection 3.3.1).

Five versions of the code are run.

1. ***Myopic exogenous learning***
   The agent only observes returns (no information acquisition before portfolio choice) and invests myopically.
2. ***Dynamic exogenous learning***
   The agent only observes returns and invests dynamically.
3. ***Static endogenous learning***
   The agent acquires signals, observes no returns and invests myopically.
4. ***Dynamic endogenous learning***
   The agent acquires signals, observes returns and invests dynamically.
5. ***Dynamic endogenous learning: no observations***

[10] Simulating the dynamic model fully would be very computationally very costly. It would require simulating forwards for all remaining periods of the investment horizon in each current period t in order to find the complete future chain of the state variables because some of them are stochastic.

The agent acquires signals only and invests dynamically.

The values of the parameters and the initial values of variables are in the two tables below. The initial beliefs are altered to give four states of the world. The first is fully described below, the second is the same but with more accurate and optimistic initial mean beliefs ($\mu_1 = 4.8$ & $\mu_1 = 5.2$), in the third the agent has the same optimism as the first but more confidence ($\sigma_1 = 0.8$ & $\sigma_2 = 0.5$) and in the fourth the agent is as confident as in the third ($\sigma_1 = 0.8$ & $\sigma_2 = 0.5$) but much more pessimistic ($\mu_1 = -2$ & $\mu_1 = -0.5$). See the true mean, $\theta = 5$, the investor is generally pessimistic.

Initial variables		
W	1	wealth
μ_1	1	belief: mean
μ_2	5	belief: mean
σ_1	2	belief: variance
σ_2	3	belief: variance

The investor's initial wealth is low so he cannot purchase many assets. Initially he has high (and different) uncertainty about both risky assets, making him unconfident. Also, he is optimistic about asset 2 because he believes its excess return is positive ($\mu_2 = 5$ exceeds $pr = 2$) but he is pessimistic about asset 1 for the inverse reason ($\mu_1 = 1$).

Parameters		
r	1	risk free rate
β	0.99	discount factor
ρ	1	risk aversion
p_1	2	price
p_2	2	price
θ	5	true: mean
Σ_f	2	true: variance
Σ_η	2	signal variance
L	2	draw limit

The risk free return ($r = 1$) exactly maintains all wealth invested in it. The true mean of both assets is the same ($\theta = 5$), they are essentially identical, so they payoff a positive amount on average. So in state 1, the agent is likely to (rightly) buy asset 2 and (wrongly) sell asset 1. He can draw up to 2 signals each period, specialising his learning by drawing both signals about one asset or generalising[11] by drawing a signal about each asset. Note that the true variance is the same as the variance of the signals.

The investment horizon is set at 75 periods. Intuitive graphs illustrating how the main variables operate over a typical (not average) set of 75 periods are in appendix section A.2.4. See figures A.6-A.15. The first two are for the myopic agent who learns exogenously whereas the second two are for the dynamic agent who learns exogenously. Wealth typically increases over time and is higher under the dynamic investment strategy. Beliefs steadily converge over time under both investment strategies and influence the quantities of risky assets the agent purchases. The next six graphs are all for the endogenous learner: Two each for the myopic agent, dynamic agent and dynamic agent who does not observe returns. The myopic agent always chooses the same learning strategy (generalisation) no matter the returns or his beliefs. He invests similarly to the myopic agent who learns exogenously. So there is little intuitive difference between the myopic agents who learn endogenously and exogenously. This is plain given that they both receive 1 signal per period, only the endogenous learner receives it before his portfolio allocation. The dynamic agent who does not observe returns, also has a constant learning choice, he always specialises in one asset. This shows that various degenerate steady state beliefs are possible, determined by initial prior beliefs and can provide a mechanism to explain widespread heterogeneity in beliefs. His beliefs only converge about that asset because he learns nothing ever about the other, no matter the investment horizon length. He would only switch if an acquired signal is strongly surprising (far from his prior belief). When he does observe returns, the dynamic agent switches learning strategies. Eventually he suspects his prior beliefs were wrong and that the two assets are the same. Then he consistently chooses a generalised learning strategy.

[11] In appendix results section A.2.6, a value of 3 indicates specialisation, 1 indicates no learning and 2 indicates generalisation in that asset.

His choice of risky assets also becomes consistent and this is when he becomes very confident in his beliefs (low variance).

Thus when the agent does not observe public information and he draws signals that support his priors, the result arises. With these two assumptions, the model can replicate heterogeneous beliefs.

3.3.1 Results

This paper finds three interesting results. The first comes from testing the prediction of paper 1 that the agent who uses a dynamic investment strategy will be less likely to default than the one who invests myopically. Default happens after the investor loses more wealth than he could raise by short selling the maximum possible quantity in the risky asset market. Unsurprisingly, the dynamic agent defaults less frequently, though both exogenous and endogenous learners can default. Brandt (2010) provides numerical examples that show that the utility of an agent who invests dynamically is higher than that of one who invests myopically. He finds that the increased utility is directly related to the presence of hedging demands. Hedging demands also generate the present result and their existence depends on the assumptions of a learning technology and multi-period framework under IID returns, as shown bypaper 1. With no learning technology hedging demands come from correlation between the risky asset returns and the investment opportunity set as Brandt explains.

The second result is that the dynamic investor's learning allocation varies whilst VN&V's myopic investor always generalises (under indifferent time preferences). The dynamic investor switches between specialisation (draw all signals from one asset) and generalisation (spread the signals evenly between the two), though specialising in one asset for the whole investment horizon only occurs under the necessary condition that returns are not observed. So the static model result of generalised learning does not apply to a dynamic model. The third result is that the agent's demand for information is unresponsive to fluctuations in returns.

When allocating his signal draws, the investor considers his beliefs about the two risky asset distributions. He dislikes imprecision. Strangely, he can specialise in an asset with more precision in the first period. For example, if $\mu_0^1 = \mu_0^2 < \theta$ (0 for initial values) and $0 < \sigma_0^1 << \sigma_0^2$ then asset 1 is superior but asset 2 could be a useful hedge. So the investor initially expects to hold some of each. Then learning about asset 2 yields a more marginal utility per signal. Then, as the investor undervalues the assets, the signals drawn likely increase his optimism about asset 2 and he will buy more of it than asset 1 in period 1. He likes assets that offer a high return and this can be through buying or selling so the absolute value of a return is important and not its sign. A very negative excess return is attractive to the investor because he can sell a large quantity to generate wealth. Therefore, the investors draws more signals about attractive (higher expected excess return and lower uncertainty) assets, rather than attempting to experiment or equalise posterior variances.

The investor can specialise in one asset, specialise in the other or generalise over his signal allocation choice. The number of draws is set at two in the simulations. Indeed, depending on the relationship between the investor's beliefs about the distributions of the two assets, he can and does take any of these three learning positions, until his beliefs converge close to the true values. Only under non-observation of returns could the investor specialise for the whole duration and this requires that the signals he receives confirm his prior beliefs. Otherwise, signals that contradict his priors could lead him to specialise in the other asset or generalise.

Result 1. *The myopic investor needs deeper pockets.*
When the investor has CARA preferences, inaccurate initial beliefs, is indifferent about the time of uncertainty resolution and investment opportunities are both stochastic & hedgable, then the agent defaults with positive probability and more under a myopic strategy.

Table A.8 in the appendix shows the probability of default for each investor type in each state. Table A.9 shows the duration of the agent's participation in the market. Paper 1 predicts that a dynamic investor will not need as deep pockets as a myopic one. This means that an agent who enters the market with inaccurate initial beliefs is more likely to survive if he considers future investment opportunities today. The simulation results support this prediction. Using a dynamic investment strategy, the exogenous learner has at least as good a chance of surviving in the market in all four states as the exogenous learner who invests myopically. In addition, he survives for more periods in states 1 & 3 and for the same number in 2. The difference in 4 is negligible and because the myopic agent invests more aggressively, the dynamic investor is more cautious (quantities bought and sold are absolutely smaller). Being more aggressive disadvantages him in states 1 & 3 where his portfolio choices amplify his inaccurate belief but in state 4 when beliefs of all agents are extremely inaccurate, aggressively shorting one asset and buying the other is a superior strategy to cautiously selling both as the dynamic agent does.

The importance of this result is to substantiate that deviations from the standard diversified Markowitz portfolio in a multi-period investment model benefit the agent. Moreover, given that the investor is uninformed, he is more likely to survive when trying to spread the risk he faces over time. This intuitively makes sense. When the agent takes a position based on uncertain beliefs, there is a strong possibility that they will lead him to make a substantial error. With positive probability this substantial mistake will lose him much wealth and cause him to default. The default limit is set beyond the point where selling as many assets as possible cannot cover his debt (negative wealth holding). Default is much more likely early in the investment horizon because the agent because reasonably informed after investing for several periods as his information set grows. Thus, hedging his position over two periods reduces the risk of default from inaccurate initial beliefs. As returns in this model are not correlated with investment opportunities, this type of hedging demand is somewhat particular to this paper and a variation on Brandt's (2010) perspective. In his case, hedging demands are for smoothing the possibility that investment opportunities will fluctuate between today and tomorrow based on information from correlation data. However, in this study, the agent spreads risk over time caused by the degree to which he is uninformed (he infers something about returns tomorrow using his beliefs today). So this is a slightly different hedging demand.

Moreover the endogenous learner defaults with a lower or the same probability as the exogenous learner and survives for at least as many periods (an investment duration at least as long). See this by comparing rows 1 & 3 and 2 & 5 in tables A.8 and A.9. The endogenous learner fares better because he is better informed and updates his information set prior allocating his portfolio. Under dynamic investing, the endogenous learner has a larger information set. However, information sets are the same size for myopic exogenous and endogenous learners. Information acquisition benefits the myopic investor not by providing more information but more timely information, it arrives before he allocates his portfolio. Effectively, the timing of the learning reduces the risk of conditioning portfolio choice on inaccurate beliefs.

The second result relates to one found by VN&V. They discovered that an agent with CARA preferences over wealth and indifferent to time preferences always chooses a generalised learning allocation. That is, there is an interior solution to their static problem of investing and learning.

Result 2. *The multi-period learning allocation varies.*
When the investor has CARA preferences, inaccurate initial beliefs, is indifferent about the time of uncertainty resolution and investment opportunities are both stochastic & hedgable, then the agent may vary his learning choice in a multi-period investment horizon.

The dynamic investor frequently deviates from the generalisation result found by VN&V. The simulations that support this result are summarised in four figures: A.17, A.18, A.19 and A.20 in the appendix. The first two relate to the endogenous dynamic version of the model and the second two, to the endogenous dynamic version without the observation of actual returns. The first figure in each pair displays the learning allocation to asset 1 in each of the four states (left to right) and the second displays the allocations to asset 2. The dotted red line on each represents the generalised choice of VN&V's myopic agent, at a constant value of 2. The coloured dots distributed between 1 & 3 represent the agent's learning choices under each investment strategy, where each represents a different simulation trial. 3 indicates a specialised learning choice in that asset and 1 means specialisation in the other. Each visible coloured dot is a learning choice by the investor that deviates from the prediction of VN&V's static model and hence an indication that myopic and dynamic choices differ. Perhaps if the investment horizon were to be made much longer then the beliefs would converge in all four states and the learning allocations would no longer deviate from the red dotted line. This could be for future work to investigate.

In the endogenous model, there are deviations in all trials throughout the investment period in states 2-4 (figures A.17 & A.18), coloured dots are visible for the whole investment period. In state 1, the agent's mean beliefs converge about period 21 - he realises the assets are identical and employs generalisation - and there is only one deviation after period 21. Yet, even in state 1 deviations persist for about one third of the investment period. So when investing dynamically the endogenous learner persistently allocates signals differently to VN&V's static agent.

Interestingly, when the endogenous dynamic investor cannot observe signals there are two possible equilibria (figures A.19 & A.20). First, like the endogenous learning version, mean beliefs can converge and the persistently generalises. This is primarily seen in state two, when initial beliefs are close for the two assets. The agent switches his learning allocations for about 30 periods, then his beliefs converge and he generalises (except in one trial, in dark blue, where variation continues). However, in states 1 & 3, he learns about one asset forever. This is the second equilibrium. In state 1 the agent learns about asset 2 (with a more attractive mean-variance combination) forever as his learning is self-reinforcing. As long as he receives signals that agree with his prior and observes nothing about asset 1 he retains his specialised strategy and asset 2 even become more attractive because his uncertainty about it deceases. This can lead to degenerate belief distributions that completely depend on the initial priors and potentially explain widespread heterogeneity of beliefs. The endogenous learner continues to observe returns and so is unlikely to fall into a degenerate distribution. The self-reinforcing equilibrium can be shifted to the generalising equilibrium by reintroducing observation of returns or by giving the agent an unexpected negative shock about his preferred asset. Observing an extreme, low signal will reduce his mean-variance combination about that asset and persuade him to learn about the other. Signals must then keep the mean-variance combination close for the two assets otherwise he will keep his specialised strategy. State 3 (figures A.19 & A.20) give examples where this is unsuccessful. The agent diversifies out of his specialised learning strategy briefly but only in one trial do the signals bring his mean beliefs close enough to persuade him to give up specialising.

The importance of this result is two-fold. Most important, it shows that the learning strategy of the endogenously learning dynamic investor differs from the myopic investor's strategy. Even as myopic and dynamic portfolio choices can differ under non-zero hedging demands, evidence is found that myopic and dynamic learning decisions differ too. Further research could investigate different preference combinations (wealth and time) to discover how widely this result applies. For the investor, this indicates that his learning choices depend on the length of time for which he has been trading assets and this is a proxy for his information set size and hence how confident he is about the accuracy of his beliefs. Regarding portfolio choices, the investor who has invested for longer may have a stabilised learning strategy which is associated with a stabilised portfolio allocation strategy in the simulations. In an uninvestigated state of the world in which the underlying value of one asset is superior to the other, this could cement a long term specialised investment strategy and eliminate the inferior asset in a general equilibrium model.

Of second importance is that a permanent uninformed equilibrium can occur (degenerate distributions can be sustained). When the agent learns endogenously without observing returns and initially strongly favours one asset over the other, he can both persistently and permanently specialise his learning in his preferred asset. With no learning about the disfavoured asset and expected signals (as in confirmation bias) a self-reinforcing partial equilibrium in which he remains in the dark about one asset is possible. So with just these two assumptions, this mechanism can explain widespread heterogeneity in agent beliefs.

The third result is that the agent's demand for information is unresponsive to changes in risky asset returns. Given Coibion et al's survey finding that decision markers in New Zealand tend to exhibit anti-cyclical (negative correlation with the business cycle) preferences for information (about inflation), this is a little surprising.

Result 3. *Learning is unresponsive to fluctuations in return.*
When the investor has CARA preferences, inaccurate initial beliefs, is indifferent about the time of uncertainty resolution and investment opportunities are both stochastic & hedgable, then his demand for information does not depend on the fluctuations of returns.

Looking at figure A.16 shows, that the information choice that the myopic investor makes is constant and so his demand for information is completely insensitive to the fluctuations of returns. Figures A.17, A.18, A.19 and A.20 show that the dynamic investor does vary his information choice when investing dynamically as Result 2 states. However, these fluctuations do not primarily depend on the realised returns of the two assets but his mean and the variance beliefs and his mean-variance preferences. There is a hint of a positive response in demand for information early in the investment horizon because signals have a strong impact on beliefs when the information set is small. Yet as the information set grows the prior belief becomes more important than new information because the Kalman gain coefficient decreases and this effect is non-existent in cases where the investor does not observe returns. So there is little evidence from the simulations of any effect of return fluctuations on the agent's learning choices. This results comes from the construction of the learning mechanism (Kalman filter), which values new information less with information set size and the stationarity of the underlying distribution. The latter means the agent learns the distribution over time. A persistent or moving average process would stop him learning the truth and such a model could better explain the anti-cyclical (regarding the business cycle) demand for information that Coibion et al find. Future work could test this result further using adaptive learning and a stochastic distribution.

In sum, with just the assumption of no access to public information and no surprise signals, the results of the simulations show that the theory can potentially offer a powerful explanation for the heterogeneity of beliefs; when the agent does not observe returns and draws signals that confirm his prior beliefs he specialises permanently in one asset. This can further explain why an agent may hold a concentrated portfolio in the long run; the superior expected payoffs that specialisation generates through reduced expected risk leads him to hold only that asset. As the underlying process is stationary, the investor can learn it meaning his portfolio allocation follows underlying values rather than short term fluctuations. An implication is that an investor's demand for information is little affected by the fluctuations of returns, unlike Coibion et al's survey results. This is something that the model is less able to contribute to explaining.

3.4 Conclusion

The standard Markowitz portfolio expects an investor to diversify his portfolio amongst imperfectly correlated risky assets. However, both Van Nieuwerburgh & Veldkamp (2010) and paper 1 find that the ability to endogenously acquire information renders this portfolio suboptimal. The latter finds that extending the model's horizon to multiple periods also renders the Merton portfolio suboptimal. Additionally, widespread heterogeneous beliefs exist amongst agents as documented by the 2014 wave of the PAT€R survey and Coibion et al (2015). This paper tests the implications of these theoretical results through numerical simulations. Three interesting results emerge. It also proposes a theory to explain heterogeneity of beliefs that relies on two key assumptions; no observation of public information and no surprise signals.

Critically, the theory highlights the effect of the multi-period investment model that a dynamic investor may forgo learning about one asset to acquire signals exclusively about the other because he believes that a superior underlying process drives it. Specialisation in that asset increases his expected utility of future wealth because he expects to hold more of that asset and learning reduces his uncertainty about its payoffs. Observing no returns and signals that confirm his priors, the agent will maintain such specialisation and develop an information set different to another agent with opposing initial priors. This theory can explain widespread heterogeneity in beliefs despite public availability of information. The theory can also explain how an investor could rationally hold a concentrated portfolio: he expects to hold the asset favoured by his prior beliefs and hence chooses to specialise in it, increasing his expected utility through uncertainty reduction, but if acquired signals confirm his priors he holds the expected portfolio, concentrated solely in that asset.

Demand for information is found to be unresponsive to return fluctuations: the agent either settles on a learning strategy after acquiring enough information to accurately know the underlying process driving returns or his acquired signals support his initial beliefs enough to persuade him to use a specialised learning strategy. Hence demand for information does not rely on returns. Given that demand is anti-cyclical (regarding the business cycle) in the survey findings of Coibion et al (2015), this is surprising. An implication is that agents in reality do not treat returns as though they are generated by an underlying, learnable process. The stationarity of the underlying process in the model generates this result and hence future work should consider modelling returns as following a persistent process with a stochastic underlying distribution.

Not only do endogenous information acquisition and the dynamic investment horizon mean that the optimal portfolio deviates from the diversified one of standard theory, it also outperforms the Markowitz portfolio. That is, there is a reduced probability of defaulting in the stock market through a coarse information set. The flip side is an (uninformed) investor with inaccurate initial priors needs deep pockets to participate in the stock market. However, the investor's probability of default is reduced if he endogenously acquires information and invests dynamically rather than myopically. Endogenous information acquisition provides both more and more timely information whilst a dynamic investment strategy induces aggression through hedging demands, increasing the sensitivity of portfolio choice to new information.

VN&V find that an agent tilts his portfolio towards riskier assets than the standard theory of diversification suggests when he acquires signals endogenously. Their agent with CARA wealth and indifferent time preferences chooses a generalised learning strategy. Yet, the present simulation results indicate that this does not extend to a dynamic investment model and an agent may actually pick a specialised learning strategy permanently. The evidence suggests that static generalised learning strategies do not hold in a dynamic model.

4

Paper 3

4.1 Introduction

The 2008 financial crisis that has so far led to the "Great Recession" is also shaking the foundations of macroeconomics. At the heart of the debate, is the role of expectations in state-of-the-art macroeconomic models, and in particular, in their financial counterparts. The standard practice has been to adopt the rational expectations (RE) paradigm, whereby households hold a (common) statistically correct and unbiased view of the future. RE have then a crucial advantage: rather than attempting the difficult task of measuring expectations they can be inferred from realisations. In a stationary environment, crises are then expected only to the extent they have happened in the past but unexpectedly[1], the crisis happened. Since the RE paradigm has proven fragile, we

[1]Before the crisis, (RE) macroeconomists were trying to understand the causes of the "Great Moderation".

need to undertake the difficult task of (i) critically assessing it (Guesnerie (2001)), (ii) measuring individual expectations (Manski (2004)), and (iii) understanding how they are formed and updated. A theory of individual expectations is crucial for economic theory in general and for financial markets in particular.

So far, the importance of (heterogeneous) subjective expectations in financial markets has been ascertained from evidence gathered (i) in laboratory experiments (Hommes (2013)), (ii) from agent-based computational algorithms (Arthur (2006)) and from survey data (Pesaran & Weale (2006)) for (i) stock market investors (Vissing-Jorgensen (2004)), for (ii) a specific population subgroup which includes non-stockholders (by age, e.g. Dominitz & Manski (2007)) and for (iii) a representative sample of the population by age and wealth (with an embedded experimental design, e.g. Hurd et al (2011), Arrondel et al (2014)).

The main shortcomings from those studies are that (i) only a subset of expectations updating rules is evaluated (chosen by the scientist), (ii) heterogeneity in individual information sets is not allowed, which implicitly assumes that households discount public announcements regarding the stock market index (Dominitz & Manski (2011)) and (iii) surveys often lack a longitudinal dimension on both individuals' expectations and information sets, which is crucial to studying (learning and) stock market expectations formation in realistic settings. Veldkamp (2011) and Vissing-Jorgensen (2004) do not allow for strategic motives in forming subjective expectations and this is at odds with theoretical, experimental and empirical evidence from the behavioural finance literature [2]. To overcome some of these limitations, in the PAT€R waves we inquired of respondents about their **perceptions** regarding the evolution of the stock market index over the three years, P_{t-3}, immediately prior to the survey (December 2014). The question posed is as follows (translated wording):

C42. 'Over the last 3 years, do you think that the stock market... -For each category write

[2]Examples include the 'law of small numbers' (the tendency of small samples to reflect the properties of the parent population), 'biased self-attribution' (the tendency to attribute past gains to one's skill and past losses to random bad luck) or the 'disposition effect' (the tendency to delay selling investments on which losses have been incurred in the hope that they will recover their losses).

down the probability of occurrence assigning a value between 0 and 100. The sum of all your answers must be equal to 100-:

-... has increased by more than 25%

-... has increased by 10 to 25%

-... has increased by less than 10%

-... has remained the same

-... has decreased by less than 10%

-... has decreased by 10 to 25%

-... has decreased by more than 25%.

Question C42 asks household i about the subjective relative likelihood of occurrence, $p^i_{t,k}$, of each of the seven possible scenarios, $k = 1, ..., 7$. Each scenario represents a possible outcome range for the percentage change in the index between $t - 3$ and t, $1 + R_t(3) \equiv P_t/P_{t-3}$.[3] The outcome ranges for R_t are identical to those of question C39, which instead inquires of respondents about the percentage change in the index between t and $t + 5$. That is $1 + R_{t+1}(5) \equiv P_{t+5}/P_t$. Question C39 captures respondents' **subjective expectations** about the return on a buy-and-hold portfolio that tracks the stock market index over a five year horizon. Accordingly, households' subjective likelihoods are given by $p^i_{t,k} \equiv \Pr{}^i[R_t \in k] = \Pr{}^i[(P_t/P_{t-5}) - 1 \in k], \forall i$ and similarly for $p^i_{t+1,k}$. From those answers, we can build the corresponding subjective cumulative distribution functions (CDF) for both perceptions and expectations of stock market returns, $F^i_{t,k} \equiv \Pr{}^i[R_t \in \cup\{k\}] = \sum_k p^i_{t,k}, \forall i$, in order to fit household specific CDFs. Under the parametric assumption of normality (standard in the financial literature), one can obtain the mean (m^i_t) and the variance (r^i_t) of the subjective normal distribution by Least Squares:

$$(m^i_t, r^i_t) \in \arg\min_{(m,r)} \sum\nolimits_k [F^i_{t,k} - F(R^i_{t,k}; m, r)]^2, \forall i$$

[3]Since ranges $k = 1$ and $k = 7$ are unbounded, we set (R_{max}, R_{min}) to match historically observed values.

where $F(R^i_{t,k};\mu,\sigma) = \Phi[(R^i_{t,k}-\mu)/\sigma]$. In addition, we disentangle the effect on households' asset demands of unobserved heterogeneity (in risk aversion and information sets) from state dependency (inertia, from transaction costs) along the lines suggested by Miniaci & Weber (2002), which is possible since we have direct measures of most of the unobserved components. Key to this identification strategy is Arrondel and Masson (2013)'s (Arrondel & Masson (2011)) who, by exploiting data from the 2007, 2009 and 2011 waves, established that households' risk and time preferences measured by a comprehensive score were not affected by the stock market crash of 2008.

To empirically identify if **strategic elements in individual stock market return expectations can explain stock market fluctuations**, we embed within Brandt (2010) the general framework of analysis by Desgranges & Gauthier (2013). Therein agents in (in)complete information strategic settings forecast returns, f^i, minimising the M.S.E. of the prediction subject to the way they believe the financial market works, the perceived law of motion (PLM)[4]. This can be seen below.

$$f^i \in \arg\min_f \left\{ \int_{R_{t+1}} [r_{t+1} - f^i(r_t,\theta)]^2 \underbrace{p^i(r_{t+1}|\, r_t,\theta)}_{\text{Individual i's predictive distribution}} dr_{t+1} \,\middle|\, \underbrace{r_{t+1} = \phi \int_0^1 f^j dj + \eta}_{\text{Perceived law of motion (PLM)}} \right\}$$

The best response forecast of individual i is then the following.

$$\underbrace{f^i}_{\text{Individual i's forecast}} = \phi \underbrace{\int_0^1 f^j dj}_{\text{Cross-sectional average forecast}} + \eta, \forall i$$

This framework captures strategic uncertainty as a source of ambiguity (in line with Hansen & Sargent (2012)) since it puts an additional 'market' constraint (PLM) in the individual forecasting problem. This gives rise to uncertainty about 'others' forecasts' f^j (ϕ) in addition to uncertainty about the fundamentals (η). Both are included in θ.

[4]The particular form of the PLM is adopted here for expositional simplicity, and not because it corresponds to the market clearing equation of a specific model. More complicated examples can be found in LeBaron (2013) or in Hommes & Wagener (2008).

With survey answers to question C39 across waves, ${}_{w}m^{i}_{t+1}$, we could use spatial econometric techniques to estimate the above equation, and identify whether individuals perceive the stock market as (i) efficient $(H_0 : \phi = 0)$, (ii) an instance where strategic substitutes prevail $(H_0 : \phi < 0)$ as in Guesnerie (2001) or in Morris & Shin (2002), and hence where increased competition may exacerbate volatility through the emergence of multiple equilibria (Veldkamp (2011)), or as (iii) an instance where strategic complements prevail $(H_0 : \phi > 0)$ as in Morris & Shin (1998). We could then hypothesise a two-step econometric estimation procedure, as in Arrondel et al (2014), to quantify the extent to which individual investment decisions depend on the first-step estimated effect on one's own forecast of the average of others' forecasts.

Morris & Shin (2002) models strategic interactions with two components, one in which agents value proximity to the true state and one in which they regard proximity to the average population action either positively under a strategic complementarity preference or negatively under a strategic substitutes preference. Agents minimise an expected loss function:

$$EL(a_i, a, s) = E[(1-r)(a_i - s)^2 + r(a_i - a)^2],$$

where a_i is the agents action, a is the average population action, s is the state. The strategic motive is described by:

- $r > 0$ - strategic complements
- $r = 0$ - no strategic motive
- $r < 0$ - strategic substitutes.

Reference to the model of Morris & Shin (2002) is relevant because strategic actions within the stock market can be seen as zero-sum. Only proximity to the true state raises social welfare. Their main finding is that strategic motives can be socially costly when signals (private or public) lead to actions far from the true state. This happens to the extent that information is imprecise and so individual and average population actions correlate little with the true state ($cov(a, s)$).

However, it has proven notoriously difficult in surveys to inquire of respondents about 'second order beliefs', or beliefs about others' beliefs. A possible route out we pursue here is to invoke Hellwig & Veldkamp's main theorem, which establishes a one-to-one relationship between individual actions (q_t^i) and individual motives in acquiring information (m_t^i) in incomplete information games with strategic substitutes or complements. To capture **strategic motives in information acquisition,** as in Veldkamp (2011), we designed a set of questions in the PAT€R 2014 wave, inquiring of respondents about their perceptions regarding the behaviour of others and how informed they believe others are. These questions (in translated wording) are:

C35. 'In your opinion, out of 100, how many people in the French population invest in the stock market (directly or via mutual funds)? If you do not know then select "I don't know"

and also, about others' information:

C36. 'In your opinion, out of 100, how many people in the French population are informed regarding the evolution of the stock market? If you do not know then select "I don't know".

Denoting by q_t^{-i} and m_t^{-i} the mean individual responses to question C35 (individual perception about the probability that an average individual invests in the stock market, when extracted at random from the French population) and C36 (individual perception about the probability that an average individual is informed about the stock market, when extracted at random from the French population) respectively, we can attempt to answer the main question by estimating the following relationship.

$$q_t^i = \max(0, \alpha_0 + \mathbf{X}_i' \alpha_1 + \alpha_q \cdot q_t^{-i} + \alpha_m \cdot m_t^{-i} + u_t^i) : u_t^i \big| \mathbf{X}_i; q_t^{-i}, m_t^{-i} \sim N(0, \sigma^2)$$

where $\mathbf{X}_i$ contains individual characteristics relevant to explaining the share of financial wealth invested directly or indirectly in the stock market, like risk tolerance, endowments, constraints or expectations m_t^i. This econometric specification shares features with 'peer-effects regressions', where respondents' perceptions relate instead to the characteristics of respondents' peers. Here, we inquire of respondents about their perceptions of average behaviour and information within the general population. For example, since $\widehat{\alpha}_q$ is found to be positive and statistically significant (see row 1 under col. 2 in Figure A.26 in the appendix exploiting the 2014 PAT€R survey wave), there is empirical evidence in support of the conjecture that individuals behave as if strategic complements prevail in the stock market. They are more likely to invest a share of their financial wealth in the stock market the higher the perceived proportion of fellow Frenchmen that also invest. Crucially, the identified effect could be an information effect in the sense of Hellwig & Veldkamp (2009), since $\widehat{\alpha}_m$ is found negative and statistically significant (see row 2 under col. 3 in Figure A.26 in the appendix when included with perceived proportion of fellow Frenchmen that invest). Yet it may not be since it is also negative and statistically insignificant when included alone. Individuals are found less likely to invest a share of their financial wealth the larger the fraction of the population they perceive to be informed about the stock market.

When we turn towards the actual mechanism behind the identified effects (own expectations), and directly examine whether subjective expectations of mean stock market returns are determined by strategic considerations, i.e. by q_t^{-i} and m_t^{-i}, Figure A.28 in the appendix reports the results under columns 2 and 3. The sign of both effects is consistent with respondents perceiving the stock market as a large game where strategic complements (included singly) prevail, i.e. the higher the expected return the higher the percentage of the population they perceive to be investing (column 2) or informed about (column 3) the stock market.

Hence, individual expectations regarding others' beliefs (interaction component) could have been a factor in triggering the recent financial crisis, in line with recent theoretical, such as Veldkamp (2011), and experimental, such as Hommes (2013), developments, as well as contributing towards its spread to the real economy, as in Hall (2010). In that sense, strategic elements in the subjective expectations of unsophisticated investors seem to provide empirical support to the destabilising effects of public information announcements, precisely for the reason advanced by Morris & Shin (2002).

This paper further investigates what determines individual beliefs about returns in the stock market and whether or not proxy variables for household information could explain the statistically significant strategic considerations in Figures A.26 & A.28. In particular these include household level of trust in others and financial advisers, inertia, sources of financial information and level of financial literacy. Subsequent sections of this study are organised as follows, section 4.2 describes the dataset, section 4.3 describes the econometric model in further detail and the results are presented in section 4.4. Also, information proxies are employed in section 4.5 and section 4.6 summarises.

4.2 Data

This project uses a unique dataset based on a survey of 3,670 French households. It offers a novel feature; households not only disclose their own perceptions and expectations of the stock market but they also state their beliefs about the actions and information of other households. This is an excellent feature that makes the analysis of this paper cutting-edge.

The 2014 wave of the PAT€R survey is used, which is representative of the population by both age and wealth. The main variables of interest are the participation of households in the stock market (both a dummy variable and a percentage of wealth are available), the perceptions of the stock market that households have (the mean and variance of their beliefs about recent stock market realisations) and their expectations for future returns (the mean and variance of their expectations about future returns. Households' perceptions of other French are also key variables using which analysis of coordinational motives is conducted. Perceptions of others' participation is the percentage of the French population that a household thinks invests in the stock market whereas perceptions of others' information is the percentage of French that the household thinks is informed about the stock market.

Vital household characteristics that are used include the age, gender & education of the respondent and household income, wealth & assets. The age of participants fits a normal distribution with a little negative skew and spans from 19 to 94 years, with an average of 54. 46% of respondents are male and 54% are female whilst 38% are educated at the level of college or higher and 60% are married. Middle categories of assets are the most popular; between 75,000 & 449,999, representing 60% of the distribution. Having less than 8,000 in assets is disproportionately common at 14% of the distribution. Being married, being male and investing in the stock market are all associated with higher asset levels yet higher education surprisingly has the opposite association. Higher categories of savings are less popular than lower levels of savings and 32% have no savings. Being married, being male, investing in the stock market and higher education all have the inverse effect on savings; they save less and have more assets. This is possibly because savings and assets are substitutes. Income has quite a uniform distribution across categories except that 20,000 to 29,999 is twice as common and 40,000 or more is half as common.

Participation in the stock market is markedly low; 22% but this is completely consistent with commonly found findings in other studies. Being male adds 4% to this figure whereas being female takes 4% from it. Having higher education, an income of 20,000 or more and an income of 40,000 or more are all associated with higher participation, the last considerably so. Being married has little impact on participation. A baffling statistic is that 70% of the sample believes that 0% (to the nearest round number) of the French population participates in the stock market. The mean of perceptions of participation is low at 7%. Males perceive a higher percentage of investors in the stock market than females. Having higher incomes 20,000 or more, higher education or being an investor all considerably raise perceptions of others' participation but marriage has little effect. Interestingly, all subgroups perceive a lower population participation than actually occurs (25%). 68% of respondents believe that 0 of 100 French are informed about the stock market, another surprising statistic. Females have a more extreme opinion (73%) than males (63%). Married folk have similar perceptions to the whole sample but those earning 20,000 or more, with higher education or with stocks have a far lower incidence of perceiving stock market ignorance. What is more, these subgroups perceive a higher average informed proportion of the population (13%, 10% and 11% each compared to 8%). Thus, the average French household appears to be woefully informed about the stock market behaviour and information of others.

Regarding perceptions of recent realisations of the stock market, interviewees appear to be poorly informed. 20% of individuals believe that the stock market did not change and only 52% thought that it grew. Males appear to be better informed than females; 58% compared to 46% correctly identified a positive recent return. Being married penalises perceptions by 1.5% yet earning 20,000 or more and investing in the stock market are all associated with a higher proportion of individuals identifying, correctly, positive returns. Surprisingly, having higher education decreases to 50%, the proportion of individuals that identified the positive return realisation. Investors correctly identified a positive return 64% of the time but those earning 40,000 or more did better, 71%. Hence, higher income and investing are all related to being better informed but the married, more educated and female subgroups are worse informed.

Expectations of future stock market returns are low. 31% of individuals believe there will be no change and only 45% believe that there will be positive returns. Compared to the population, 3% more males expect a positive return whereas 3% fewer females do. So females are less optimistic. Married individuals are slightly more pessimistic (44% expect a positive return) but investors, those earning 20,000 or more and those with higher education are all more optimistic that the overall sample (58%, 52% & 55% each). Those earning 40,000 or more are the most optimistic, 2% more than the subsample of investors at 60%. So higher incomes and holding stocks are associated with both being better informed and more optimistic, perhaps implying a link between the two, but those with higher education despite being worse informed are actually more optimistic than the average.

The pertinent conclusions from the data summary are that the picture for the key variables (%p, %i and participation) changes greatly when the sample is restricted to include either only investors or only those earning 40,000 or more. So some of the peculiarities (low stock market participation and unbelievably low perceptions of others for both %p and %i) can be explained a bit by subdividing the data. Being in the top income category generally has the strongest effect of the subgroups and the higher earners tend to be much better informed. However, education gives a mixed picture. Those with college or higher education are not very well informed (unlike investors and the high earners) but they are optimistic (as the investors and high earners are). So it seems that their optimism could be driven by something different to accurate beliefs (unlike investors and the high earners).

4.3 Econometric model

The main question of this study is how do strategic motives determine household participation in the stock market and if so, in which direction? Yet it is important also to consider what else determines participation. Following are the econometric specifications used along with some explanation of the variables included.

Estimation is with a probit model for participation (actions) in the stock market, which is itself a discrete choice of whether or not to own a portfolio. The OLS technique is used to analyse expectations & perceptions and both of them are continuous variables with beliefs about recent and future returns ranging from serious decline to strong growth of 30% (the highest category on the survey) in each direction.

The econometric specifications for: 1. actions, 2. expectations & 3. perceptions are:

1.

$$Pr(Par_i = 1|\%P, \%I, E, X, WIS) = \Phi(\beta_0+\beta_1\%P+\beta_2\%I+\beta_3P+\beta_4X+\beta_5WIS)$$

2.

$$Exp_i = \beta_0 + \beta_1\%P + \beta_2\%I + \beta_4X$$

3.
$$Per_i = \beta_0 + \beta_1\%P + \beta_2\%I + \beta_4X$$

where

- %P = perceptions of others' participation
- %I = perceptions of others' information
- P = perceptions (mean & s.d.)
- X = characteristics: gender, education, risk aversion, marriage
- WIS = endowment: wealth, income, savings.

Perceptions of others' participation and actions are the key variables in the study for addressing the research question and from whose coefficients this project draws its conclusions. Expected future returns in the stock market play no significant role in determining household participation and suggest unusual behaviour that households prefer both higher risk and higher returns. Therefore they are omitted as there is no apparent way of interpreting them economically and they have no econometric significance. This is despite the fact that expectations ought to include perceptions as well as further information that households deem relevant about future returns. Perceived returns however, are typically econometrically significant and economically they can impact actions as they represent the most recent information received, which should be the most pertinent information the household has. So expectations are omitted and perceived returns are included. Actually, expected & perceived returns are highly correlated (68.9%) meaning that it is not worthwhile to include both in the econometric specification concurrently.

The previous chapter showed that age demonstrates a significant effect on how well households are informed and information critically contributes to the participation of households. Simply observe the persistently significant coefficients at the 1% level of perceived returns and their standard deviation. Quite rationally a higher perceived return (better information) induces greater household participation whilst lower certainty about returns (noisier information i.e. a higher standard deviation of perceived returns) discourages participation. Age's positive coefficient indicates that the old participate more in the stock market than the young and this is typically because they are better informed. Also the negligible effect of age's square suggests that the relationship is linear, stock market participation does not diminish in old age. Males generally participate more in the stock market although the married are more reluctant to do so.

The variable CARA is a proxy for the risk aversion of a household. It is developed from a question posed to households concerning their willingness to participate in a lottery. Naturally, more risk averse households will have a lower tendency to participate in the stock market even if their characteristics over endowments (wealth, income & savings), education, age and others are identical. So risk aversion is an important regressor theoretically. Moreover, statistically it demonstrates its own importance at it remains significant at the 1% or 5% levels and, as expected, has a negative impact on households participation meaning the risk averse really do participate less given the same characteristics.

Endowments are represented by the three variables: savings, income and wealth and each is split into four categories, the lowest always being the null (omitted) category. Not only can endowments be expected to have a significant impact on stock market participation directly through the sheer ability to purchase stocks and bear the costs of doing so but also through their positive relationship with education, marriage and age. Education is included given that it could account for how well a household understands the stock market and its exclusion could cause omitted variable bias when including endowment variables. Endowments are generally positively related to stock market participation. That is, higher earners, higher savers and the wealthier participate more. However, the impact of savings levels off: having positive savings increases participation but having high savings does not. This is possibly because those in the highest category substitute assets for savings so they build savings rather than invest in the stock market. Also the low earners participate less in the stock market than non-earners, likely because the old are retired but are also stock market participants - participation increases with age but being in the non-earning category (Income$<$12,000) has a convex relationship with age: the young are likely to earn more in middle age but will return to it when old. Education is always positively related to participation despite the higher educated being generally worse informed about the stock market than the average. However, even though the positive relationship exists, only being college or higher educated has a significant impact.

A significant issue in the dataset is the non-responses of participants to questions, which significantly affect the explanatory variables. We tackle this by the use of non-response dummy variables. Right hand side variables with missing values are given a numerical value (-1) and a non-response dummy variable is assigned to that variable, which takes a value of one when the explanatory variable is -1 and 0 otherwise. This technique ensures that all information about the explanatory variables is squeezed out for understanding the participation of households in the stock market without falsely attributing explanatory power to missing values.

4.4 Regression analysis

The main question of this study is do coordination motives determine household participation in the stock market and if so, in which direction? Hence a basic regression on whether or not households invest in the stock market is run to check the role of coordinational motives, households' perspectives on the participation and information of others in the population. This can be seen in Table A.26 Appendix A.3.2. Crucially, perceptions of others' are significant and this forms the basis of the major finding of this paper.

4.4.1 Households act upon others' perceived information

The base regression specification is that described by 1 in section 4.3. Adding coordination motives individually (specifications 2 & 3) to it sees others' participation significantly affect household participation but others' information not. So in a simple sense households are motivated to coordinate together and evidence has been found to affirmatively answer the study's question in the positive sense. Interestingly, both increase in significance when included together (specification 4) and this indicates that they are dependent in some sense. That is, there is an effect of others' participation that runs through others' information and vice versa. Adding an interaction term between the two confirms this as the significance and some of the impact of others' actions on household participation disappears. Yet others' information retains its significance at the 10% level and its impact increases marginally. This indicates that there is an impact on household participation of believing that other households participate and are informed. That others' perceived information survives the addition of the interaction term at the 10% level gives evidence that the household really cares about others' actions only because of the perceived information content contained therein. This is understandable when one realises that coordination motives have a correlation of 78.7%. The interaction term itself is insignificant and has little effect from which the conclusion arises that there is no significant effect on participation in the stock market of coordinating with both others' actions & information yet it is through perceptions of others' information that perceptions of their actions derives significance. Hence others' actions are important because they convey something about what others know.

Regression analysis on stock market participation unearths a strange finding; when strategic coordination over information and actions are present, motives are complementary in actions but substitutionary in information, this is persistent across specifications. So perceived population participation and information indicate distinct strategies: perceptions of others' participation is positive, a complementary strategy, whilst perceptions of others' information is negative, suggesting a substitutionary strategy. This indicates that if other households are better informed and hence participate more (better informed households participate more), a household will simultaneously decrease its own participation due to the improved information of others and increase its participation due to others' increased participation. So the household interprets the same signal in two opposite directions.

4.4.2 A separation theorem for signs disagreement

Consider the separation theorem of Hellwig & Veldkamp (2009). Under a preference for complementarity in participation, a household wishes to take the same action as others requiring it to have the same information set and signals as them. Conversely, under a preference for substitution in participation, a household wishes to take a different action to others requiring it to observe different signals and have a different information set to them. The following model represents this.

$$a_i = (1 - r)E_i[s] + rE_i[a] \tag{4.4.1}$$

where $a \equiv \int_j a_j$.

A household chooses action a_i conditional on its information set. It forms expectations about the true state of the world, s, (the future stock market return) and the average action of others, a. Its preference for complementary or substitutionary participation is described by the parameter r. Econometrically, if $r \neq 0$ (estimated by α_r/α_m) then there is evidence of coordination motives and in particular $r < 0$ indicates a substitutionary motive for the household whilst $r > 0$ suggest it prefers a complementary strategic approach.

This theorem helps interpret the conflicting signs of perceptions of others' actions & information by providing a way of modelling strategic motives that can subsequently be indentified in the dataset. Then it is possible to apply a finance story to the distinct strategic motives observed using a standard formula for the return on a risky asset and to distinguish between dividend and future price effects using the separation theorem. The formula is below.

$$r_{t+1} = \frac{p_{t+1} + d_{t+1}}{p_t} \tag{4.4.2}$$

The return on a risky asset tomorrow decreases with its price today but increases as its price (at which to sell) and dividend tomorrow rise. Suppose two cases, firstly that the asset is "cheap", a low price today, p_t, compared to the price it will sell for tomorrow, p_{t+1}, and the dividend that will be received, d_{t+1}. That is a high return tomorrow, r_{t+1}, can be earned. A household would like to buy the asset and would require an information set and signals that are different to other households. Otherwise households with the same information would also buy the asset causing the price today to rise and eating away tomorrow's return. So a household wishing to trade cheap assets wants to know what others do not and is, as a result, dissuaded from participating when others are better informed as fewer such opportunities exist.

Secondly, suppose that a household expects the price tomorrow to rise ceteris paribus. That is because it expects other households to buy the asset. Hellwig & Veldkamp (2009)'s theorem indicates that other households acting similarly have the same information set and if the household wants to copy their purchase it must hold the same information set too. It's strategic complementarity is based on the expectation that the price will be higher tomorrow, pushed up by others buying, and so the household can earn a return from holding it. Therefore, to benefit from this sort of opportunity, the household wants to know what others know and participates more when other households themselves participate more.

Thus Hellwig & Veldkamp (2009) provides a theory to interpret the empirical disagreement in the signs of strategic motives. Households respond positively to others' participation because they want to exploit higher expected prices tomorrow to make a return but respond negatively to others being informed which could exhaust opportunities to invest in cheap assets today whose prices will rise.

The market microstructure literature also offers a possible explanation for the signs disagreement. It typically predicts that an investor will participate less in the stock market as others become better informed because he believes he will lose wealth to other, better informed agents. Yet it also expects him to participate more as others participate more because he understands that they are benefiting from rising returns and hence he could exploit general growth in the market.

4.4.3 Are household perceptions & expectations consistent with signs disagreement?

Households' perceptions are not significantly affected by perceptions of other's participation but they are by perceptions of others' information in simple specifications (2 & 3), see figure A.29. Specifications (4) & (5) include both coordinational motives together. Here, perceptions of others' participation has a positive coefficient meaning that as others participate more, the household is better informed and so is more likely to participate, consistent with the positive coefficient in the regressions on actions. Perceptions of others' information has a negative coefficient and this reduces how informed the household is and, ipso facto, the household is less likely to participate, also consistent with the participation results. Expectations are affected by coordinational motives too. In figure A.28, specifications (4) & (5), only perceived actions of others has a significant influence when both motives are included. The motives also maintain their contrasting signs. Expectations typically have a positive (but insignificant) effect on the chances of participating. So as others are perceived to participate more, expectations rise and so does participation whereas when others are perceived to be better informed, expectations fall and participation decreases. Again, these effects flow consistently from expectations & perceptions to participation.

Hence household behaviour appears to be self-consistent but there is little clue so far about why others' participation and information have opposing effects on stock market participation. We propose that information effects within the household can explain both the mismatch and the significance of strategic motives. This project next looks into this possibility and what determines the household's expectations and perceptions.

4.4.4 Households coordinate over expectations & perceptions

Households form expectations (figure A.28) using others' participation and information. When included individually (2 & 3) they are both positive and significant. However, when both are included together (4 & 5) only others' actions remains significant and others' information becomes negative. It is somewhat surprising that others' actions retain significance at the expense of others' information, completely the opposite to regressions on participation and they are strongly significant too (at the 5% level). In fact it appears that the significance of others' information in (3) is truly through others' actions so when others' actions is introduced, it loses its significance meaning the household does not really care about others' information when forming expectations of future returns but only about what others actually do. Further, when the interaction term is added, there is a small transference of impact from population information to population participation. This indicates that there is a small impact of others' information through their actions. This is the opposite direction to the participation regressions; households care only about others' information because of what it reveals about their actions. Thus in expectations it is others' actions that the household finds pertinent whilst others' information is important in its own actions.

Although expectations ultimately do not have a significant impact on actions their positive effect on actions rationally follows from these regressions on expectations: Increased population participation elevates the household's expectations, which encourages its participation whilst higher population information diminishes its expectations and discourages its participation as a result. Age has no significant impact upon expectations and although its effect is positive, this is in a linear fashion (age squared has almost no impact). That age has a positive impact is hardly surprising given that it can convey increased human capital it can certainly also convey greater knowledge and understanding of stock markets. Note also that there is a strong and positive relationship between perceptions and expectations (39.44% correlation) that shows the better informed are generally more optimistic. However, such an observation is a bit unanticipated because the benefits of age, if through knowledge and understanding, could decrease when an individual retires (losing access to knowledge sources and networks) and experiences reduced mental faculties. In that case a concave relationship between age and expectations would exist

Married individuals have significantly lower (5% significance level) expectations, as with perceptions. Hence married individuals are both worse informed and more pessimistic. It is not apparent why married folk would be more pessimistic and worse informed because although they are more likely to be older (over 50% are married by 37) they are just slightly more risk averse (CARA = 35 for married and 34 for unmarried individuals) and lower earners. Risk aversion (CARA) has a significant and negative impact on both expectations and perceptions. Expressed differently, the risk averse are both worse informed and more pessimistic. This translates into their actions; they are less likely to participate in the stock market so their behaviour is consistent in the data. Endowments (wealth, income and savings) are not used to explain expectations due to possible endogenous effects. Those who are wealthier and have higher incomes could have earned them from the stock market and so expect higher returns than those who have not.

Strategic motives play a significant role in the formation of household perceptions (figure A.29). Perceptions of others' participation has no significant impact (even when the sole strategic motive, specification 2) but others' information does. So households consider others' information to be important for what they learn but not others' participation. This is akin to actions but opposite to expectations formation. Specification 4 shows that perceptions of others' information is significant at the 10% level but when the strategic motives interaction term is introduced in specification 5 this rises to the 5% level. Concurrently the coefficient of others' perceived actions falls to a negligible value, a transference of impact & significance. This indicates that there is some effect of information contained in others' perceived actions that is attributed to perceived population participation when the interaction term is excluded. Further evidence is that perceived information of others has a similarly significant coefficient in specification (3) when others' perceived actions are omitted. Thus, it seems that the household derives an impression about what others know from their perceived actions. Again, the household appears to care truly about only one strategic motive at once, in actions and perceptions it is others' information yet in expectations it is their participation. The puzzle is that the two coordination effects are again opposites; households pick a complementary strategy over actions but a substitutionary one over information. Thus information gleaned directly from perceiving others' information is interpreted in an opposite way to that gleaned indirectly from their actions. Yet the interaction term may be further informative: upon its inclusion, all information content is now seemingly correctly attributed to perceptions of others' information and a disagreement in coefficients' signs is no longer important. So when others are better informed, the household becomes worse informed itself and participates less in the stock market.

Other factors include education, which surprisingly has no significant impact on how informed the household is, although higher education has a positive effect on perceptions. Age is, as with expectations, insignificant, positive and linear, akin to its effect on expectations and the married and risk averse are again significantly worse informed.

Thus the evidence shows that a household forms its expectations and perceptions of returns, as well as in its participation decision, strategically. Contrariwise, it cares not about both concurrently but strategic motives can be distilled to find that others' information is pertinent for determining participation and perceptions but others' actions matter in forming future returns expectations. So strategic motives are important for households but why this is requires more analysis. This project turns next to investigate if information effects could explain them.

4.5 Additional information variables

Section (4.4) showed evidence for significant strategic motives. This section addresses why households care about them and proposes further information variables as an explanation. A second question that this group of variables can contribute to answering is why strategic motives operate in opposite directions.

The group of additional information variables proposed can be split into four subsets. The first describes the trust a households has generally in others, how much it trusts financial advisers (increasing scale) and that it has no identifiable reason for stock market non-participation. This subset aims to capture household trust in the actions and information perceived about others. The lack of a reason for participating acts as a proxy for the household unwilling to trust others' information & actions for making rational inferences about returns due to its own negative stock market predilections - without a rational reason the household avoids the stock market. Quite simply, if a household is trusting then it will put more importance on the role of strategic motives in its own actions and perceptions because it believes its observations of others to be reliable. So strategic motives would lose significance upon trust's inclusion.

The second subset is about the sources of information that households use. This includes how often friends, family, financial advisers, general media, and specialised media are consulted for information, how often a household relies on friends, family, banker/financial adviser, its own financial knowledge, and the media to make a decision in the stock market, and how often the household consults each type of media (written press, audiovisual, online social networks, financial establishment websites, financial authorities' websites, sites describing investment opportunities and others). This set aims to discover the information sources a household uses and if they and the frequency of use can explain strategic motives. Perhaps the household with certain information sources or frequency of use does not use coordination motives and could be separated from the one that does. For example, frequent consultation of financial advisers and specialised media, efficient sources of information, may be a substitute for observing others' participation and information. Alternatively, perceptions of others may be key facts that households research and this subset could reveal the sources they use to be informed about others.

Next is the inertia set that simply measures how many times the household has purchased financial assets recently. It aims to test if the frequency of the household's activity in the stock market determines its perceptions of others and the importance of those perceptions for its own actions and perceived returns. Inertia could be a good proxy for the quality and quantity of information that a household receives. The more transactions it makes the more times, perhaps, it interacts with a financial expert or observes recent and important stock market information. Inertia could therefore pick up an aspect of how well informed a household is that perceived returns does not. So if inertia is an important additional information variable then strategic motives would lose significance as their information content would be explained.

Last is the knowledge level of the household (self-described financial literacy in five categories: strong, average, weak, very weak and non-existent). This set aims to reveal how well a household processes the information that it has collected. A household's knowledge level can give a fuller and longer term impression of how well informed it is than perceived returns (most recent information update) because the latter could be noisy for some arbitrary reason. Hence if a household seems uninformed on the basis of perceived returns but actually understands financial markets well then it may have accurate perceptions of others (households typically strongly underestimate others' participation and information) and hence make good participation decisions. Hence strong strategic motives could really reflect robust information sets. So if the household's knowledge level is important then strategic motives could lose significance.

4.5.1 Econometric specifications with additional information variables

Following are the econometric specifications with the additional information variables introduced. Estimation is with a probit model for participation (actions) in the stock market and perceptions are still analysed using OLS.

The econometric specifications with information for: 1. actions & 2. perceptions are:

1. $Pr(Par_i = 1|\%P, \%I, E, X, WIS) = \Phi(\beta_0 + \beta_1\%P + \beta_2\%I + \beta_3 P + \beta_4 X + \beta_5 WIS$
   $+ \beta_6 T + \beta_7 S + \beta_8 I + \beta_9 K)$

2. $Per_i = \beta_0 + \beta_1\%P + \beta_2\%I + \beta_4 X + \beta_6 T + \beta_7 S + \beta_8 I + \beta_9 K$

where

- T = Trust in others and markets
- S = Sources of information (friends, family, media etc)
- I = Inertia, no. transactions in period $t-1$
- K = Knowledge level, self-described financial literacy

with %P, %I, P, X & WIS identified in section 4.3.

Now a new econometric model has been outlined for the inclusion of the additional information variables, their effects can be analysed. The next section describes and explains their impact on strategic motives in regressions on participation and perceptions.

4.5.2 Coordination motives are driven by information

The additional information variables are added systematically to the regression on participation with both strategic motives but no interaction term (specification 4 in table A.26). They are also added to the regression on perceptions with both strategic motives and the interaction term (specification 5 in table A.29). The additional information variables are added one subset at a time; firstly trust (Tr), then sources of information (So), then inertia (In) & lastly knowledge (Kn) before two combination are used; trust & sources of information (TS) and inertia & knowledge (IK) before finally all are used together.

Trust & inertia explain coordination in participation

Specification 2 of table A.27 includes the trust subset and immediately sees that significance of others' participation (5% level) and information (10% level) evaporate. Inertia has exactly the same influence (specification 4 of table A.26). This suggests that households account for the actions and information of others in part because they are trusting. Examination of the trust subset shows it be highly significant. Trust in others is actually a discouragement to participating for the household. The effect of strategic motives is therefore not that households trust others and so glean information from their actions and information. Instead, trusting others may provide discouraging perceptions and expectations of returns. As it trusts its financial adviser more the household is more likely to invest. So financial advisers may provide encouraging counsel otherwise attained by strategic motives. However, a powerful factor is the lack of a determined reason for stock market non-participation, which gives a strong disincentive. This means that a household can take a lead from others beyond its own rational information processes (perceptions account for information). Additionally, perceived returns loses significance at the 1% level under the trust group's addition. This indicates that some of the household's perceptions are explained by its trust in others or its financial adviser.

Inertia is added to the econometric model in specification 4 of table A.27. It is also able to explain the significance of strategic motives and is highly significant, at the 1% level. The higher the number of transactions made in the recent past, the more likely the household is to participate in the stock market. Yet why does the participation level in the stock market yesterday influence the household's strategic motivations? This is not a case in which it has significantly more information from participating yesterday because perceived returns remain relevant at the 1% level. However, looking at the standard deviation of perceptions and CARA (proxy variables for risk aversion and perceptions of risk) can offer a clue. The former loses significance at the 10% level whilst the latter does at the 1% level. So participation yesterday can explain reliance on others' participation and information through the uncertainty reduction that strategic motives provide. Observing others does not provide more information but lowers uncertainty about what the household perceives.

Further inspection of the results reveals that the significant and hump-shaped impact of the household's age disappears under inertia's inclusion. This is an effect also seen in specification 3 when sources of information are used. Thus age has a concave impact on participation because sources from which to gather information are better when the household is middle-aged than young or old and because middle-aged households make more stock market transactions than young and old ones.

Strategic motives in perceptions is robust to additional information

The additional information subset is used to investigate perceptions with the results recorded in table A.30. Participation of others is never significant but no combination of additional information subsets can explain the significance of others' information at the 10%, though trust, sources of information and inertia can by themselves provide an explanation at the 5% level. So others' information is robust to the inclusion of additional information variables at the 10% level. The interaction term transfers significance from others' actions through their information content to others' information meaning the coefficient remains more significant than without the interaction term.

Risk aversion is another important but more robust determinant of the household's perceptions. CARA consistently remains significant at the 1% level meaning the risk averse are significantly worse informed. Age continues to have no significant impact on the household's information but marriage remains weakly significant, that is at the 10% level. In fact, it gains significance when trust is included, which means that there is an effect of marriage on the household's perceived returns that relies on the level of trust that it has. Indeed, all trust variables are significant. Thus married folk are worse informed and the effect is significantly enhanced by their trust characteristics. Therefore perceptions of population information is an important determinant of the household's perceptions, robust to additional information effects, but being married and risk averse both significantly reduce the perceptions that the household has.

The important sources of information

The sources of information subset is useful for identifying where households glean information from and which sources are important in the participation and perceptions decisions. Specification 3 of table A.30 demonstrates which sources of information affect household perceptions. Media has a significant impact with the general sort lowering perceptions and the specialised sort raising them as both are used more frequently. That general and specialist media would have opposing effects on perceptions could indicate the kind of information provided or that the household prefers its chosen type because it confirms a preselected perception. Alternatively, the household that searches for "cheap" assets may choose one media type and the one that wants assets that will rise in price tomorrow may choose the other. These suggestions require further investigation. Frequently relying on its own financial knowledge and the media both significantly increase the perceptions of the household. Which sources a household relies on for decisions and information can explain the importance of others' information at the 5% level so the strategic motive houses part of their impact, when excluded.

In making its participation in the stock market decision, the household relies on both friends (negatively) and its own financial knowledge (positively). Consulting a financial adviser and specialised media significantly determine its decision with participation rising with the frequency of consultation. Amongst types of media used to gather information, online social websites is the only significant one and the more frequently the household consults it, the less likely it is to invest in the stock market. In all, perceptions' significance is unaffected but CARA's is reduced so such individuals are perhaps not better informed but less risk averse. So the media and financial knowledge already gained are important in the determination of household perceptions whereas the household is seen to lean on its social network (friends and online social websites) for information but not on the general media. Sources of a more social nature tend to have a negative impact on participation and perceptions whereas more dispassionate sources (financial advisers & specialised media) encourage both along with the household's own knowledge.

4.6 Summary

This project investigates the role of strategic interactions between households in the stock market using novel data from the 2014 PAT€R survey wave. We empirically assess household participation in the stock market based upon their perceptions of how other households themselves participate and how well those others are informed about returns as well as using a proxy for a household's own information set (perceived recent market returns) and characteristic variables (age, married etc). Evidence is found for significant strategic motives between households and this is of a substitutionary nature when considering others' information. Adding an interaction term between the strategic motives sees the role of others' information grow to the detriment of others' participation, suggesting that others' actions are useful in that they provide insights about others' information. Hence the effects of others' information and actions appear to be opposite. This may appear peculiar but the market microstructure literature typically expects an investor to participate less the better informed others are as he believes better informed agents will outcompete him. Yet it also expects an investor to participate more as others participate more because he understands that they are benefiting from rising returns and hence he could also profit from a rising tide in the market.

Additional information variables representing how trusting a household is and its inertia in the stock market both provide ways to explain the significance of strategic motives. The latter particularly suggests that lower inertia (more trades) reduces household uncertainty about returns and strategic motives appear to otherwise house this effect.

Formation of household expectations and perceptions are both also investigated, with significant evidence of strategic motives found too. In the case of perceptions, additional information variables: trust, sources of information & inertia only partly explain the significance of strategy. So strategic motives along with marriage and risk aversion remain robust determinants of household perceptions. Studying sources of information reveals that households tend to rely on their own knowledge (positively) and the media (mixed) for participation and perceptions choices whilst social sources (friends & online social sites) have a negative impact and dispassionate sources (media & financial advisers) have a positive impact on perceptions and the chances of participating.

Using the findings of our empirical studies we make a simple conclusion. As there is substantial heterogeneity in expectations and perceptions of returns and that the median investor is relatively unsophisticated in the empirical literature, we claim that as strategic substitutes prevail in participation choices even amongst them, a portion of the excess volatility observed in stock markets may be driven by strategic expectations motives.

5

Conclusion

This thesis makes a combined contribution to the fields of Bounded Rationality within Macroeconomics and Portfolio Theory within Financial Economics. It endogenously models the portfolio and information choices of households to explain, theoretically, two related stylised facts: under diversified portfolios and widespread heterogeneous beliefs. It also empirically investigates the existence and effects of strategic motives between households in the stock market.

Contrary to standard portfolio theory investors hold highly concentrated portfolios. Van Nieuwerburgh & Veldkamp (2010) explain this through a feedback effect, endogenously acquired information reduces uncertainty about an assets payoffs, persuading the investor to hold more of it and conveying benefits to specialisation as further information will inform about more of the expected portfolio thus concentrating his portfolio.

Paper one employs a dynamic framework to test this rationale and find the conditions under which the standard diversified portfolio is recovered. That the results diverge from the static model solutions under a late resolution of uncertainty illustrates the relevance of the extension to an infinite horizon. As hedging demands arise in the steady state under an exogenous learning technology, the state variables are sufficient for a dynamic investment strategy despite time independent returns. Further, a learning factor scales the portfolio choice in the transition dynamics under an endogenous learning technology. The latter disappears in the steady state meaning that whilst learning, the agent invests more aggressively assured of greater certainty about returns both today and tomorrow. So not only do dynamic investment strategies matter but so does the ability to acquire signals endogenously.

A strict preference for early resolution of uncertainty is sufficient in Van Nieuwerburgh & Veldkamp to replicate concentrated portfolios via specialised learning. With the same preference structure and a dynamic framework the numerical solution to paper one suggests that strict unobservable preferences are not necessary. A strongly superior expected Sharpe ratio on one asset than another can induce the investor to specialise and concentrates his portfolio. This indicates that underdiversification can be addressed empirically through the reported stock market beliefs of households without relying on unobserved preferences. The result shows that dynamic as well as myopic decision-making can account for concentrated portfolio choice through the learning strategy employed. Interestingly, in a dynamic setting this result fortifies Van Nieuwerburgh & Veldkamp's feedback effect between the gathering of noisy information and expected portfolio holdings that produces a concentrated static portfolio.

Endogenous information acquisition renders standard portfolio theory suboptimal in Van Nieuwerburgh & Veldkamp (2010) and an extension to multiple periods renders the Merton portfolio suboptimal in the first project. Additionally, widespread heterogeneous beliefs are documented by the 2014 wave of the PAT€R survey and Coibion et al (2015). The second paper tests these findings using numerical simulations.

This project can explain belief heterogeneity. In a multi-period investment model a dynamic investor may specialise in one asset when he, a prior, sees a superior process driving it. Specialisation increases his expected utility of future wealth because he expects to hold more of that asset and learning decreases his expected risk. A distinct information set from another agent with different initial priors develops under non-observation of returns and signals that confirm prior beliefs.

The agent's demand for information is unresponsive to return fluctuations as he either learns the underlying process driving returns or his acquired signals support his initial beliefs. Both lead to a constant learning strategy. The result contrasts the anti-cyclical (regarding the business cycle) demand in Coibion et al (2015) and suggests that agents in reality do not treat returns as though they are generated by an underlying, learnable process. The optimal portfolio under endogenous information acquisition and dynamic investment outperforms the Markowitz and Merton portfolios. The investor defaults with a reduced probability through conditioning on a coarse information set. VN&V's agent with CARA wealth and indifferent time preferences chooses a generalised learning strategy. Yet, the present simulations reveal an agent who may actually pick a specialised learning strategy permanently giving evidence that static generalised learning strategies do not hold in a dynamic model.

Project three investigates the role of strategic interactions between households in the stock market using novel data from the 2014 PAT€R survey wave. It empirically assesses household stock market participation using their perceptions of how other households themselves participate and how well those others are informed about returns as well as using a proxy for a household's own information set (perceived recent market returns) and characteristic variables (age, married etc). Evidence shows that significant strategic motives exist between households and this is of a substitutionary nature when they account for the information sets others hold. Adding an interaction term between the strategic motives sees the role of others' information grow as the role of others' participation dissipates. This indicates that others' actions are useful because they can provide insights about others' information. Additional information variables representing how trusting a household is and its inertia in the stock market both provide ways to explain the significance of strategic motives. The latter particularly suggests that lower inertia (more trades) reduces household uncertainty about returns and strategic motives appear to otherwise house this effect.

As expectations and perceptions of returns show substantial heterogeneity and that the median investor is relatively unsophisticated, we claims that as even less advanced individuals use strategic substitutes to guide participation choices, a portion of the excess volatility observed in stock markets may be driven by strategic expectations motives.

The presence of strategic motives in household participation decisions regarding the stock market shows that consideration must be given to the impact of shocks to individual beliefs. As strategic substitutionary behaviour is prevalent, decision makers must be aware that shocks and noisy information may lead to increased volatility. This thesis has found evidence that learning strategies vary with time with the implication that the pattern by which a household's information set develops can itself depend on time. Thus decision makers should be away that not only do actions depend on the outcome of an information choice by an agent (quantities demanded) but that information choices themselves depend on the outcomes of actions taken. Thus simultaneous relationships should be acknowledge, in the opinion of this thesis, in decision making to avoid the potential biases inherent in assuming a single-direction causality. A final caution to decision makers concerns apparently risky portfolios that seem concentrated beyond what is sensible. Increasing returns to information processing may provide incentives over multiple investment periods for investors to hold such positions. However this does not insulate them against aggregate level risks nor against unlearnable adverse shocks.

Through the endogenous modelling of agent-based learning Bounded Rationality has contributed to several areas of Economics; it matches observations from laboratory experiments with human participants, helps solve the "forward guidance puzzle", can model agents that do not perceive future tax repayments a la Ricardian Equivalence, replicate how agents save "too late" for retirement as they hurriedly slash expenditure and it can match the sensitivity of inflation expectations and the term structure to shocks as found in data. However, this author would hope to see future applications in further areas particularly concerning household finance. This thesis could certainly be extended, with some concerted effort, in using stochastic underlying processes. Also, applying Bounded Rationality to the risk premium puzzle could provide a new insight to that long-standing focus of Economics.

APPENDIX A

Appendix H

A.1 Appendix: Paper 1

This appendix contains four sections. First is the step by step method for obtaining the optimal risky asset choice in the exogenous learning model with CARA preferences, along with the conditions for returning it to the myopic benchmark and a proof that an agent's mean beliefs will converge in the limit. Then the model with endogenous learning and a preference for late resolution of uncertainty is solved for the optimal portfolio and posterior variance choices along with the model's RE benchmark. Next the endogenous model with indifference to the time of uncertainty resolution is solved for the steady state portfolio choice and the numerical approximation results to its decision rules over portfolio and posterior variance choices are given. Lastly, the Kalman Filter used in this paper is derived too.

A.1.1 When learning is exogenous: CARA preferences

The Bellman equation is the following.

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q\}}||E_t^-\{u(W_{t+1})|\mu_t^-, \Sigma_t^-\}$$
$$+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$$

Substitute in the wealth process and the functional form.

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q\}}||E_t^-\{-exp(-\rho(rW_t + q_t'(f_{t+1} - p_t r)))|\mu_t^-, \Sigma_t^-\}$$
$$+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$$

Take the expectation, treating $exp(W_{t+1})$ as a lognormal variable.
$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q\}}|| - exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r) + \frac{1}{2}\rho^2 q_t'\Sigma_t^- q_t))$
$+ \beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$

Take the first order conditions in terms of the portfolio allocation, q_t.
$V_q(W_t, \mu_t^-, \Sigma_t^-) = -(-\rho(\mu_t^- - p_t r) + \rho^2 q_t'\Sigma_t^-)exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r) + \frac{1}{2}\rho^2 q_t'\Sigma_t^- q_t))$
$+ \beta E_t^-\{V_q(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$

The form of $V_q(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)$ is unknown so use the Benveniste-Scheinkmann condition to find it. The Benveniste-Scheinkmann condition is the envelope condition that is specified in theorem 4.10 in Stokey & Lucas (1989) page 84. Firstly, take the first order conditions in terms of the state variable, W_t.
$V(W_t, \mu_t^-, \Sigma_t^-) = (\rho r)exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r) + \frac{1}{2}q_t'\Sigma_t^- q_t))$

Move this first order condition forwards by one period.
$V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) = (\rho r)exp(-\rho(rW_{t+1} + q_{t+1}'(\mu_{t+1}^- - p_{t+1} r) + \frac{1}{2}q_{t+1}'\Sigma_{t+1}^- q_{t+1}))$

Substitute this into the first order condition in terms of q_t to find the Euler equation.
$(-\rho(\mu_t^- - p_t r) + \rho^2 q_t' \Sigma_t^-) exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r) + \frac{1}{2}\rho^2 q_t' \Sigma_t^- q_t))$
$= \beta E_t^- \{(\rho r) exp(-\rho(rW_{t+1} + q_{t+1}'(\mu_{t+1}^- - p_{t+1} r) + \frac{1}{2} q_{t+1}' \Sigma_{t+1}^- q_{t+1}))\}$

Impose the steady state. Firstly, $\sigma_t^- = 0$, which means that $\Sigma_t^- = \Sigma_\eta$ in the steady state because the agent knows everything about the risky assets, particularly the true variance. Next, note that the agent does not expect the price to vary. Also, set the steady state wealth to $\bar{W}$, quantity to $\bar{q}$ and mean belief to $\bar{\mu}$.
$(-\rho(\bar{\mu} - p_t r) + \rho^2 \bar{q} \Sigma_\eta) exp(-\rho(r\bar{W} + \bar{q}'(\bar{\mu} - p_t r) + \frac{1}{2}\rho^2 \bar{q}' \Sigma_\eta \bar{q}))$
$= \beta E_t^- \{(\rho r) exp(-\rho(r\bar{W} + \bar{q}'(\bar{\mu} - p_{t+1} r) + \frac{1}{2}\bar{q}' \Sigma_\eta \bar{q}))\}$

Take the expectation. Let the expectations of prices be static, $E[p_t], E[p_{t+1}] = \bar{p}$
$(-\rho(\bar{\mu} - \bar{p} r) + \rho^2 \bar{q} \Sigma_\eta) exp(-\rho(r\bar{W} + \bar{q}'(\bar{\mu} - \bar{p} r) + \frac{1}{2}\rho^2 \bar{q}' \Sigma_\eta \bar{q}))$
$= \beta(\rho r) exp(-\rho(r\bar{W} + \bar{q}'(\bar{\mu} - \bar{p} r) + \frac{1}{2}\bar{q}' \Sigma_\eta \bar{q}))$

Notice that the exponential terms can be cancelled because their arguments are identical.
$-\rho(\bar{\mu} - \bar{p} r) + \rho^2 \bar{q} \Sigma_\eta = \beta(\rho r)$

Solve for the steady state portfolio allocation.
$-(\bar{\mu} - \bar{p} r) + \rho \bar{q} \Sigma_\eta = \beta r$
$\rho \bar{q} \Sigma_\eta = \bar{\mu} - \bar{p} r + \beta r$
$\bar{q} = \frac{\bar{\mu} - \bar{p} r}{\rho \Sigma_\eta} + \frac{\beta r}{\rho \Sigma_\eta}$

Will mean beliefs converge?

When using the steady state, it is important to ask whether or not the agent learns the true mean ($\mu_t^- = \theta$) before his uncertainty reaches zero ($\sigma_t^- = 0$), that is when his beliefs converge. Here, firstly the fact that convergence in beliefs will never take place, in other words, 0 is an asymptote for prior uncertainty (σ_t^-) is shown, meaning that learning will not converge before the agent can learn the true mean. Secondly, it is shown that the agent's beliefs about both the mean and uncertainty converge at the same rate.

Does $\sigma_t^- \to 0$ before $\mu_t^- \to \theta$?

Consider the formula by which σ_t^- is updated.

$$\sigma_{t+1}^- = \frac{\sigma_t^- \Sigma_\eta}{\sigma_t^- + \Sigma_\eta}$$

When does $\sigma_{t+1}^- = 0$?

$$0 = \frac{\sigma_t^- \Sigma_\eta}{\sigma_t^- + \Sigma_\eta}$$

So

$$0 = \sigma_t^- \Sigma_\eta.$$

Hence either $^{(1)}$ $\Sigma_\eta = 0$, which is not true because the underlying distribution has a non-zero variance meaning that returns vary over time, or $^{(2)}$ $\sigma_t^- = 0$. This is also not true because this would require that $\sigma_0^- = 0$ and this would prevent the agent from ever having any incentive to learn from the beginning. He would be supremely confident in his beliefs from the start.

Therefore, $\sigma_0^- = 0$ will never occur, meaning that it is bounded below and here at

0. This also means that the agent will always keep on learning about the mean because learning will never absolutely converge. $\therefore$ as $\sigma_t^- \to 0$, $\mu_t^- \to \theta$.

Solve forwards for updating uncertainty and the Kalman Gain

Consider the formula for updating the agent's uncertainty.

$\sigma_{t+1}^- = \frac{\sigma_t^- \Sigma_\eta}{\sigma_t^- + \Sigma_\eta} = K_t \Sigma_\eta$

Now solve this expression for the prior uncertainty today.

$\sigma_{t+1}^-(\sigma_t^- + \Sigma_\eta) = \sigma_t^- \Sigma_\eta$

$\sigma_{t+1}^- \sigma_t^- + \sigma_{t+1}^- \Sigma_\eta = \sigma_t^- \Sigma_\eta$

$\sigma_{t+1}^- \Sigma_\eta = \sigma_t^- \Sigma_\eta - \sigma_{t+1}^- \sigma_t^-$

$\sigma_{t+1}^- \Sigma_\eta = \sigma_t^- (\Sigma_\eta - \sigma_{t+1}^-)$

$$\sigma_t^- = \frac{\sigma_{t+1}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+1}^-} \tag{A.1.1}$$

Now move the expression forwards one period.

$$\sigma_{t+1}^- = \frac{\sigma_{t+2}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+2}^-} \tag{A.1.2}$$

Now move the expression forwards one period again.

$$\sigma_{t+2}^- = \frac{\sigma_{t+3}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+3}^-} \tag{A.1.3}$$

Substitute (A.1.2) into (A.1.1).

$$\sigma_t^- = \frac{\frac{\sigma_{t+2}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+2}^-}\Sigma_\eta}{\Sigma_\eta - \frac{\sigma_{t+2}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+2}^-}}$$

Simplify.

$$\sigma_t^- = \frac{\frac{\sigma_{t+2}^- \Sigma_\eta^2}{\Sigma_\eta - \sigma_{t+2}^-}}{\frac{\Sigma_\eta^2 - \Sigma_\eta \sigma_{t+2}^- - \sigma_{t+2}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+2}^-}}$$

$$\sigma_t^- = \frac{\sigma_{t+2}^- \Sigma_\eta^2}{\Sigma_\eta^2 - 2\Sigma_\eta \sigma_{t+2}^-}$$

$$\sigma_t^- = \frac{\sigma_{t+2}^- \Sigma_\eta}{\Sigma_\eta - 2\sigma_{t+2}^-} \tag{A.1.4}$$

Substitute (A.1.3) into (A.1.4).

$$\sigma_t^- = \frac{\frac{\sigma_{t+3}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+3}^-}\Sigma_\eta}{\Sigma_\eta - 2\frac{\sigma_{t+3}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+3}^-}}$$

Now simplify.

$$\sigma_t^- = \frac{\frac{\sigma_{t+3}^- \Sigma_\eta^2}{\Sigma_\eta - \sigma_{t+3}^-}}{\frac{\Sigma_\eta^2 - \Sigma_\eta \sigma_{t+3}^- - 2\sigma_{t+3}^- \Sigma_\eta}{\Sigma_\eta - \sigma_{t+3}^-}}$$

$$\sigma_t^- = \frac{\sigma_{t+3}^- \Sigma_\eta^2}{\Sigma_\eta^2 - \Sigma_\eta \sigma_{t+3}^- - 2\sigma_{t+3}^- \Sigma_\eta}$$

$$\sigma_t^- = \frac{\sigma_{t+3}^- \Sigma_\eta}{\Sigma_\eta - 3\sigma_{t+3}^-}$$

Solve for future prior uncertainty.

$$\sigma_t^-(\Sigma_\eta - 3\sigma_{t+3}^-) = \sigma_{t+3}^- \Sigma_\eta$$

$$\sigma_t^- \Sigma_\eta - 3\sigma_t^- \sigma_{t+3}^- = \sigma_{t+3}^- \Sigma_\eta$$

$$\sigma_t^- \Sigma_\eta = \sigma_{t+3}^- \Sigma_\eta + 3\sigma_t^- \sigma_{t+3}^-$$

$$\sigma_t^- \Sigma_\eta = \sigma_{t+3}^-(\Sigma_\eta + 3\sigma_t^-)$$

$\sigma^-_{t+3} = \frac{\sigma^-_t \Sigma_\eta}{\Sigma_\eta + 3\sigma^-_t}$

Generalise the expression found.

$$\sigma^-_{t+T} = \frac{\sigma^-_t \Sigma_\eta}{\Sigma_\eta + T\sigma^-_t} \tag{A.1.5}$$

Now consider the formula for the Kalman Gain.

$K_t = \frac{\sigma^-_t}{\sigma^-_t + \Sigma_\eta}$

Move this forwards one period.

$$K_{t+1} = \frac{\sigma^-_{t+1}}{\sigma^-_{t+1} + \Sigma_\eta} \tag{A.1.6}$$

Move it forwards again.

$$K_{t+2} = \frac{\sigma^-_{t+2}}{\sigma^-_{t+2} + \Sigma_\eta} \tag{A.1.7}$$

Substitute the expression for σ^-_{t+2} into K_{t+2} above.

$K_{t+2} = \frac{\frac{\sigma^-_{t+1}\Sigma\eta}{\sigma^-_{t+1}+\Sigma\eta}}{\frac{\sigma^-_{t+1}\Sigma\eta}{\sigma^-_{t+1}+\Sigma\eta}+\Sigma_\eta}$

Simplify.

$$K_{t+2} = \frac{\frac{\sigma^{-}_{t+1}\Sigma_{\eta}}{\sigma^{-}_{t+1}+\Sigma_{\eta}}}{\frac{\sigma^{-}_{t+1}\Sigma_{\eta}+\sigma^{-}_{t+1}\Sigma_{\eta}+\Sigma^{2}_{\eta}}{\sigma^{-}_{t+1}+\Sigma_{\eta}}}$$

$$K_{t+2} = \frac{\sigma^{-}_{t+1}\Sigma_{\eta}}{2\sigma^{-}_{t+1}\Sigma_{\eta}+\Sigma^{2}_{\eta}}$$

$$K_{t+2} = \frac{\sigma^{-}_{t+1}}{2\sigma^{-}_{t+1}+\Sigma_{\eta}}$$

Substitute the expression for σ^{-}_{t+1} into K_{t+2} above.

$$K_{t+2} = \frac{\frac{\sigma^{-}_{t}\Sigma_{\eta}}{\sigma^{-}_{t}+\Sigma_{\eta}}}{2\frac{\sigma^{-}_{t}\Sigma_{\eta}}{\sigma^{-}_{t}+\Sigma_{\eta}}+\Sigma_{\eta}}$$

Simplify.

$$K_{t+2} = \frac{\frac{\sigma^{-}_{t}\Sigma_{\eta}}{\sigma^{-}_{t}+\Sigma_{\eta}}}{\frac{2\sigma^{-}_{t}\Sigma_{\eta}+\sigma^{-}_{t}\Sigma_{\eta}+\Sigma^{2}_{\eta}}{\sigma^{-}_{t}+\Sigma_{\eta}}}$$

$$K_{t+2} = \frac{\sigma^{-}_{t}\Sigma_{\eta}}{3\sigma^{-}_{t}\Sigma_{\eta}+\Sigma^{2}_{\eta}}$$

$$K_{t+2} = \frac{\sigma^{-}_{t}}{3\sigma^{-}_{t}+\Sigma_{\eta}}$$

Generalise.

$$K_{t+T} = \frac{\sigma^{-}_{t}}{(T+1)\sigma^{-}_{t}+\Sigma_{\eta}}$$

Compare this expression to the forward expression for Sigma.

$$\sigma^{-}_{t+T} = \frac{\sigma^{-}_{t}\Sigma_{\eta}}{\Sigma_{\eta}+T\sigma^{-}_{t}}$$

Ultimately, both will converge to zero as $T \to \infty$ and both at the same rate because all other components of both expressions are constants and the multiplicative ones are identical. Hence learning will not converge before the agent can learn the true mean of the underlying returns distribution.

A.1.2 When u^- is concave

The outer and inner utility functions are given by the following two expressions.
$u^-(x) = ln(-x)$ and $u^+(W_{t+1}) = -exp(-\rho W_{t+1})$

Working on the Bellman equation

The Bellman Equation is:

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q\}}||E_t^-\{ln(- - E_t^+\{exp(-\rho(rW_t + q_t(f_{t+1} - p_t r)))|\mu_t^+, \Sigma_t^+\})$$
$$|\mu_t^-, \Sigma_t^-\} + \beta E_t^-\{V_t(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$$

As the distribution of $-\rho W_{t+1}$ is normal, the MGF can be used for taking the inner expectation, in terms of posterior beliefs. So the Bellman Equation becomes

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q,\Sigma^+\}}||E_t^-\{ln(exp(-\rho(rW_t + q_t(\mu_t^+ - p_t r))$$
$$+\frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t))|\mu_t^-, \Sigma_t^-\} + \beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}.$$

After some simplification, the result becomes

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q,\Sigma^+\}}||E^-\{-\rho(rW_t + q_t(\mu_t^+ - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t|\mu_t^-, \Sigma_t^-\}$$
$$+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}, \Sigma_{t+1})|\mu_t^-, \Sigma_t^-\}.$$

The outer expectation, E^-, over prior expectations can now be taken. The posterior mean, μ_t^+, and posterior variance, Σ_t^+, can be replaced with the Kalman filter formulas. The resultant Bellman equation follows.

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q,\Sigma^+\}}|| - \rho(rW_t + q_t(E_t^-\{\mu_t^+\} - p_t r)) + \frac{1}{2}\rho^2 q_t' E_t^-\{\Sigma_t^+\} q_t$$
$$+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}, \Sigma_{t+1})|\mu_t^-, \Sigma_t^-\}.$$

The Kalman filter updates prior to posterior expectations. That means that $E^{-}\{\mu^{+}\}$ and $E^{-}\{\Sigma^{+}\}$ can be found. Substitute in the Kalman filter process for updating the expected return and then that for the signal observed. The expectation of the white noise is zero and that of returns next period conditional on the prior expectation is the prior expectation.

$E^{-}\{\mu_t^{+}\} = E_t^{-}\{(1 - \frac{\sigma_t^{-}}{\sigma_t^{-}+\Sigma_\eta})\mu_t^{-} + (\frac{\sigma_t^{-}}{\sigma_t^{-}+\Sigma_\eta})\eta_t\}$
$= E^{-}\{(1 - \frac{\sigma_t^{-}}{\sigma_t^{-}+\Sigma_\eta})\mu_t^{-} + \frac{\sigma_t^{-}}{\sigma_t^{-}+\Sigma_\eta}(f_{t+1} + e_\eta)\}$
$= (1 - \frac{\sigma_t^{-}}{\sigma_t^{-}+\Sigma_\eta})\mu_t^{-} + (\frac{\sigma_t^{-}}{\sigma_t^{-}+\Sigma_\eta})\mu_t^{-}$
$= \mu_t^{-}$

Now take the expectation of the Kalman filter formula for updating the variance.
$E^{-}\{\Sigma^{+}\} = E^{-}\{\frac{\sigma_t^{-}\Sigma_\eta}{\sigma_t^{-}+\Sigma_\eta} + \Sigma_\eta\}$
$= \frac{\sigma_t^{-}\Sigma_\eta}{\sigma_t^{-}+\Sigma_\eta} + \Sigma_\eta$
$= \Sigma_t^{+}$
The first equality is by substitution and the second is by the fact that there is nothing random in the formula for updating the uncertainty or the underlying variance. Seeing as the agent knows the variance of the signal shocks, Σ_η, there is actually no random component in Σ^{+}.

$$V(W_t, \mu_t^{-}, \Sigma_t^{-}) = max_{\{q,\Sigma^{+}\}}|| - \rho(rW_t + q_t(\mu_t^{-} - p_t r)) + \frac{1}{2}\rho^2 q_t'(\Sigma_t^{+})q_t$$
$$+\beta E_t^{-}\{V(W_{t+1}, \mu_{t+1}^{-}, \Sigma_{t+1}^{-})|\mu_t^{-}, \Sigma_t^{-}\}.$$

It may seem peculiar that, after taking the prior expectations, the Bellman equation does not depend on the prior variance but instead the posterior. The reason is that the formula for updating the agent's uncertainty based on the number of signals gathered is $\sigma_t^{+} = \frac{\sigma_t^{-}\Sigma_\eta}{\Sigma_\eta + T\sigma_t^{-}}$. Hence, it is possible to see that if no signals are gathered, $T = 0$, then the formula will collapse to the prior variance.

Finding the optimal portfolio

Take the first order conditions on q, that is q in the current period.
$FOC_q: \quad V_q(W_t, \mu_t^-, \Sigma_t^-) = -\rho(\mu_t^- - p_t r) + q_t \rho^2 \Sigma_t^+$
$+ \beta E_t^- \{V_q(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^-\} = 0$

The Benveniste-Scheinkmann condition can be used to find the unknown form of $V_q(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)$. Take the first order conditions on the state variable W.
$FOC_W: \quad V_W(W_t, \mu_t^-, \Sigma_t^-) = -\rho r$

In cases where the utility function in the initial period were nonlinear in the state variable, the Benveniste-Scheinkmann condition would depend on some variable and not just parameters, which is the situation here. The next step would be to move the condition forwards by one period so that it could be substituted into the FOC. However, there is no need for this here as both ρ and r are constants hence the Euler equation can be found directly.
$V_q(W_t, \mu_t^-, \Sigma_t^-) = -\rho(\mu_t^- - p_t r) + q_t \rho^2 \Sigma_t^+ + \beta E_t^- \{-\rho r\} = 0$

This simply becomes
$V_q(W_t, \mu_t^-, \Sigma_t^-) = -\rho(\mu_t^- - p_t r) + q_t \rho^2 \Sigma_t^+ - \beta \rho r = 0.$

The Euler equation can now be solved for the optimal portfolio, q^*.
$q_t \rho^2 \Sigma_t^+ = \rho(\mu_t^- - p_t r) + \beta \rho r.$
$q^* = \frac{\rho(\mu_t^- - p_t r) + \beta \rho r}{\rho^2 \Sigma_t^+}.$
This can be separated into two terms as follows.
$q^* = \frac{\mu_t^- - p_t r}{\rho \Sigma_t^+} + \frac{\beta r}{\rho \Sigma_t^+}$

To recover q^{div}, the hedging demand must be zero. That requires the following to hold
$\frac{\beta r}{\rho \Sigma_t^-} = 0.$

Hence $\beta r = 0.$

This can be broken down into two further conditions:
$\beta = 0$ and $r = 0$.

It is also necessary for the agent to gather no signals, $T = 0$. Then $\Sigma_t^+ = \Sigma_t^-$ (because $\sigma_t^+ = \sigma_t^-$ will be true) will be true and the relevant myopic benchmark solution will be found.

Finding the optimal learning allocation

As the steady state portfolio, q^*, has now been found, it can be used to derive the optimal allocation of the learning constraint from the Bellman equation, the decision rule. The aim is to find Σ^+, the posterior variances for each asset in the current period. That is, the variance of the investor's beliefs after he allocates his signals between the assets. So substitute q^* into the Bellman equation.

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q,\Sigma^+\}}|| - \rho(rW_t + q(\mu_t^- - p_t r)) + \frac{1}{2}(\rho q)^2(\Sigma_t^+) \\ + \beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}.$$

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q,\Sigma^+\}}|| - \rho(rW_t + (\frac{\mu_t^- - p_t r}{\rho\Sigma_t^+} + \frac{\beta r}{\rho\Sigma_t^+})(\mu_t^- - p_t r)) \\ + \frac{1}{2}(\rho[\frac{\mu_t^- - p_t r}{\rho\Sigma_t^+} + \frac{\beta r}{\rho\Sigma_t^+}])^2\Sigma_t^+ + \beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}.$$

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q,\Sigma^+\}}|| - \rho rW_t - (\frac{\mu_t^- - p_t r + \beta r}{\Sigma_t^+})(\mu_t^- - p_t r) \\ + \frac{1}{2}(\frac{\mu_t^- - p_t r + \beta r}{\Sigma_t^+})^2\Sigma_t^+ + \beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}.$$

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q,\Sigma^+\}}|| - \rho r W_t - (\mu_t^- - p_t r + \beta r)(\mu_t^- - p_t r)(\Sigma_t^+)^{-1}$$
$$+ \frac{1}{2}(\mu_t^- - p_t r + \beta r)^2(\Sigma_t^+)^{-1} + \beta E_t^- \{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}.$$

Now the first order conditions on the precision of beliefs, $(\Sigma_t^+)^{-1}$, can be taken.
$FOC_{(\Sigma_t^+)^{-1}}: \quad V_{(\Sigma_t^+)^{-1}}(W_t, \mu_t^-, \Sigma_t^-) = -(\mu_t^- - p_t r + \beta r)(\mu_t^- - p_t r)$
$+ \frac{1}{2}(\mu_t^- - p_t r + \beta r)^2 + \beta E_t^- \{V_{\Sigma^{+1}}(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\} = 0$

The first order condition does not depend on $(\Sigma_t^+)^{-1}$ which illustrates the linear nature of the dependency of the Bellman equation on the precision of the investor's beliefs.

Now the first order conditions on the variance, Σ_t^+, can be taken.
$FOC_{\Sigma_t^+}: \quad V_{\Sigma_t^+}(W_t, \mu_t, \Sigma_t) = (\mu_t^- - pr + \beta r)(\mu_t^- - pr)(\Sigma_t^+)^{-2}$
$- \frac{1}{2}(\mu_t^- - pr + \beta r)^2(\Sigma_t^+)^{-2} + \beta E\{V_{\Sigma^{+1}}(W_{t+1}, \mu_{t+1}, \Sigma_{t+1})|\mu_t, \Sigma_t\} = 0$

The Benveniste-Scheinkmann condition can be used to find the unknown form of $V_{\Sigma_t^+}(W_{t+1}, \mu_{t+1}, \Sigma_{t+1})$. Take the first order conditions on the state variable W.
$FOC_{\Sigma_t^+}: \quad V_{\Sigma_t^+}(W_t, \mu_t, \Sigma_t) = -\rho r$

The Euler equation is:
$(\mu_t^- - pr + \beta r)(\mu_t^- - pr)(\Sigma_t^+)^{-2}$
$- \frac{1}{2}(\mu_t^- - pr + \beta r)^2(\Sigma_t^+)^{-2} + \beta E\{-\rho r\} = 0$

There is nothing random in the Benveniste-Scheinkmann condition.
$(\mu_t^- - pr + \beta r)(\mu_t^- - pr)(\Sigma_t^+)^{-2} - \frac{1}{2}(\mu_t^- - pr + \beta r)^2(\Sigma_t^+)^{-2} = \beta \rho r$

Multiply both sides by the posterior variance squared.
$(\mu_t^- - pr + \beta r)(\mu_t^- - pr) - \frac{1}{2}(\mu_t^- - pr + \beta r)^2 = \beta \rho r(\Sigma_t^+)^2$

Solve for the posterior variance.

$$(\Sigma_t^+)^2 = \frac{(\mu_t^- - pr + \beta r)(\mu_t^- - pr) - \frac{1}{2}(\mu_t^- - pr + \beta r)^2}{\beta \rho r}$$

$$\Sigma_t^+ = \sqrt{\frac{(\mu_t^- - pr + \beta r)(\mu_t^- - pr) - \frac{1}{2}(\mu_t^- - pr + \beta r)^2}{\beta \rho r}}$$

To find the optimal posterior uncertainty, simply subtract the true variance from both sides.

$$\sigma_t^+ + \Sigma_\eta = \sqrt{\frac{(\mu_t^- - pr + \beta r)(\mu_t^- - pr) - \frac{1}{2}(\mu_t^- - pr + \beta r)^2}{\beta \rho r}}$$

$$\sigma_t^+ = \sqrt{\frac{(\mu_t^- - pr + \beta r)(\mu_t^- - pr) - \frac{1}{2}(\mu_t^- - pr + \beta r)^2}{\beta \rho r}} - \Sigma_\eta$$

It would be necessary to bound σ_t^+ below at 0 because the lower bound of its first term is zero so it could fall as low as $-\Sigma_\eta$ without being bounded below. It does not make sense for the uncertainty to be less than zero.

Visualising the optimal posterior variance

A graphical representation of the posterior variance is below. It is a numerical example with the excess return as a variable and parameters set to be: $\beta = 0.99$, $\rho = 4$ and $r = 1$. The posterior uncertainty is the blue curve and the posterior precision the red. As the excess return tends to zero, the investor must be compensated with a higher posterior precision or a lower posterior variance.

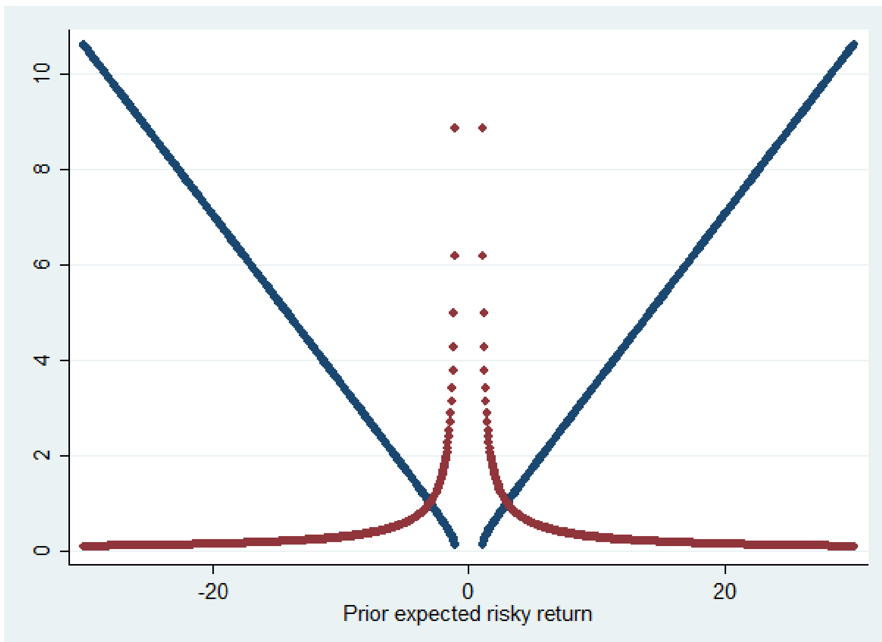

Figure A.1

Working on the Bellman equation under Rational Expectations

The Bellman Equation is:

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q\}}||E_t^-\{ln(- - E_t^+\{exp(-\rho(rW_t + q_t(f_{t+1} - p_t r)))|\mu_t^+, \Sigma_t^+\}) \\ |\mu_t^-, \Sigma_t^-\} + \beta E_t^-\{V_t(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$$

As the distribution of $-\rho W_{t+1}$ is normal, the MGF can be used for taking the inner expectation. So, under Rational Expectations the Bellman Equation becomes

$$V(W_t, \bar{\mu}, \Sigma_\eta) = max_{\{q,\Sigma^+\}}||E_t^-\{ln(exp(-\rho(rW_t + q_t(\bar{\mu} - \bar{p}r)) \\ + \frac{1}{2}\rho^2 q_t' \Sigma_\eta q_t))|\bar{\mu}, \Sigma_\eta\} + \beta E_t^-\{V(W_{t+1}, \bar{\mu}, \Sigma_\eta)|\bar{\mu}, \Sigma_\eta\}.$$

After some simplification, the result becomes

$$V(W_t, \bar{\mu}, \Sigma_\eta) = max_{\{q,\Sigma^+\}}||E_t^-\{-\rho(rW_t + q_t(\bar{\mu} - \bar{p}r)) + \frac{1}{2}\rho^2 q_t' \Sigma_\eta q_t)|\bar{\mu}, \Sigma_\eta\} \\ + \beta E_t^-\{V(W_{t+1}, \bar{\mu}, \Sigma_\eta)|\bar{\mu}, \Sigma_\eta\}.$$

The outer expectation, E^-, over posterior expectations, is rather redundant in this case when the agent has Rational Expectations so it can be removed.

$$V(W_t, \bar{\mu}, \Sigma_\eta) = max_{\{q\}}|| - \rho(rW_t + q_t(\bar{\mu} - \bar{p}r)) + \frac{1}{2}\rho^2 q_t' \Sigma_\eta q_t \\ + \beta E_t^-\{V(W_{t+1}, \bar{\mu}, \Sigma_\eta)|\bar{\mu}, \Sigma_\eta\}.$$

There is now an explicit expression for the Bellman equation under Rational Expectations. Notice that there is no longer maximisation in terms of the posterior variance because learning is irrelevant to the agent with Rational Expectations. The learner is “poorer” than the Rational Expectations agent in terms of learning.

Finding the optimal portfolio under Rational Expectations

Take the first order conditions on q, that is q in the current period.

$FOC_q: \quad V_q(W_t, \bar{\mu}, \Sigma_\eta) = -\rho(\bar{\mu} - \bar{p}r) + q_t\rho^2\Sigma_\eta + \beta E_t^-\{V_q(W_{t+1}, \bar{\mu}, \Sigma_\eta)|\bar{\mu}, \Sigma_\eta\} = 0$

The Benveniste-Scheinkmann condition can be used to find the unknown form of $V_q(W_{t+1}, \bar{\mu}, \Sigma_\eta)$. Take the first order conditions on the state variable W.

$FOC_W: \quad V_W(W_t, \bar{\mu}, \Sigma_\eta) = -\rho r$

Substitute the Benveniste-Scheinkmann condition into the FOC.

$V_q(W_t, \bar{\mu}, \Sigma_\eta) = -\rho(\bar{\mu} - \bar{p}r) + q_t\rho^2\Sigma_\eta + \beta E_t^-\{-\rho r\} = 0$

This simply becomes

$V_q(W_t, \bar{\mu}, \Sigma_\eta) = -\rho(\bar{\mu} - \bar{p}r) + q_t\rho^2\Sigma_\eta - \beta\rho r = 0.$

The Euler equation can now be solved for the optimal portfolio, q^*, under Rational Expectations.

$q_t\rho^2\Sigma_\eta = \rho(\bar{\mu} - \bar{p}r) + \beta\rho r.$

$q^* = \frac{\rho(\bar{\mu}-\bar{p}r)+\beta\rho r}{\rho^2\Sigma_\eta}.$

This can be separated into two terms as follows.

$q^* = \frac{\bar{\mu}-\bar{p}r}{\rho\Sigma_\eta} + \frac{\beta r}{\rho\Sigma_\eta}$

In matrix terms, for multiple risky assets, it is as follows.

$q^* = (\bar{\mu} - \bar{p}r)(\rho\Sigma_\eta)^{-1} + \beta r(\rho\Sigma_\eta)^{-1}$

A.1.3 When learning is endogenous: indifference to the time of resolution

This part of the appendix derives the Bellman equation when the agent is indifferent to the time of resolution of uncertainty and has CARA preferences over wealth. The steady state solution is also derived here.

Finding the Bellman equation

The outer and inner utility functions are given by the following two expressions:
$u^-(x) = y$ and $u^+(W) = -exp(-\rho W)$.

Using these, the Bellman equation is the following.

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{W_{t+1}\}}||E_t^-\{u^-(E_t^+\{u^+(W_{t+1})|\mu_t^+, \Sigma_t^+\})|\mu_t^-, \Sigma_t^-\}$$
$$+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$$

This expression is in terms of expectations that, so far, are not explicit. The utility functions are also not explicit. Substitute the outer and inner utility functions into the Bellman equation.

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{W_{t+1}\}}||E_t^-\{E_t^+\{-exp(-\rho W_{t+1})|\mu_t^+, \Sigma_t^+\}|\mu_t^-, \Sigma_t^-\}$$
$$+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$$

In order to take the expectations next, substitute the expression for wealth into the Bellman equation.

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t\}}||E_t^-\{E_t^+\{-exp(-\rho(rW_t + q_t'(f_{t+1} - p_t r)))|\mu_t^+, \Sigma_t^+\}|\mu_t^-, \Sigma_t^-\}$$
$$+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}$$

Take the first expectation, in terms of the posterior beliefs, using the mean of a lognormal distribution.

$$
\begin{aligned}
&V(W_t, \mu_t^-, \Sigma_t^-) = \\
&max_{\{q_t,(\Sigma_t^+)^{-1}\}}||E_t^-\{-exp(-\rho(rW_t + q_t'(\mu_t^+ - p_t r)) + \frac{1}{2}\rho^2 q_t'\Sigma_t^+ q_t)|\mu_t^-, \Sigma_t^-\} \\
&+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}
\end{aligned}
$$

Take the second expectation, in terms of the prior beliefs, using the mean of a log-normal distribution.

$$
\begin{aligned}
V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t,(\Sigma_t^+)^{-1}\}}|| - exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)+) + \frac{1}{2}\rho^2 q_t'\Sigma_t^+ q_t \\
+\frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\
+\beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\}
\end{aligned}
$$

The prior expectation of the posterior mean is the prior mean. The signal is the true underlying parameter plus some noise. The agent expects, in prior terms, the first to be the prior mean and the second to be zero:

$$
\begin{aligned}
&E_t^-[\mu_t^+] = E_t^-[\mu_t^- + K_t(\eta_t - \mu_t^-)] \\
&\quad = E_t^-[\mu_t^- + K_t(\eta_t - \mu_t^-)] \\
&\quad = E_t^-[\mu_t^- + K_t(\theta + e_{\eta,t} - \mu_t^-)] \\
&\quad = \mu_t^- + K_t(\mu_t^- - \mu_t^-) \\
&\quad = \mu_t^-.
\end{aligned}
$$

The prior expectation of the posterior variance is the posterior variance as there is nothing random in the formula for updating the variance. However, it collapses to the prior variance if no signals are allocated to the asset in question.

$$
\begin{aligned}
&E_t^-[\Sigma_{ii,t}^+] = E_t^-[\tfrac{\sigma_t^- \Sigma_\eta}{T\sigma_t^- + \Sigma_\eta} + \Sigma_\eta] \\
&\quad = \tfrac{\sigma_t^- \Sigma_\eta}{T\sigma_t^- + \Sigma_\eta}.
\end{aligned}
$$

If the agent sets $T = 0$, that is, he gathers no signals about asset ii, then the posterior variance is the following.

$E_t^-[\Sigma_{ii,t}^+] = \frac{\sigma_t^- \Sigma_\eta}{\Sigma_\eta} + \Sigma_\eta$
$\quad = \Sigma_t^-$

Taking the variance of the argument of the exponential in the penultimate expression gives the following.

$V_t^-[-\rho(rW_t + q_t'(\mu_t^+ - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t] = V_t^-[-\rho q_t' \mu_t^+]$
$\quad = \rho^2 q_t' V_t^-[\mu_t^+] q_t$
$\quad = \rho^2 q_t' V_t^-[\mu_t^- + K_t(\eta_t - \mu_t^-)] q_t$
$\quad = \rho^2 q_t' V_t^-[K_t \eta_t] q_t$
$\quad = \rho^2 q_t' K_t V_t^-[\eta_t] K_t' q_t$
$\quad = \rho^2 q_t' K_t \Sigma_\eta K_t' q_t$

The steady state solution

The Bellman equation is the following.

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t, (\Sigma_t^+)^{-1}\}} || - exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t \\ + \frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\ + \beta E_t^- \{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^-\} \end{aligned}$$

Take the first order condition in terms of the quantity of risky assets purchased.

$$\begin{aligned} V_{q_t}(W_t, \mu_t^-, \Sigma_t^-) = -(-\rho(\mu_t^- - p_t r) + \rho^2 q_t' \Sigma_t^+ \\ + \rho^2 q_t' K_t \Sigma_t^- K_t') exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t \\ + \frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\ + \beta E_t^- \{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^-\} = 0 \end{aligned}$$

Take the first order condition in terms of the quantity of risky assets purchased when the problem is 2D; scalars and vectors instead of vectors and matrices.

$$\begin{aligned} V_{q_t^1}(W_t,\mu_t^-,\Sigma_t^-) &= -(-\rho(\mu_t^{-,1}-p_t^1 r)+\rho^2 q_t'^1\Sigma_t^{+,1}+\rho^2 q_t'(K_t^1)^2\Sigma_t^{-,1}) \\ exp(-\rho(rW_t+q_t'^1(\mu_t^{-,1}-p_t^1 r)+q_t'^2(\mu_t^{-,2}-p_t^2 r)) &+ \frac{1}{2}\rho^2(q_t'^1)^2\Sigma_t^{+,1}+\frac{1}{2}\rho^2(q_t'^2)^2\Sigma_t^{+,2} \\ &+\frac{1}{2}\rho^2((K_t^1)^2\Sigma_t^{-,1}(q_t^1)^2+(K_t^2)^2\Sigma_t^{-,2}(q_t^2)^2) \\ &+\beta E_t^-\{V(W_{t+1},\mu_{t+1}^-,\Sigma_{t+1}^-)|\mu_t^-,\Sigma_t^-\}=0 \end{aligned}$$

Take the first order condition in terms of the posterior variance.

$$\begin{aligned} V_{\Sigma_t^{+,1}}(W_t,\mu_t^-,\Sigma_t^-) &= -(\frac{1}{2}\rho^2(q_t'^1)^2+\frac{1}{2}\rho^2(K_t^1)^2(q_t^1)^2) \\ exp(-\rho(rW_t+q_t'^1(\mu_t^{-,1}-p_t^1 r)+q_t'^2(\mu_t^{-,2}-p_t^2 r)) &+ \frac{1}{2}\rho^2(q_t'^1)^2\Sigma_t^{+,1}+\frac{1}{2}\rho^2(q_t'^2)^2\Sigma_t^{+,2} \\ &+\frac{1}{2}\rho^2((K_t^1)^2\Sigma_t^{-,1}(q_t^1)^2+(K_t^2)^2\Sigma_t^{-,2}(q_t^2)^2) \\ &+\beta E_t^-\{V(W_{t+1},\mu_{t+1}^-,\Sigma_{t+1}^-)|\mu_t^-,\Sigma_t^-\}=0 \end{aligned}$$

Find the Benveniste-Scheinkmann condition by firstly taking the first order condition in terms of the state variables; wealth today.

$$\begin{aligned} V_{W_t}(W_t,\mu_t^-,\Sigma_t^-) = -(-\rho r)exp(-\rho(rW_t+q_t'(\mu_t^- - p_t r)) &+ \frac{1}{2}\rho^2 q_t'\Sigma_t^+ q_t \\ &+\frac{1}{2}\rho^2 q_t' K_t\Sigma_t^- K_t' q_t) \end{aligned}$$

Move the Benveniste-Scheinkmann condition forwards by one period.

$$\begin{aligned} V_{W_{t+1}}(W_{t+1},\mu_{t+1}^-,\Sigma_{t+1}^-) &= -(-\rho r)exp(-\rho(rW_{t+1}+q_{t+1}'(\mu_{t+1}^- - p_{t+1} r)) \\ &+\frac{1}{2}\rho^2 q_{t+1}'\Sigma_{t+1}^+ q_{t+1}+\frac{1}{2}\rho^2 q_{t+1}' K_{t+1}\Sigma_{t+1}^- K_{t+1}' q_{t+1}) \end{aligned}$$

Substitute the Benveniste-Scheinkman condition into the first order condition in terms of q_t.

$$
\begin{aligned}
V_{q_t}(W_t, \mu_t^-, \Sigma_t^-) = -(-\rho(\mu_t^- - p_t r) + \rho^2 q_t' \Sigma_t^+ \\
+\rho^2 q_t' K_t \Sigma_t^- K_t') exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t \\
+\frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\
+\beta E_t^- \{(\rho r) exp(-\rho(rW_{t+1} + q_{t+1}'(\mu_{t+1}^- - p_{t+1} r)) + \frac{1}{2}\rho^2 q_{t+1}' \Sigma_{t+1}^+ q_{t+1} \\
+\frac{1}{2}\rho^2 q_{t+1}' K_{t+1} \Sigma_{t+1}^- K_{t+1}' q_{t+1}) | \mu_t^-, \Sigma_t^- \} = 0
\end{aligned}
$$

Simplify.

$$
\begin{aligned}
(-\rho(\mu_t^- - p_t r) + \rho^2 q_t' \Sigma_t^+ \\
+\rho^2 q_t' K_t \Sigma_t^- K_t') exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t \\
+\frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\
= \beta E_t^- \{(\rho r) exp(-\rho(rW_{t+1} + q_{t+1}'(\mu_{t+1}^- - p_{t+1} r)) + \frac{1}{2}\rho^2 q_{t+1}' \Sigma_{t+1}^+ q_{t+1} \\
+\frac{1}{2}\rho^2 q_{t+1}' K_{t+1} \Sigma_{t+1}^- K_{t+1}' q_{t+1}) | \mu_t^-, \Sigma_t^- \}
\end{aligned}
$$

Impose the steady state. Firstly, $\sigma_t^- = 0$, which means that $\Sigma_t^- = \Sigma_\eta$ in the steady state because the agent knows everything about the risky assets, particularly the true variance. Additionally, $K_t \to 0$ in the limit so set $K_t = 0$. Next, note that the agent does not expect the price to vary. Also, set the steady state wealth to $\bar{W}$, quantity to $\bar{q}$ and mean belief to $\bar{\mu}$.

$$
\begin{aligned}
(-\rho(\bar{\mu} - \bar{p} r) + \rho^2 \bar{q}' \Sigma_\eta) \\
exp(-\rho(r\bar{W} + \bar{q}'(\bar{\mu} - \bar{p} r)) + \frac{1}{2}\rho^2 \bar{q}' \Sigma_\eta \bar{q}) \\
= \beta E_t^- \{(\rho r) exp(-\rho(r\bar{W} + \bar{q}'(\bar{\mu} - \bar{p} r)) + \frac{1}{2}\rho^2 \bar{q}' \Sigma_\eta \bar{q}) | \bar{\mu}, \Sigma_\eta \}
\end{aligned}
$$

As nothing is stochastic any longer, the expectation becomes redundant.

$$\begin{aligned}
&(-\rho(\bar{\mu} - \bar{p}r) + \rho^2 \bar{q}'\Sigma_\eta) \\
&exp(-\rho(r\bar{W} + \bar{q}'(\bar{\mu} - \bar{p}r)) + \frac{1}{2}\rho^2 \bar{q}'\Sigma_\eta \bar{q}) \\
&= \beta(\rho r) exp(-\rho(r\bar{W} + \bar{q}'(\bar{\mu} - \bar{p}r)) + \frac{1}{2}\rho^2 \bar{q}'\Sigma_\eta \bar{q})
\end{aligned}$$

The exponential terms are the same and so can be cancelled.

$$-\rho(\bar{\mu} - \bar{p}r) + \rho^2 \bar{q}'\Sigma_\eta = \beta\rho r$$

Solve for the steady state portfolio choice, $\bar{q}$.

$$\rho\bar{q}'\Sigma_\eta = \bar{\mu} - \bar{p}r + \beta r$$

The solution with one risky asset is:

$$\bar{q} = \frac{\bar{\mu} - \bar{p}r}{\rho\Sigma_\eta} + \frac{\beta r}{\rho\Sigma_\eta}.$$

The solution with multiple risky assets is:

$$\bar{q} = (\bar{\mu} - \bar{p}r)(\rho\Sigma_\eta)^{-1} + \beta r(\rho\Sigma_\eta)^{-1}.$$

Is the Bellman equation concave or convex in the posterior variance?

The Bellman equation is the following.

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t,(\Sigma_t^+)^{-1}\}} || - exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t \\ + \frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\ + \beta E_t^- \{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^-\} \end{aligned}$$

Take the first order condition in terms of the posterior variance.

$$\begin{aligned} V_{\Sigma_t^+}(W_t, \mu_t^-, \Sigma_t^-) = -(\frac{1}{2}\rho^2 q_t' q_t) \\ exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t + \frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\ + \beta E_t^- \{V_{\Sigma_t^+}(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^-\} = 0 \end{aligned}$$

Take the second order condition in terms of the posterior variance.

$$\begin{aligned} V_{(\Sigma_t^+)^2}(W_t, \mu_t^-, \Sigma_t^-) = -(\frac{1}{2}\rho^2 q_t' q_t)^2 \\ exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t + \frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \\ + \beta E_t^- \{V_{(\Sigma_t^+)^2}(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-) | \mu_t^-, \Sigma_t^-\} = 0 \end{aligned}$$

Find the Benveniste-Scheinkmann condition by firstly taking the first order condition in terms of the state variables; wealth today.

$$\begin{aligned} V_{W_t}(W_t, \mu_t^-, \Sigma_t^-) = -(-\rho r) exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t \\ + \frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) \end{aligned}$$

Then find the second order condition in terms of the state variables; wealth today.

$$V_{(W_t)^2}(W_t,\mu_t^-,\Sigma_t^-) = -(\rho r)^2 exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t'\Sigma_t^+ q_t$$
$$+\frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t)$$

Move the Benveniste-Scheinkmann condition forwards by one period.

$$V_{W_{t+1}}(W_{t+1},\mu_{t+1}^-,\Sigma_{t+1}^-) = -(\rho r)^2 exp(-\rho(rW_{t+1} + q_{t+1}'(\mu_{t+1}^- - p_{t+1} r))$$
$$+\frac{1}{2}\rho^2 q_{t+1}'\Sigma_{t+1}^+ q_{t+1} + \frac{1}{2}\rho^2 q_{t+1}' K_{t+1}\Sigma_{t+1}^- K_{t+1}' q_{t+1})$$

Substitute the Benveniste-Scheinkmann condition into the first order condition.

$$V_{\Sigma_t^+}(W_t,\mu_t^-,\Sigma_t^-) = -(\frac{1}{2}\rho^2 q_t' q_t)$$
$$exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t'\Sigma_t^+ q_t$$
$$+\frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t)$$
$$-\beta(\rho r)^2 E_t^-\{exp(-\rho(rW_{t+1} + q_{t+1}'(\mu_{t+1}^- - p_{t+1} r)) + \frac{1}{2}\rho^2 q_{t+1}'\Sigma_{t+1}^+ q_{t+1}$$
$$+\frac{1}{2}\rho^2 q_{t+1}' K_{t+1}\Sigma_{t+1}^- K_{t+1}' q_{t+1})|\mu_t^-,\Sigma_t^-\} = 0$$

Take the expectation of the Benveniste-Scheinkmann condition.

$$E_t^-\{exp(-\rho(rW_{t+1} + q_{t+1}'(\mu_{t+1}^- - p_{t+1} r)) + \frac{1}{2}\rho^2 q_{t+1}'\Sigma_{t+1}^+ q_{t+1}$$
$$+\frac{1}{2}\rho^2 q_{t+1}' K_{t+1}\Sigma_{t+1}^- K_{t+1}' q_{t+1})|\mu_t^-,\Sigma_t^-\}$$

The prior expectation of the posterior mean is the prior mean. The prior expectation of the posterior variance is the posterior variance as there is nothing random in the formula for updating the variance. The prior mean tomorrow is the Kalman filter formula for updating the posterior mean today with the observed risky return vector, f_{t+1}.

$E_t^-[-\rho(rW_{t+1} + q'_{t+1}(\mu_{t+1}^- - p_{t+1}r)) + \frac{1}{2}\rho^2 q'_{t+1}\Sigma_{t+1}^+ q_{t+1} + \frac{1}{2}\rho^2 q'_{t+1}K_{t+1}\Sigma_{t+1}^- K'_{t+1}q_{t+1}]$
$= -\rho(rE_t^-[rW_t + q_t(f_{t+1} - p_t r)] + q'_{t+1}(E_t^-[\mu_{t+1}^-] - p_{t+1}r)) + \frac{1}{2}\rho^2 q'_{t+1}\Sigma_{t+1}^+ q_{t+1} + \frac{1}{2}\rho^2 q'_{t+1}K_{t+1}\Sigma_{t+1}^- K'_{t+1}q_{t+1}$
$= -\rho(r(rW_t + q_t(E_t^-[f_{t+1}] - p_t r)) + q'_{t+1}(E_t^-[\mu_t^+ + K_t(f_{t+1} - \mu_t^+)] - p_{t+1}r)) + \frac{1}{2}\rho^2 q'_{t+1}\Sigma_{t+1}^+ q_{t+1} + \frac{1}{2}\rho^2 q'_{t+1}K_{t+1}\Sigma_{t+1}^- K'_{t+1}q_{t+1}$
$= -\rho(r(rW_t + q_t(\mu_t^- - p_t r)) + q'_{t+1}(\mu_t^- + K_t(\mu_t^- - \mu_t^-) - p_{t+1}r)) + \frac{1}{2}\rho^2 q'_{t+1}\Sigma_{t+1}^+ q_{t+1} + \frac{1}{2}\rho^2 q'_{t+1}K_{t+1}\Sigma_{t+1}^- K'_{t+1}q_{t+1}$
$= -\rho(r(rW_t + q_t(\mu_t^- - p_t r)) + q'_{t+1}(\mu_t^- - p_{t+1}r)) + \frac{1}{2}\rho^2 q'_{t+1}\Sigma_{t+1}^+ q_{t+1} + \frac{1}{2}\rho^2 q'_{t+1}K_{t+1}\Sigma_{t+1}^- K'_{t+1}q_{t+1}$

The variance of the signal shock, Σ_η, is known whilst the variance of tomorrow's return, f_{t+1}, is the prior variance.

$V_t^-[-\rho(rW_{t+1} + q'_{t+1}(\mu_{t+1}^- - p_{t+1}r)) + \frac{1}{2}\rho^2 q'_{t+1}\Sigma_{t+1}^+ q_{t+1} + \frac{1}{2}\rho^2 q'_{t+1}K_{t+1}\Sigma_{t+1}^- K'_{t+1}q_{t+1}]$
$= (\rho r)^2 V_t^-[rW_t + q_t(f_{t+1} - p_t r)] + (\rho q'_{t+1})^2 V_t^-[\mu_{t+1}^-]$
$= (\rho r q_t)^2 V_t^-[f_{t+1}] + (\rho q'_{t+1})^2 V_t^-[\mu_t^+ + K_t(f_{t+1} - \mu_t^+)]$
$= (\rho r q_t)^2 V_t^-[f_{t+1}] + (\rho q'_{t+1})^2 V_t^-[\mu_t^- + K_t(\eta_t - \mu_t^-) + K_t(f_{t+1} - \mu_t^- + K_t(\eta_t - \mu_t^-))]$
$= (\rho r q_t)^2 \Sigma_t^- + (\rho q'_{t+1})^2 (K_t\Sigma_\eta K'_t + K_t\Sigma_t^- K'_t + K_t K_t \Sigma_\eta K'_t K'_t)$

So the expectation of the Benveniste-Scheinkmann condition is as follows.

$$exp(-\rho(r(rW_t + q_t(\mu_t^- - p_t r)) + q'_{t+1}(\mu_t^- - p_{t+1}r)) + \frac{1}{2}\rho^2 q'_{t+1}\Sigma_{t+1}^+ q_{t+1}$$
$$+\frac{1}{2}\rho^2 q'_{t+1}K_{t+1}\Sigma_{t+1}^- K'_{t+1}q_{t+1} + \frac{1}{2}[(\rho r q_t)^2\Sigma_t^- + (\rho q'_{t+1})^2(K_t\Sigma_\eta K'_t + K_t\Sigma_t^- K'_t + K_t K_t\Sigma_\eta K'_t K'_t)])$$

Thus the SOC is as below.

$$V_{\Sigma_t^+}(W_t, \mu_t^-, \Sigma_t^-) = -(\frac{1}{2}\rho^2 q_t' q_t)^2 exp\{-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t$$
$$+\frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t\} - \beta(\rho r)^2 exp\{-\rho(r(rW_t + q_t(\mu_t^- - p_t r)) + q_{t+1}'(\mu_t^- - p_{t+1}r))$$
$$+\frac{1}{2}\rho^2 q_{t+1}' \Sigma_{t+1}^+ q_{t+1} + \frac{1}{2}\rho^2 q_{t+1}' K_{t+1} \Sigma_{t+1}^- K_{t+1}' q_{t+1}$$
$$+\frac{1}{2}[(\rho r q_t)^2 \Sigma_t^- + (\rho q_{t+1}')^2 (K_t \Sigma_\eta K_t' + K_t \Sigma_t^- K_t' + K_t K_t \Sigma_\eta K_t' K_t')]\}$$

The bracketed coefficients of both exponential terms above are squared and hence positive and beta is also positive. Therefore, both terms are negative, given that both are multiplied by minus one. This means that the second order condition is always concave.

Numerical solution to the linear-CARA Bellman equation

The numerical results to the Bellman equation of the endogenous learning technology with CARA preferences are presented below. They describe the agent's optimal actions over the four choice variables: the quantity of assets held for each asset and the quantity of signals acquired for both. Due to high computational costs, the numerical results are based on very limited grids for the state variables. The parameters and grids used are detailed next.

Parameters		
r	0.5	risk free rate
ρ	1	risk aversion coefficient
θ	1.5	true: mean
Σ_η	0.9	true variance
p_i	2	risky asset price
β	0.99	discount rate

State variable grids		
W	2	wealth
μ_i	0.00, 0.25, 0.50, 0.75	belief: mean
	1.00, 1.25, 1.50, 1.75, 2.00	
σ_i	1.000, 1.375, 1.750, 2.125	belief: uncertainty
	2.500, 2.875, 3.250, 3.625, 4.000	

Actually, the two assets are identical. This is not problematic because it is the beliefs that the agent holds that matter about the underlying values rather than those values themselves. Notice that the excess return on average is positive as $\theta > p_i$. Wealth only takes two different values but only two are necessary for showing whether or not wealth matters. Changing the values of other key parameters such as risk aversion and prices is beyond the scope of this project. Additionally, the quantity of risky asset purchased, q_i, takes 51 evenly spaced values between -2 and 2. The last key item to report from the numerical calculation is the transition matrix. It describes the probability of transiting from one element of the μ_i grid to another element (becoming more or less optimistic), dependent on the signal received tomorrow. It does not represent the probability of the belief about the variance changing because that is deterministic not stochastic. The probability of transiting to a given state does not depend on the state today (signals received today are independent of those yesterday) and hence each row is identical. The probability of some mean belief occurring is calculated as the normal probability density given the true mean $\&$ variance scaled by the sum of all normal probability densities for all possible mean beliefs.

Transition matrix								
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315
0.0383	0.0585	0.0828	0.1085	0.1315	0.1477	0.1535	0.1477	0.1315

The choice of assets held is allowed to be negative and hence the agent is permitted to short sell assets. In this, he is constrained; the absolute value of the assets purchased may not exceed the wealth he inherited from yesterday. This constraint has two purposes. Firstly, it stops the agent from borrowing more, on the risky assets, than he holds in wealth, this is a kind of debt constraint. Secondly, it bounds the wealth grid so that the state variables are bounded and hence the Bellman equation implemented is a contraction mapping. The agent cannot borrow at the risk free rate. The following two-part table shows the reader the wealth endowment and the variance beliefs on which the agent's choices over the decision variables depend for all possible states of the world.

Now the numerical approximate decision rules for the choice variables are presented. The state of the world, the state variable endowment that the agent wakes up today with, is listed above each set of decision variable results. It corresponds to table A.1 above, which list the values of the state variables in those states. Each column then gives the optimal choice that the agent should make in each state of the world: the first two columns show the amount of each risky asset purchased and the second two show the number of signals acquired about each asset. The matrix for each decision variable has dimensions 9x9. This structure represents the 16 possible mean belief combinations that the agent could wake up with today. The belief about the mean of asset 1 increases (along its grid) rightwards and that of asset 2; downwards, so as an example the top right entry represents the state in which the agent is as optimistic as possible about asset 1 and as pessimistic as possible about asset 2.

State	W	σ_1	σ_2
1	2	1.000	1.000
2	2	1.000	1.375
3	2	1.000	1.750
4	2	1.000	2.125
5	2	1.000	2.500
6	2	1.000	2.875
7	2	1.000	3.250
8	2	1.000	3.625
9	2	1.000	4.000
10	2	1.375	1.000
11	2	1.375	1.375
12	2	1.375	1.750
13	2	1.375	2.125
14	2	1.375	2.500
15	2	1.375	2.875
16	2	1.375	3.250
17	2	1.375	3.625
18	2	1.375	4.000
19	2	1.750	1.000
20	2	1.750	1.375
21	2	1.750	1.750
22	2	1.750	2.125
23	2	1.750	2.500
24	2	1.750	2.875
25	2	1.750	3.250
26	2	1.750	3.625
27	2	1.750	4.000
28	2	2.125	1.000
29	2	2.125	1.375
30	2	2.125	1.750
31	2	2.125	2.125
32	2	2.125	2.500
33	2	2.125	2.875
34	2	2.125	3.250
35	2	2.125	3.625
36	2	2.125	4.000
37	2	2.500	1.000
38	2	2.500	1.375
39	2	2.500	1.750
40	2	2.500	2.125
41	2	2.500	2.500

State	W	σ_1	σ_2
42	2	2.500	2.875
43	2	2.500	3.250
44	2	2.500	3.625
45	2	2.500	4.000
46	2	2.875	1.000
47	2	2.875	1.375
48	2	2.875	1.750
49	2	2.875	2.125
50	2	2.875	2.500
51	2	2.875	2.875
52	2	2.875	3.250
53	2	2.875	3.625
54	2	2.875	4.000
55	2	3.250	1.000
56	2	3.250	1.375
57	2	3.250	1.750
58	2	3.250	2.125
59	2	3.250	2.500
60	2	3.250	2.875
61	2	3.250	3.250
62	2	3.250	3.625
63	2	3.250	4.000
64	2	3.625	1.000
65	2	3.625	1.375
66	2	3.625	1.750
67	2	3.625	2.125
68	2	3.625	2.500
69	2	3.625	2.875
70	2	3.625	3.250
71	2	3.625	3.625
72	2	3.625	4.000
73	2	4.000	1.000
74	2	4.000	1.375
75	2	4.000	1.750
76	2	4.000	2.125
77	2	4.000	2.500
78	2	4.000	2.875
79	2	4.000	3.250
80	2	4.000	3.625
81	2	4.000	4.000

Table A.1: State variable values

State = 1																		
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
$\mu_1 \rightarrow$	0.00	0.25	0.50	0.75	1.00	1.25	1.50	1.75	2.00	0.00	0.25	0.50	0.75	1.00	1.25	1.50	1.75	2.00
0.00	-0.48	-0.16	0.00	0.00	0.00	0.00	0.00	0.16	0.48	-0.48	-0.80	-0.96	-0.96	-0.96	-0.96	-0.96	-0.80	-0.48
0.25	-0.80	-0.48	-0.16	0.00	0.00	0.00	0.16	0.48	0.80	-0.16	-0.48	-0.80	-0.96	-0.96	-0.96	-0.80	-0.48	-0.16
0.50	-0.96	-0.80	-0.48	-0.16	0.00	0.16	0.48	0.80	0.96	0.00	-0.16	-0.48	-0.80	-0.88	-0.80	-0.48	-0.16	0.00
0.75	-0.96	-0.96	-0.80	-0.32	0.00	0.32	0.80	0.96	0.96	0.00	0.00	-0.16	-0.32	-0.40	-0.32	-0.16	0.00	0.00
1.00	-0.96	-0.96	-0.88	-0.40	0.00	0.40	0.88	0.96	0.96	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.96	-0.96	-0.80	-0.32	0.00	0.32	0.80	0.96	0.96	0.00	0.00	0.16	0.32	0.40	0.32	0.16	0.00	0.00
1.50	-0.96	-0.80	-0.48	-0.16	0.00	0.16	0.48	0.80	0.96	0.00	0.16	0.48	0.80	0.88	0.80	0.48	0.16	0.00
1.75	-0.80	-0.48	-0.16	0.00	0.00	0.00	0.16	0.48	0.80	0.16	0.48	0.80	0.96	0.96	0.96	0.80	0.48	0.16
2.00	-0.48	-0.16	0.00	0.00	0.00	0.00	0.00	0.16	0.48	0.48	0.80	0.96	0.96	0.96	0.96	0.96	0.80	0.48
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.25	2	1	0	0	0	0	0	1	2	0	1	2	2	2	2	2	1	0
0.50	2	2	1	0	0	0	1	2	2	0	0	1	2	2	2	1	0	0
0.75	2	2	2	1	0	1	2	2	2	0	0	0	1	2	1	0	0	0
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	2	2	2	1	0	1	2	2	2	0	0	0	1	2	1	0	0	0
1.50	2	2	1	0	0	0	1	2	2	0	0	1	2	2	2	1	0	0
1.75	2	1	0	0	0	0	0	1	2	0	1	2	2	2	2	2	1	0
2.00	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1

State = 2																		
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
0.00	-0.40	-0.16	0.00	0.00	0.00	0.00	0.00	0.16	0.40	-0.56	-0.80	-0.96	-0.96	-0.96	-0.96	-0.96	-0.80	-0.56
0.25	-0.72	-0.40	-0.16	0.00	0.00	0.00	0.16	0.40	0.72	-0.24	-0.56	-0.80	-0.96	-0.96	-0.96	-0.80	-0.56	-0.24
0.50	-0.80	-0.72	-0.40	-0.16	0.00	0.16	0.40	0.72	0.80	-0.16	-0.24	-0.56	-0.80	-0.88	-0.80	-0.56	-0.24	-0.16
0.75	-0.96	-0.80	-0.56	-0.24	0.00	0.24	0.56	0.80	0.96	0.00	-0.16	-0.16	-0.32	-0.40	-0.32	-0.16	-0.16	0.00
1.00	-0.96	-0.88	-0.56	-0.32	0.00	0.32	0.56	0.88	0.96	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.96	-0.80	-0.56	-0.24	0.00	0.24	0.56	0.80	0.96	0.00	0.16	0.16	0.32	0.40	0.32	0.16	0.16	0.00
1.50	-0.80	-0.72	-0.40	-0.16	0.00	0.16	0.40	0.72	0.80	0.16	0.24	0.56	0.80	0.88	0.80	0.56	0.24	0.16
1.75	-0.72	-0.40	-0.16	0.00	0.00	0.00	0.16	0.40	0.72	0.24	0.56	0.80	0.96	0.96	0.96	0.80	0.56	0.24
2.00	-0.40	-0.16	0.00	0.00	0.00	0.00	0.00	0.16	0.40	0.56	0.80	0.96	0.96	0.96	0.96	0.96	0.80	0.56
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.25	2	1	0	0	0	0	0	1	2	0	1	2	2	2	2	2	1	0
0.50	2	2	1	0	0	0	1	2	2	0	0	1	2	2	2	1	0	0
0.75	2	2	2	1	0	1	2	2	2	0	0	0	1	2	1	0	0	0
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	2	2	2	1	0	1	2	2	2	0	0	0	1	2	1	0	0	0
1.50	2	2	1	0	0	0	1	2	2	0	0	1	2	2	2	1	0	0
1.75	2	1	0	0	0	0	0	1	2	0	1	2	2	2	2	2	1	0
2.00	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1

State = 3																		
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
0.00	-0.40	-0.16	0.00	0.00	0.00	0.00	0.00	0.16	0.40	-0.56	-0.80	-0.96	-0.96	-0.96	-0.96	-0.96	-0.80	-0.56
0.25	-0.72	-0.40	-0.16	0.00	0.00	0.00	0.16	0.40	0.72	-0.24	-0.56	-0.80	-0.96	-0.96	-0.96	-0.80	-0.56	-0.24
0.50	-0.80	-0.72	-0.40	-0.16	0.00	0.16	0.40	0.72	0.80	-0.16	-0.24	-0.56	-0.80	-0.88	-0.80	-0.56	-0.24	-0.16
0.75	-0.96	-0.80	-0.56	-0.24	0.00	0.24	0.56	0.80	0.96	0.00	-0.16	-0.16	-0.32	-0.40	-0.32	-0.16	-0.16	0.00
1.00	-0.96	-0.88	-0.56	-0.32	0.00	0.32	0.56	0.88	0.96	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.96	-0.80	-0.56	-0.24	0.00	0.24	0.56	0.80	0.96	0.00	0.16	0.16	0.32	0.40	0.32	0.16	0.16	0.00
1.50	-0.80	-0.72	-0.40	-0.16	0.00	0.16	0.40	0.72	0.80	0.16	0.24	0.56	0.80	0.88	0.80	0.56	0.24	0.16
1.75	-0.72	-0.40	-0.16	0.00	0.00	0.00	0.16	0.40	0.72	0.24	0.56	0.80	0.96	0.96	0.96	0.80	0.56	0.24
2.00	-0.40	-0.16	0.00	0.00	0.00	0.00	0.00	0.16	0.40	0.56	0.80	0.96	0.96	0.96	0.96	0.96	0.80	0.56
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
0.25	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.50	2	1	0	0	0	0	0	1	2	0	1	2	2	2	2	2	1	0
0.75	2	2	2	0	0	0	2	2	2	0	0	0	2	2	2	0	0	0
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	2	2	2	0	0	0	2	2	2	0	0	0	2	2	2	0	0	0
1.50	2	1	0	0	0	0	0	1	2	0	1	2	2	2	2	2	1	0
1.75	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
2.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2

Table A.2: Numerical solutions 1/3

	State = 4																	
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
$\mu_1 \rightarrow$	0.00	0.25	0.50	0.75	1.00	1.25	1.50	1.75	2.00	0.00	0.25	0.50	0.75	1.00	1.25	1.50	1.75	2.00
0.00	-0.16	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.16	-0.80	-0.88	-0.96	-0.96	-0.96	-0.96	-0.96	-0.88	-0.80
0.25	-0.40	-0.16	-0.08	0.00	0.00	0.00	0.08	0.16	0.40	-0.56	-0.80	-0.88	-0.96	-0.96	-0.96	-0.88	-0.80	-0.56
0.50	-0.64	-0.40	-0.16	-0.08	0.00	0.08	0.16	0.40	0.64	-0.32	-0.56	-0.80	-0.88	-0.88	-0.88	-0.80	-0.56	-0.32
0.75	-0.72	-0.56	-0.32	-0.08	0.00	0.08	0.32	0.56	0.72	-0.16	-0.16	-0.32	-0.40	-0.40	-0.40	-0.32	-0.16	-0.16
1.00	-0.72	-0.56	-0.32	-0.16	0.00	0.16	0.32	0.56	0.72	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.72	-0.56	-0.32	-0.08	0.00	0.08	0.32	0.56	0.72	0.16	0.16	0.32	0.40	0.40	0.40	0.32	0.16	0.16
1.50	-0.64	-0.40	-0.16	-0.08	0.00	0.08	0.16	0.40	0.64	0.32	0.56	0.80	0.88	0.88	0.88	0.80	0.56	0.32
1.75	-0.40	-0.16	-0.08	0.00	0.00	0.00	0.08	0.16	0.40	0.56	0.80	0.88	0.96	0.96	0.96	0.88	0.80	0.56
2.00	-0.16	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.16	0.80	0.88	0.96	0.96	0.96	0.96	0.96	0.88	0.80
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
0.25	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.50	2	1	0	0	0	0	0	1	2	0	1	2	2	2	2	2	1	0
0.75	2	2	1	0	0	0	1	2	2	0	0	1	2	2	2	1	0	0
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	2	2	1	0	0	0	1	2	2	0	0	1	2	2	2	1	0	0
1.50	2	1	0	0	0	0	0	1	2	0	1	2	2	2	2	2	1	0
1.75	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
2.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2

	State = 5																	
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
0.00	-0.16	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.16	-0.80	-0.88	-0.96	-0.96	-0.96	-0.96	-0.96	-0.88	-0.80
0.25	-0.32	-0.16	-0.08	0.00	0.00	0.00	0.08	0.16	0.32	-0.64	-0.80	-0.88	-0.96	-0.96	-0.96	-0.88	-0.80	-0.64
0.50	-0.48	-0.32	-0.16	-0.08	0.00	0.08	0.16	0.32	0.48	-0.48	-0.64	-0.80	-0.88	-0.88	-0.88	-0.80	-0.64	-0.48
0.75	-0.56	-0.40	-0.24	-0.08	0.00	0.08	0.24	0.40	0.56	-0.16	-0.16	-0.32	-0.40	-0.40	-0.40	-0.32	-0.16	-0.16
1.00	-0.56	-0.40	-0.32	-0.16	0.00	0.16	0.32	0.40	0.56	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.56	-0.40	-0.24	-0.08	0.00	0.08	0.24	0.40	0.56	0.16	0.16	0.32	0.40	0.40	0.40	0.32	0.16	0.16
1.50	-0.48	-0.32	-0.16	-0.08	0.00	0.08	0.16	0.32	0.48	0.48	0.64	0.80	0.88	0.88	0.88	0.80	0.64	0.48
1.75	-0.32	-0.16	-0.08	0.00	0.00	0.00	0.08	0.16	0.32	0.64	0.80	0.88	0.96	0.96	0.96	0.88	0.80	0.64
2.00	-0.16	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.16	0.80	0.88	0.96	0.96	0.96	0.96	0.96	0.88	0.80
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
0.25	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.50	1	1	0	0	0	0	0	1	1	1	1	2	2	2	2	2	1	1
0.75	2	2	1	0	0	0	1	2	2	0	0	1	2	2	2	1	0	0
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	2	2	1	0	0	0	1	2	2	0	0	1	2	2	2	1	0	0
1.50	1	1	0	0	0	0	0	1	1	1	1	2	2	2	2	2	1	1
1.75	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
2.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2

	State = 6																	
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
0.00	-0.08	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.08	-0.88	-0.88	-0.96	-0.96	-0.96	-0.96	-0.96	-0.88	-0.88
0.25	-0.32	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.32	-0.64	-0.88	-0.88	-0.96	-0.96	-0.96	-0.88	-0.88	-0.64
0.50	-0.40	-0.32	-0.08	-0.08	0.00	0.08	0.08	0.32	0.40	-0.56	-0.64	-0.88	-0.88	-0.88	-0.88	-0.88	-0.64	-0.56
0.75	-0.48	-0.32	-0.24	-0.08	0.00	0.08	0.24	0.32	0.48	-0.16	-0.32	-0.32	-0.40	-0.40	-0.40	-0.32	-0.32	-0.16
1.00	-0.48	-0.40	-0.24	-0.16	0.00	0.16	0.24	0.40	0.48	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.48	-0.32	-0.24	-0.08	0.00	0.08	0.24	0.32	0.48	0.16	0.32	0.32	0.40	0.40	0.40	0.32	0.32	0.16
1.50	-0.40	-0.32	-0.08	-0.08	0.00	0.08	0.08	0.32	0.40	0.56	0.64	0.88	0.88	0.88	0.88	0.88	0.64	0.56
1.75	-0.32	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.32	0.64	0.88	0.88	0.96	0.96	0.96	0.88	0.88	0.64
2.00	-0.08	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.08	0.88	0.88	0.96	0.96	0.96	0.96	0.96	0.88	0.88
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
0.25	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.50	1	1	0	0	0	0	0	1	1	1	1	2	2	2	2	2	1	1
0.75	2	1	1	0	0	0	1	1	2	0	1	1	2	2	2	1	1	0
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	2	1	1	0	0	0	1	1	2	0	1	1	2	2	2	1	1	0
1.50	1	1	0	0	0	0	0	1	1	1	1	2	2	2	2	2	1	1
1.75	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
2.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2

	State = 7																	
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
0.00	-0.08	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.08	-0.88	-0.88	-0.96	-0.96	-0.96	-0.96	-0.96	-0.88	-0.88
0.25	-0.32	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.32	-0.64	-0.88	-0.88	-0.96	-0.96	-0.96	-0.88	-0.88	-0.64
0.50	-0.32	-0.24	-0.08	-0.08	0.00	0.08	0.08	0.24	0.32	-0.64	-0.64	-0.88	-0.88	-0.88	-0.88	-0.88	-0.64	-0.64
0.75	-0.40	-0.24	-0.16	-0.08	0.00	0.08	0.16	0.24	0.40	-0.16	-0.32	-0.32	-0.40	-0.40	-0.40	-0.32	-0.32	-0.16
1.00	-0.40	-0.32	-0.24	-0.08	0.00	0.08	0.24	0.32	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.40	-0.24	-0.16	-0.08	0.00	0.08	0.16	0.24	0.40	0.16	0.32	0.32	0.40	0.40	0.40	0.32	0.32	0.16
1.50	-0.32	-0.24	-0.08	-0.08	0.00	0.08	0.08	0.24	0.32	0.64	0.64	0.88	0.88	0.88	0.88	0.88	0.64	0.64
1.75	-0.32	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.32	0.64	0.88	0.88	0.96	0.96	0.96	0.88	0.88	0.64
2.00	-0.08	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.08	0.88	0.88	0.96	0.96	0.96	0.96	0.96	0.88	0.88
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
0.25	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.50	1	1	0	0	0	0	0	1	1	1	1	2	2	2	2	2	1	1
0.75	2	1	1	0	0	0	1	1	2	0	1	1	2	2	2	1	1	0
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	2	1	1	0	0	0	1	1	2	0	1	1	2	2	2	1	1	0
1.50	1	1	0	0	0	0	0	1	1	1	1	2	2	2	2	2	1	1
1.75	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
2.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2

Table A.3: Numerical solutions 2/3

$State = 8$																		
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
$\mu_1 \rightarrow$	0.00	0.25	0.50	0.75	1.00	1.25	1.50	1.75	2.00	0.00	0.25	0.50	0.75	1.00	1.25	1.50	1.75	2.00
0.00	-0.08	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.08	-0.88	-0.88	-0.96	-0.96	-0.96	-0.96	-0.96	-0.88	-0.88
0.25	-0.24	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.24	-0.72	-0.88	-0.88	-0.96	-0.96	-0.96	-0.88	-0.88	-0.72
0.50	-0.32	-0.24	-0.08	-0.08	0.00	0.08	0.08	0.24	0.32	-0.64	-0.64	-0.88	-0.88	-0.88	-0.88	-0.88	-0.64	-0.64
0.75	-0.32	-0.24	-0.16	-0.08	0.00	0.08	0.16	0.24	0.32	-0.32	-0.32	-0.32	-0.40	-0.40	-0.40	-0.32	-0.32	-0.32
1.00	-0.40	-0.24	-0.16	-0.08	0.00	0.08	0.16	0.24	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.32	-0.24	-0.16	-0.08	0.00	0.08	0.16	0.24	0.32	0.32	0.32	0.32	0.40	0.40	0.40	0.32	0.32	0.32
1.50	-0.32	-0.24	-0.08	-0.08	0.00	0.08	0.08	0.24	0.32	0.64	0.64	0.88	0.88	0.88	0.88	0.88	0.64	0.64
1.75	-0.24	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.24	0.72	0.88	0.88	0.96	0.96	0.96	0.88	0.88	0.72
2.00	-0.08	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.08	0.88	0.88	0.96	0.96	0.96	0.96	0.96	0.88	0.88
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
0.25	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.50	1	1	0	0	0	0	0	1	1	1	1	2	2	2	2	2	1	1
0.75	1	1	1	0	0	0	1	1	1	1	1	1	2	2	2	1	1	1
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	1	1	1	0	0	0	1	1	1	1	1	1	2	2	2	1	1	1
1.50	1	1	0	0	0	0	0	1	1	1	1	2	2	2	2	2	1	1
1.75	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
2.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
$State = 9$																		
$\mu_2 \downarrow$	$\mathbf{Q_1}$									$\mathbf{Q_2}$								
0.00	-0.08	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.08	-0.88	-0.88	-0.96	-0.96	-0.96	-0.96	-0.96	-0.88	-0.88
0.25	-0.08	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.08	-0.88	-0.88	-0.88	-0.96	-0.96	-0.96	-0.88	-0.88	-0.88
0.50	-0.32	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.32	-0.64	-0.88	-0.88	-0.88	-0.88	-0.88	-0.88	-0.88	-0.64
0.75	-0.32	-0.24	-0.16	0.00	0.00	0.00	0.16	0.24	0.32	-0.32	-0.32	-0.32	-0.40	-0.40	-0.40	-0.32	-0.32	-0.32
1.00	-0.32	-0.24	-0.16	-0.08	0.00	0.08	0.16	0.24	0.32	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1.25	-0.32	-0.24	-0.16	0.00	0.00	0.00	0.16	0.24	0.32	0.32	0.32	0.32	0.40	0.40	0.40	0.32	0.32	0.32
1.50	-0.32	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.32	0.64	0.88	0.88	0.88	0.88	0.88	0.88	0.88	0.64
1.75	-0.08	-0.08	-0.08	0.00	0.00	0.00	0.08	0.08	0.08	0.88	0.88	0.88	0.96	0.96	0.96	0.88	0.88	0.88
2.00	-0.08	-0.08	0.00	0.00	0.00	0.00	0.00	0.08	0.08	0.88	0.88	0.96	0.96	0.96	0.96	0.96	0.88	0.88
$\mu_2 \downarrow$	$\mathbf{Signal_1}$									$\mathbf{Signal_2}$								
0.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
0.25	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
0.50	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
0.75	1	1	1	0	0	0	1	1	1	1	1	1	2	2	2	1	1	1
1.00	2	2	2	2	0	2	2	2	2	0	0	0	0	0	0	0	0	0
1.25	1	1	1	0	0	0	1	1	1	1	1	1	2	2	2	1	1	1
1.50	1	0	0	0	0	0	0	0	1	1	2	2	2	2	2	2	2	1
1.75	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2
2.00	0	0	0	0	0	0	0	0	0	2	2	2	2	2	2	2	2	2

Table A.4: Numerical solutions 3/3

A.1.4 Kalman Filter

General Form

This subsection gives a brief introduction to the Kalman Filter and the formulas in it are mainly based on Ljungqvist & Sargent (2004). Their equations forecast the next period value of the unknown variable based on all observations currently held. Instead, this paper estimates the unknown variable today based on all observations held and so Ljungqvist & Sargent's formulas are adapted for that purpose. The Kalman Filter is a recursive algorithm for calculating, $\hat{x}_t = E[x_t|y^t]$:

$$\hat{x}_{t+1} = (1 - K_t G)\hat{x}_t + K_t y_t.$$

This is the filtering equation. The parameter K_t and the variance are given by the following formulas:

$$K_t = \Sigma_t G'(G\Sigma_t G' + R)^{-1}$$

$$\Sigma_{t+1} = \Sigma_t - \Sigma_t G'(G\Sigma_t G' + R)^{-1} G\Sigma_t$$

The Kalman Gain, K_t, is the weighting applied to the two components of the filtering equation; the signal and the prior belief which contains all previous posterior expectations of the mean. This derivation utilises the text Pasricha (2006) as well as Ljungqvist & Sargent (2004) and they are just two amongst many that contain information on the Kalman Filter.

The hidden state vector that the system wishes to model is:

$$x_{t+1} = Ax_t + w_{t+1}.$$

The process for the noisy signals on the hidden state variable is given by:

$$y_t = Gx_t + v_t.$$

The v_t and w_s are both Gaussian and iid as well as being orthogonal to each other (independent).

The processes in this paper

The process for the hidden state variable, f_t, in my model is:

$$f_{t+1} = \theta + \epsilon_{t+1},$$

where $\epsilon \sim N(0,1)$. So the risky asset return is the underlying value plus some shock drawn today from the distribution of shocks to the risky asset returns.
The process for the signal is given by:

$$\eta_t = f_t + e_{\eta,t},$$

where $e_\eta \sim N(0, \Sigma_\eta)$. So the signal is the true risky asset return today plus some shock drawn today from the distribution of shocks to the signals.

Applying the Kalman Filter to this paper's model

With regard to the filtering formula for x_t, it is true that $E[w_{t+1}^2] = 1$, $G = 1$ and $E[v_t^2] = \Sigma_\eta$. So in modelling the unobserved state variable, the algorithm can be transformed as follows:
$\hat{x}_{t+1} = (1 - K_t)\hat{x}_t + K_t y_t$
$\mu_{t+1} = (1 - K_t)\mu_t + K_t f_t$.
This is the formula for updating the belief about the mean when the agent learns exogenously (using observed risky asset returns). The same formula when the agent learns endogenously (using endogenously acquired signals) is given by the following.
$\mu_{t+1} = (1 - K_t)\mu_t + K_t \eta_t$.

The Kalman Gain can be transformed in this way:
$K_t = \Sigma_t G'(G\Sigma_t G' + R)^{-1}$
$K_t = \Sigma_t(\Sigma_t + \Sigma_\eta)^{-1}$.
$K_t = \frac{\Sigma_t}{\Sigma_t+\Sigma_\eta}$.
The last equation is in the case that there is only one asset and so each item in the equation will be a scalar and not a matrix. In this project, the filtering variance is the uncertainty that the agent has about his mean belief. Hence the notation for his variance belief is a capital sigma but his uncertainty is a small sigma.
$K_t = \frac{\sigma_t}{\sigma_t+\sigma_\eta}$.

The expected variance can be manipulated in the way that proceeds:
$\Sigma_{t+1} = \Sigma_t - \Sigma_t G'(G\Sigma_t G' + R)^{-1} G\Sigma_t$
$\Sigma_{t+1} = \Sigma_t - \Sigma_t(\Sigma_t + \Sigma_\eta)^{-1}\Sigma_t$

Now this is put in the notation of this paper.
$\sigma_{t+1} = \sigma_t - \sigma_t(\sigma_t + \Sigma_\eta)^{-1}\sigma_t$

If there were only one asset then the formula could be simplified to be the following. Also, each element of the above matrix can be found by using the resulting expression. To confirm this, simply write out the matrices above explicitly for two or more assets and use the formula above.
$\sigma_{t+1} = \sigma_t - \frac{\sigma_t^2}{\sigma_t+\Sigma_\eta}$
$\sigma_{t+1} = \frac{\sigma_t^2+\sigma_t\Sigma_\eta-\sigma_t^2}{\sigma_t+\Sigma_\eta}$
$\sigma_{t+1} = \frac{\sigma_t\Sigma_\eta}{\sigma_t+\Sigma_\eta}$

A.2 Appendix: Paper 2

This appendix contains six sections. The first two contain data on past stock market returns (CAC40) and facts about heterogeneity in household beliefs about the stock market along with the two relevant questions from the survey. The third part derives the maximisation problem and the value function. Part four shows some typical time paths for each kind of investor whilst part five presents the probability of default for each investor in each state whilst the sixth part gives graphics showing deviations of the investors learning strategies from the myopic static benchmark.

A.2.1 Stock market returns & investor beliefs

This subsection provides data on recent & historical stock market returns (French CAC 40 index) as well as investor perspectives on these statistics (PAT€R survey of French households, 2014 wave) and the questions posed to elicit this information. Returns are the opening values of the market between 1 December 2011 and 1 December 2014 for recent returns and between 1 March 1990 and 1 March 2016 for historical returns. Table (A.5) describes the recent (2011-2014) and historical (1990-2016) CAC40 returns.

Recent & historical returns

Period	Range	Return	Mean	St.Dev.	Max	Min
3 years	1202.84	38.17%	3849.00	466.58	4523.90	3028.30
26 years	2507.93	136.60%	3598.73	1294.79	6648.64	1504.00

Table A.5: Recent & historical returns by subsample

Table A.6 describes perceptions of recent past returns. Table A.7 describes expected future returns.

Perceived Returns

Sample	Obs	Mean	St.Dev.	Max	Min
Whole	2328	0.0361	0.1204	-0.375	0.375
Subsamples					
High income	425	0.0710	0.1341	-0.375	0.375
Male	1210	0.0518	0.1243	-0.375	0.375
Married	1475	0.0329	0.1206	-0.375	0.375
College or higher	1049	0.0466	0.1279	-0.375	0.375

Table A.6: Perceived returns by subsample

Expected Returns

Sample	Obs	Mean	St.Dev.	Max	Min
Whole	2535	0.0162	0.0894	-0.625	0.625
Subsamples					
High income	439	0.0343	0.0798	-0.305	0.400
Male	1262	0.0228	0.0903	-0.625	0.625
Married	1593	0.0133	0.0849	-0.625	0.625
College or higher	1111	0.0223	0.0869	-0.525	0.625

Table A.7: Expected returns by subsample

A.2.2 Heterogeneity in investor beliefs

The following graphs (A.2, A.3, A.4 & A.5) are histograms of household perceptions of recent returns (past 3 years) and future expectations (next 5 years) by subsample. A red line (on perceived returns) indicates the actual return.

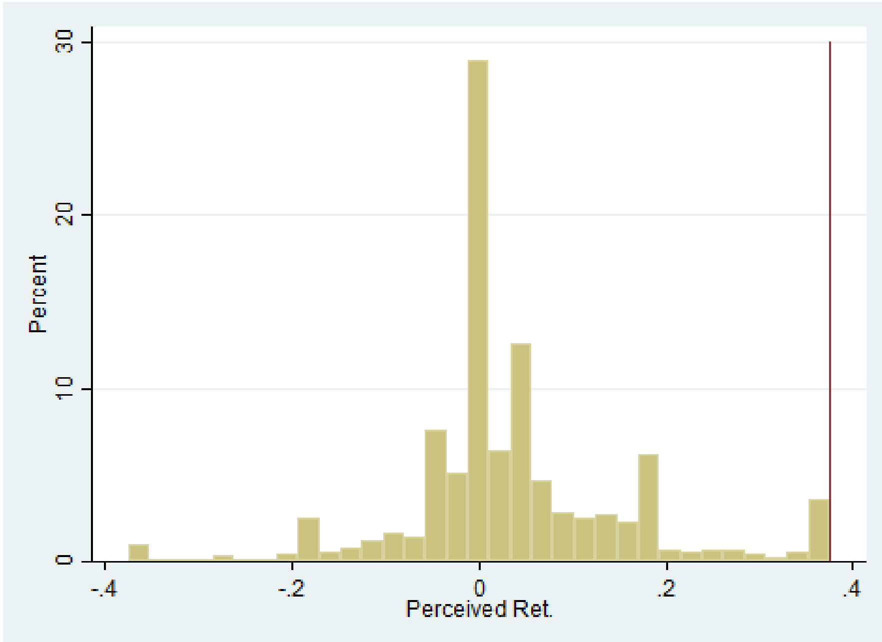

Figure A.2: Perceived returns, whole sample

Perceptions

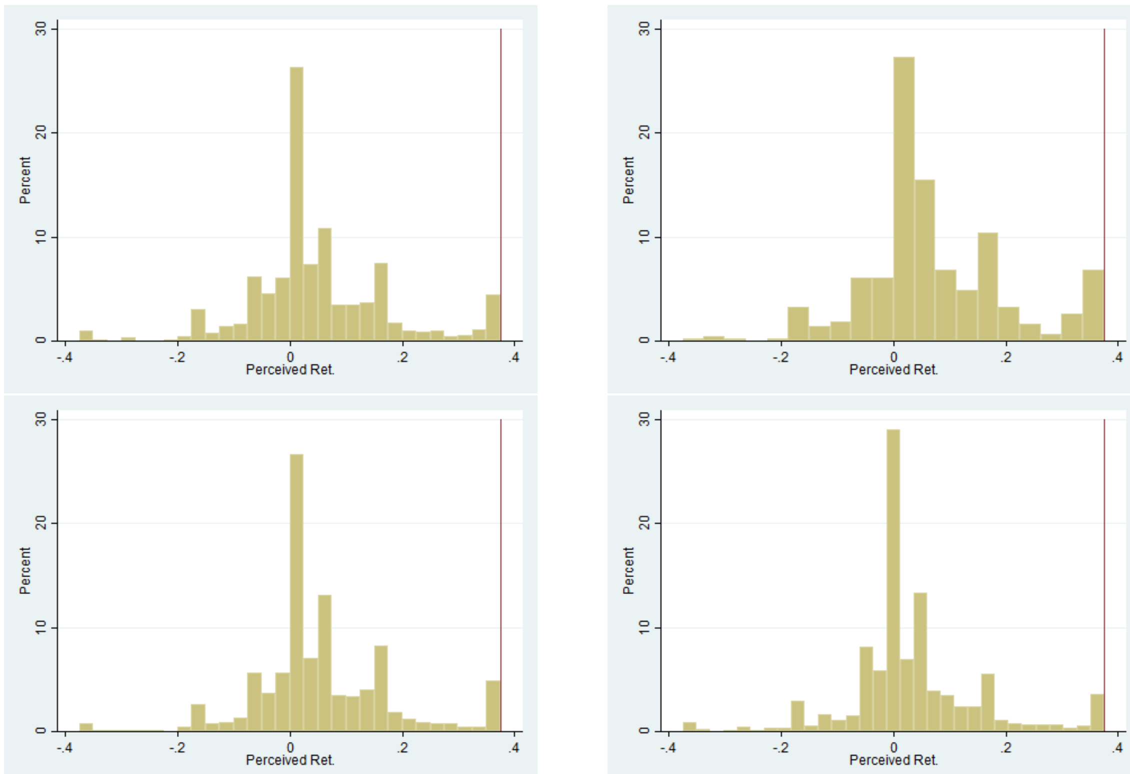

Figure A.3: Perceived returns of the highly educated, high earners, the married & males, in clockwise order

Expectations

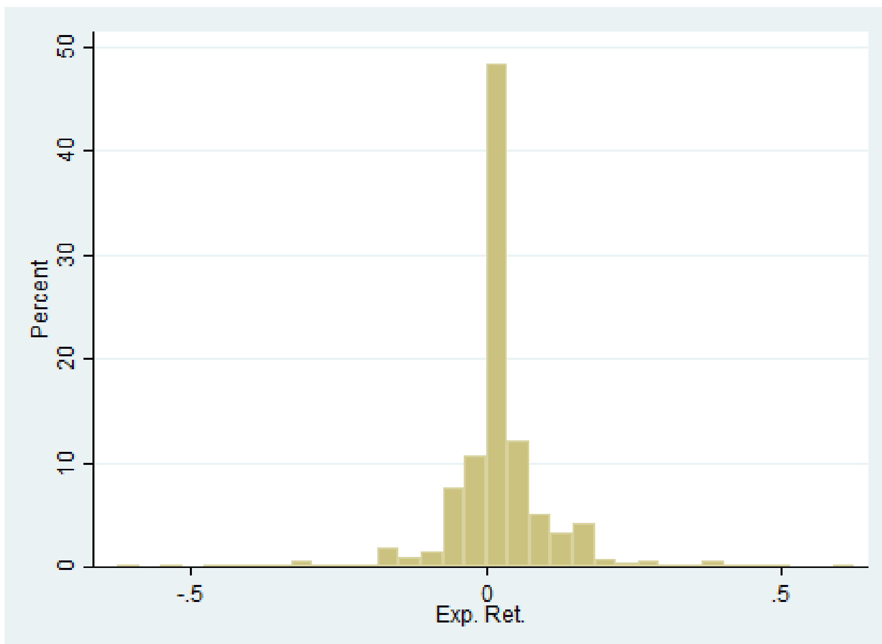

Figure A.4: Expected returns, whole sample

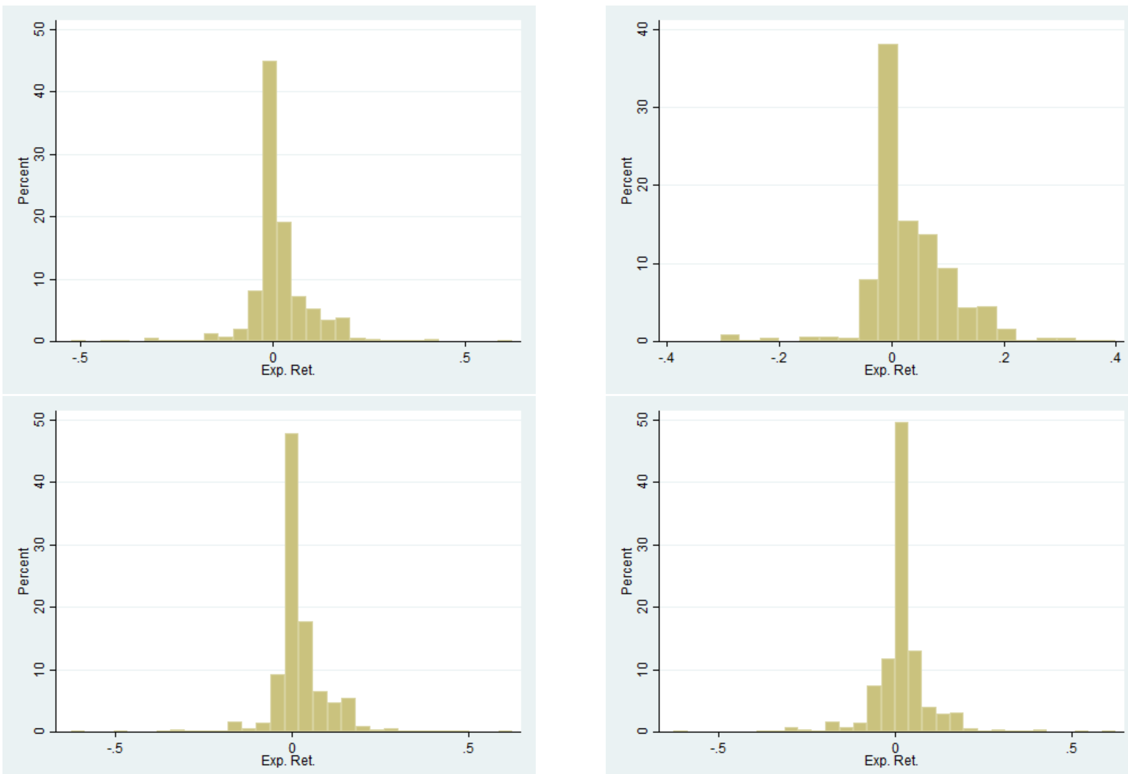

Figure A.5: Expected returns of the highly educated, high earners, the married & males, in clockwise order

Questionnaire

This part of the appendix presents the questions posed to participants in the questionnaire and links them to the variables used in the econometric analysis. The text is translated into English from French.

1. Perceptions of the stock market, Question C42

   C42. Over the last 3 years, do you think that the stock market... -For each category write down the probability of occurrence assigning a value between 0 and 100. The sum of all your answers must be equal to 100-:

   -... has increased by more than 25%

   -... has increased by 10 to 25%

   -... has increased by less than 10%

   -... has remained the same

   -... has decreased by less than 10%

   -... has decreased by 10 to 25%

   -... has decreased by more than 25%

2. Expectations of the stock market, Question C39

   C39. In 5 years' time, do you think that the stock market... -For each category write down the probability of occurrence assigning a value between 0 and 100. The sum of all your answers must be equal to 100-:

   -... will have increased by more than 25%

   -... will have increased by 10 to 25%

   -... will have increased by less than 10%

   -... will have remained the same

   -... will have decreased by less than 10%

-... will have decreased by 10 to 25%

-... will have decreased by more than 25%

A.2.3 Taking the expectation of the objective function

This subsection of the appendix derives (3.2.7) from (3.2.6), which is subject to (3.2.2), (3.2.3), (3.2.4) and (3.2.5). Take (3.2.6) and substitute in (3.2.2).

$$max_{q_t}||E_t[\sum_{t=1}^{T} -exp(-\rho(rW_t + q_t'(f_{t+1} - p_t r)))|\mu_t, \Sigma_t].$$

Use the facts that the expectation of a sum is the same as the sum of the expectations and that $E[aX] = aE[X]$, where $a = -1$ and X is a random variable.

$$max_{q_t}||\sum_{t=1}^{T} -E_t[exp(-\rho(rW_t + q_t'(f_{t+1} - p_t r)))|\mu_t, \Sigma_t].$$

f is normally distributed and it is known that if $ln(X)$ follows a normal distribution, where X is a random variable, then X is log normally distributed so $exp(-\rho W_{t+1})$ is log normally distributed. The expectation of a log normally distributed random variable, X, is $E[X] = exp(\mu + \frac{\sigma^2}{2})$ where μ is the mean of X and σ^2 is its variance. This formula can be applied to the above expression. The investor's expectation of the risky assets returns is given by $E_t[f_{t+1}|\mu_{t-1}, \Sigma_{t-1}] = \mu_t$ and his belief about the variance of the risky asset returns is $V_t[f_{t+1}|\mu_{t-1}, \Sigma_{t-1}] = \sigma_t + \Sigma_\epsilon$

$$max_{q_t}||\sum_{t=1}^{T} -exp(-\rho(rW_t + q_t'(\mu_t - p_t r)) + \frac{1}{2}\rho^2 q_t'(\sigma_t + \Sigma_\epsilon)q_t).$$

It is simpler to bundle the true variance of risky asset returns, Σ_ϵ, and the inaccuracy of beliefs, σ_t, into Σ_t because they both represent variation in returns to the investor. The expression found in the text, (3.2.7), is next.

$$max_{q_t}||\sum_{t=1}^{T} -exp(-\rho(rW_t + q_t'(\mu_t - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t q_t).$$

The value function that the agent faces can also be found. The outer (time) and inner (wealth) utility functions are given by the following two expressions:
$u^-(x) = y$ and $u^+(W) = -exp(-\rho W)$.

These give the following Bellman equation.

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{W_{t+1}\}}||E_t^-\{E_t^+\{-exp(-\rho W_{t+1})|\mu_t^+, \Sigma_t^+\}|\mu_t^-, \Sigma_t^-\} \\ +\beta E_t^-\{V(W_{t+1}, \mu_{t+1}^-, \Sigma_{t+1}^-)|\mu_t^-, \Sigma_t^-\} \end{aligned}$$

This expression is in terms of expectations so is not explicit. Substitute in the wealth process, (3.2.2).

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t\}}||E_t^-\{E_t^+\{-exp(-\rho(rW_t + q_t'(f_{t+1} - p_t r)))|\mu_t^+, \Sigma_t^+\}|\mu_t^-, \Sigma_t^-\} \\ +\beta E_t^+\{-exp(-\rho(rW_{t+1} + q_{t+1}'(f_{t+2} - p_{t+1} r)))|\mu_t^+, \Sigma_t^+\} \end{aligned}$$

Take the posterior expectations.

$$\begin{aligned} V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t\}}||E_t^-\{-exp(-\rho(rW_t + q_t'(\mu_t^+ - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t)|\mu_t^-, \Sigma_t^-\} \\ +\beta E_t^+\{-exp(-\rho(r(rW_t + q_t'(f_{t+1} - p_t r)) + q_{t+1}'(f_{t+2} - p_{t+1} r)))|\mu_t^+, \Sigma_t^+\} \end{aligned}$$

Take the prior expectations.

$$V(W_t, \mu_t^-, \Sigma_t^-) = max_{\{q_t\}} || - exp(-\rho(rW_t + q_t'(\mu_t^- - p_t r)) + \frac{1}{2}\rho^2 q_t' \Sigma_t^+ q_t + \frac{1}{2}\rho^2 q_t' K_t \Sigma_t^- K_t' q_t) - \beta exp(-\rho(r(rW_t + q_t'(\mu_t^+ - p_t r)) + q_{t+1}'(\mu_t^+ - p_{t+1} r)) + \frac{1}{2}\rho^2 (r^2 q_t' \Sigma_t^+ q_t + q_{t+1}' \Sigma_t^+ q_{t+1}))$$

Now the explicit value function is obtained, it can be parameterised for both numerical solutions and simulations.

A.2.4 Simulations

This subsection displays figures that show some example graphical results from the simulations. There are results from four versions of the code:

1. The agent learns only by observing actual returns (exogenous learning) and invests myopically.
2. The agent learns only by observing actual returns (exogenous learning) and invests dynamically.
3. The agent learns by acquiring signals and observing actual returns (endogenous & exogenous learning) and invests myopically.
4. The agent learns by acquiring signals and observing actual returns (endogenous & exogenous learning) and invests dynamically.
5. The agent learns by acquiring signals only (endogenous learning with observations switched off) and invests dynamically.

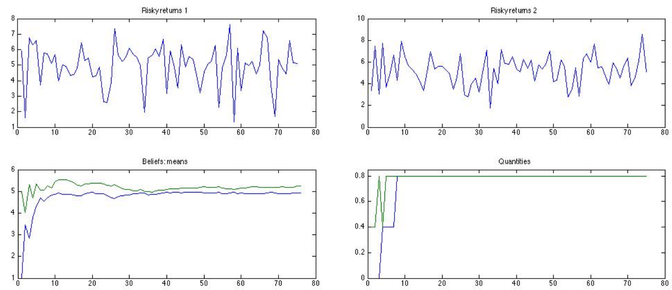

Figure A.6: Exogenous Learning 1

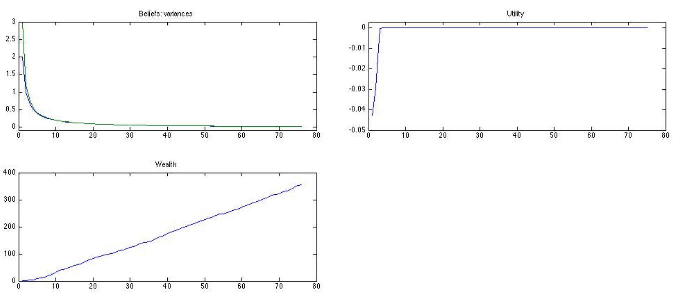

Figure A.7: Exogenous Learning 2

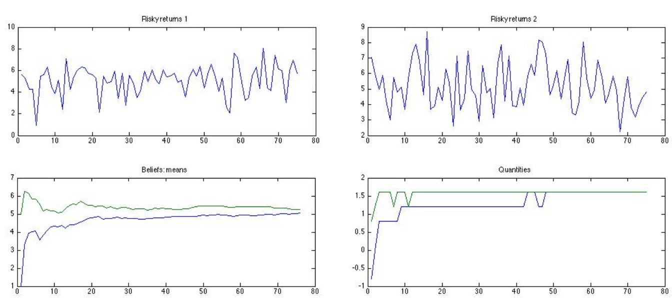

Figure A.8: Dynamic Exogenous Learning 1

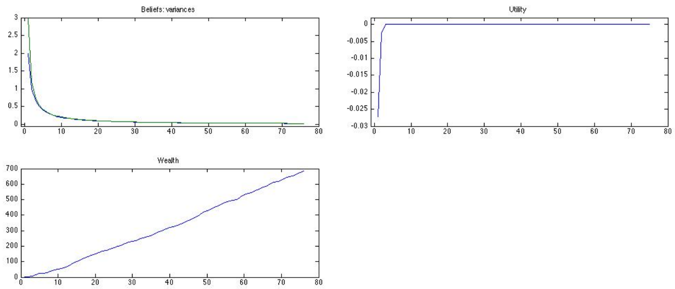

Figure A.9: Dynamic Exogenous Learning 2

(static endog)/A.jpg

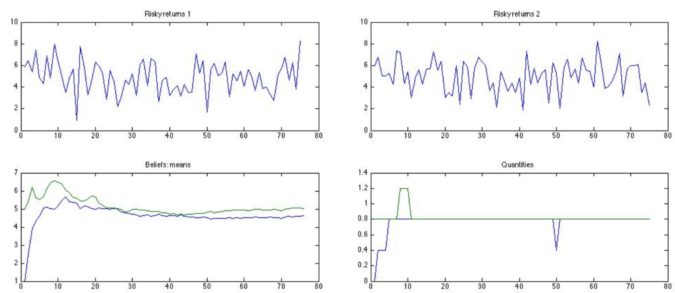

Figure A.10: Myopic Endogenous Learning 1

(static endog)/B.jpg

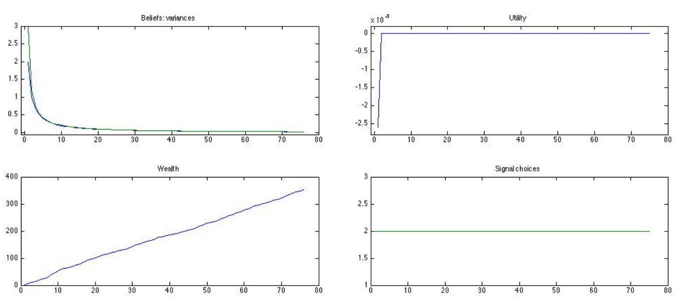

Figure A.11: Myopic Endogenous Learning 2

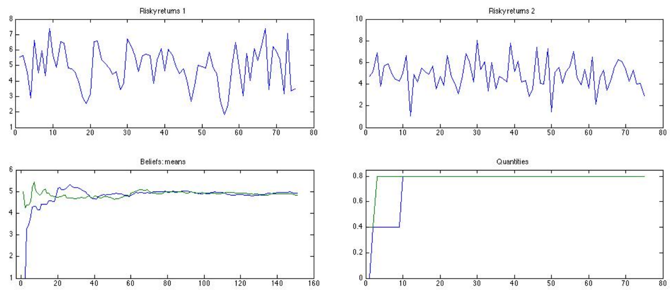

Figure A.12: Endogenous Learning 1

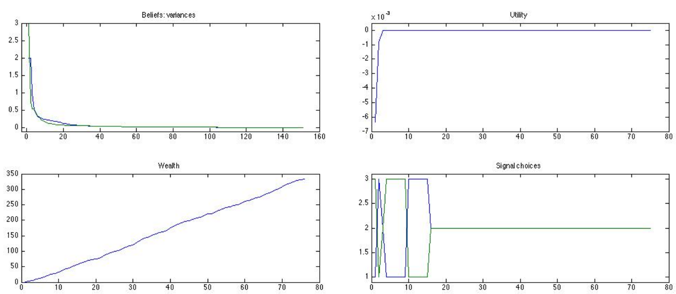

Figure A.13: Endogenous Learning 1

Switched Off/A.jpg

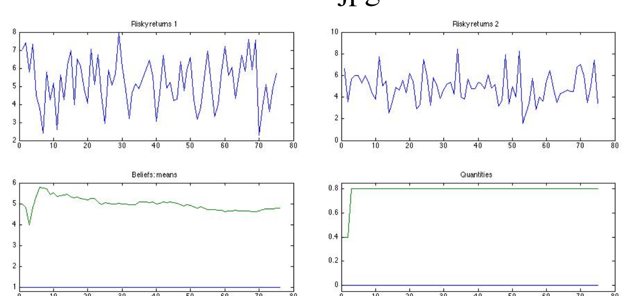

Figure A.14: Endogenous: no observation 1

Switched Off/B.jpg

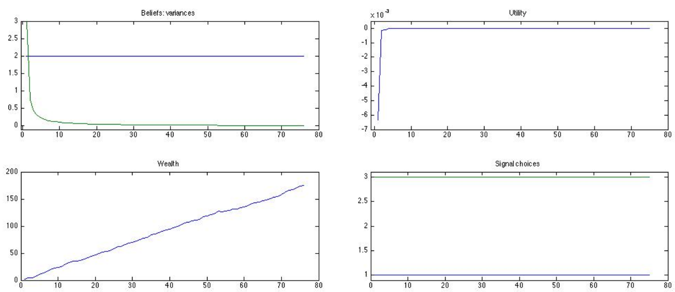

Figure A.15: Endogenous: no observation 2

A.2.5 Default probabilities

This subsection presents statistics on the probability of default and the investment duration for each agent in each state of the world using the four tables following. The different states are briefly described below.

- State 1:
  The investor is pessimistic about asset 1 and optimistic about asset 2 whilst not confident about either.
  $\mu_1 = 1, \mu_2 = 5, \sigma_1 = 2, \sigma_2 = 3$
- State 2:
  The investor is optimistic about both assets whilst not confident about either.
  $\mu_1 = 4.8, \mu_2 = 5.2, \sigma_1 = 2, \sigma_2 = 3$
- State 3:
  The investor is pessimistic about asset 1 and optimistic about asset 2 whilst confident about both.
  $\mu_1 = 1, \mu_2 = 5, \sigma_1 = 0.5, \sigma_2 = 0.8$
- State 4: The investor is very pessimistic about both assets whilst confident about both.
  $\mu_1 = -0.5, \mu_2 = -2, \sigma_1 = 0.5, \sigma_2 = 0.8$

Default probability					
Model type		**State 1**	**State 2**	**State 3**	**State 4**
Exogenous:	static	40.00%	0.00%	44.00%	100.00%
Exogenous:	dynamic	6.00%	0.00%	10.00%	100.00%
Endogenous:	static	4.00%	0.00%	20.00%	100.00%
Endogenous		0.00%	0.00%	6.00%	94.00%
Endogenous:	no observations	4.00%	0.00%	14.00%	100.00%

Table A.8: Default probabilities

Investment duration				
Model type	**State 1**	**State 2**	**State 3**	**State 4**
Exogenous: static	46.22	75.00	45.87	1.88
Exogenous: dynamic	70.78	75.00	70.27	1.96
Endogenous: static	72.12	75.00	62.97	1.26
Endogenous	75.00	75.00	72.60	6.18
Endogenous: no observations	72.18	75.00	65.53	1.70

Table A.9: Investment duration

A.2.6 Data on results 2 & 3

This section contains a number of graphs that record the learning choices of the investor across different trials for the different learning technologies. On the x-axis of each graph is the period number and on the y-axis is the learning choice, with 2 being generalisation between the two assets, 3 meaning all signals are devoted to the asset in question (specialisation) and 1 being no learning allocated to the asset. The different trials are distinguished by the different coloured dots. Only endogenous learners allocate signals so only the three relevant technologies are represented below, firstly for the myopic investor, secondly for the dynamic investor and lastly for the dynamic investor who does not observe returns. VN&V's static result of generalisation is represented by the red line at the value of two on each graph. Divergences illustrate how frequently the agent deviates from the static result in his learning choice. The myopic investor employs a constant strategy: generalising in states 1 & 2 and specialising (asset 1) in states 3 & 4 (figure A.16). Figures A.17-A.20 show that the dynamic investor's learning choices change over time. The top left graph in each represents the first state, the next shows the second state, the bottom left; the third and the last shows the learning choices in state 4.

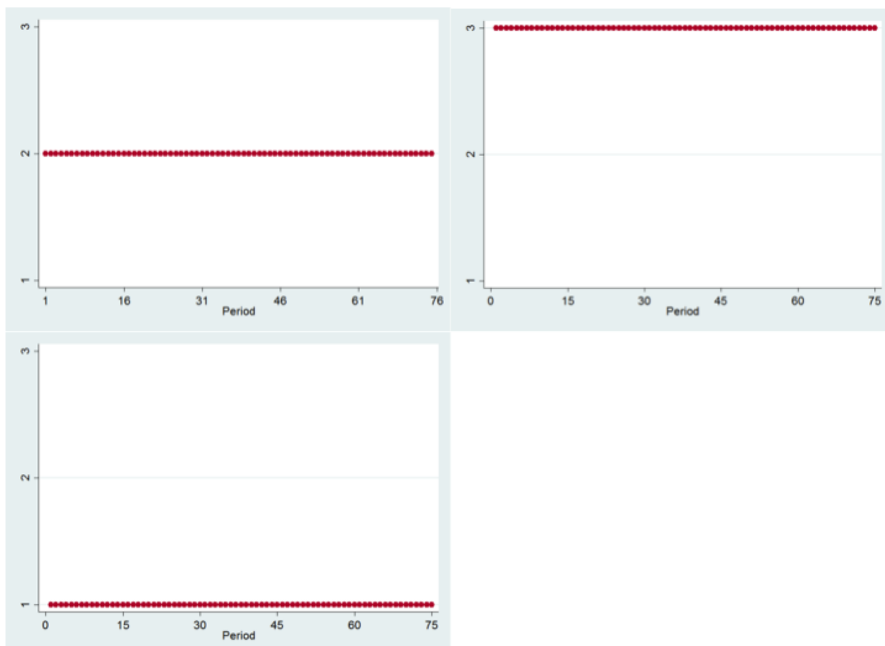

Figure A.16: Static endogenous: generalisation in states 1&2 (top left), asset 1 in states 3 & 4 (top right) and asset 2 in states 3 & 4 (bottom left)

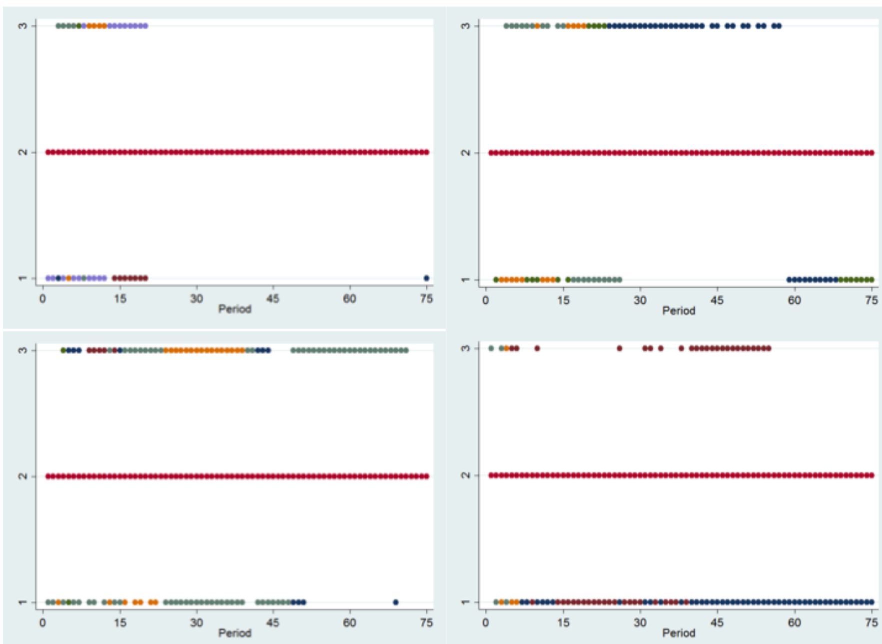

Figure A.17: Endogenous: asset 1, states 1-4 (left to right)

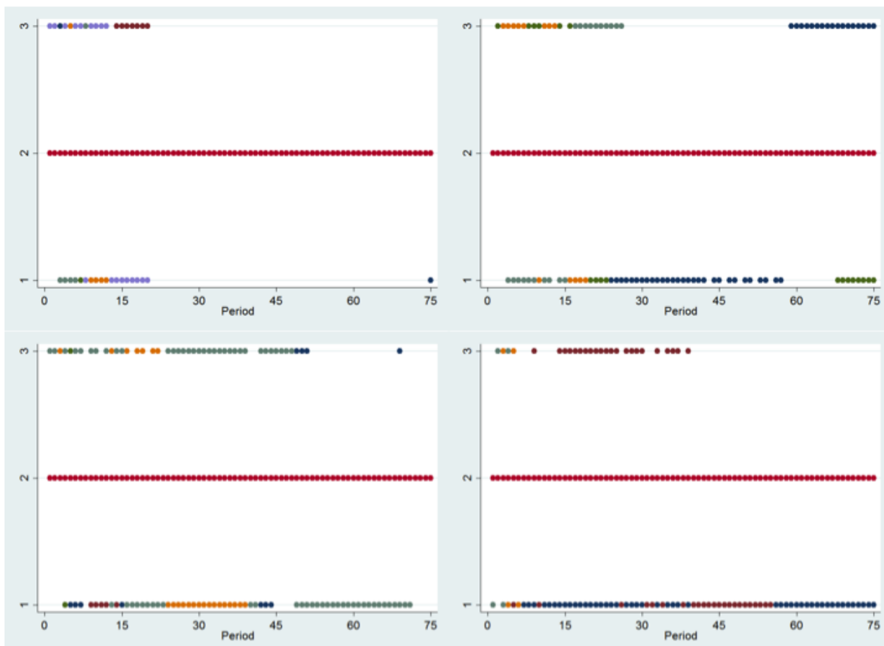

Figure A.18: Endogenous: asset 2, states 1-4 (left to right)

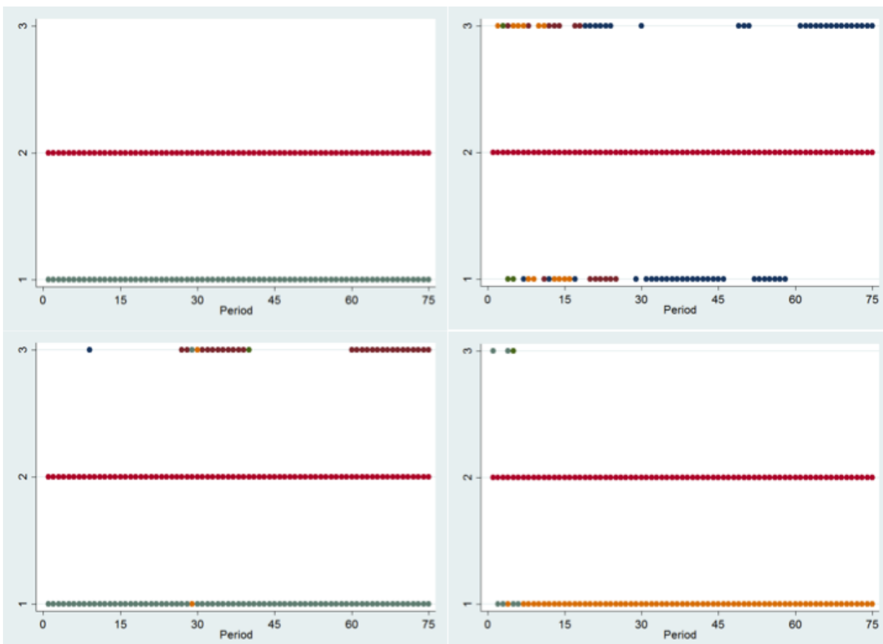

Figure A.19: Endogenous no observations: asset 1, states 1-4 (left to right)

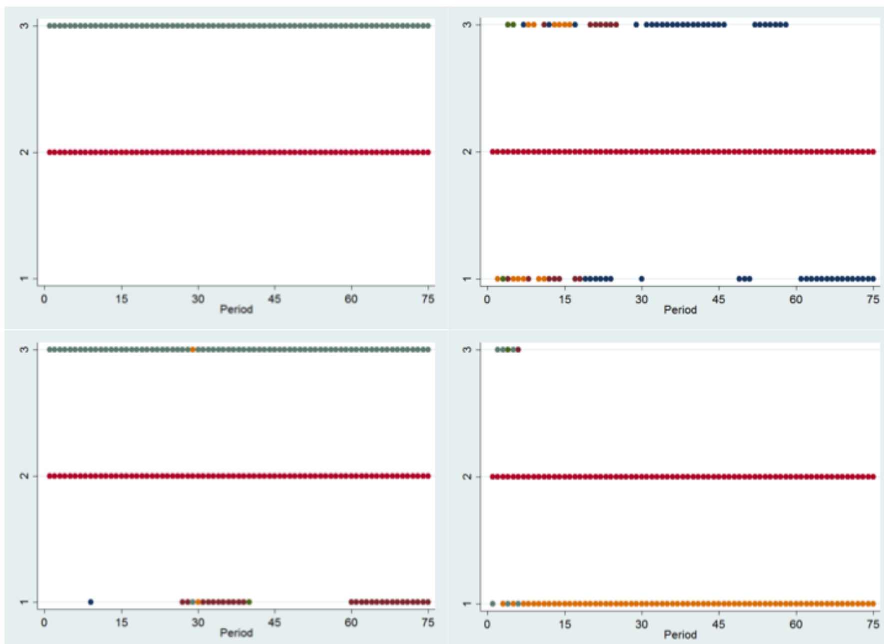

Figure A.20: Endogenous no observations: asset 2, states 1-4 (left to right)

A.3 Appendix: Paper 3

This appendix contains the regression results for the econometric analysis performed. It also contains some summary statistics on the dataset and copies of questions from the survey (translated from French to English) that apply to the variables used.

A.3.1 Data

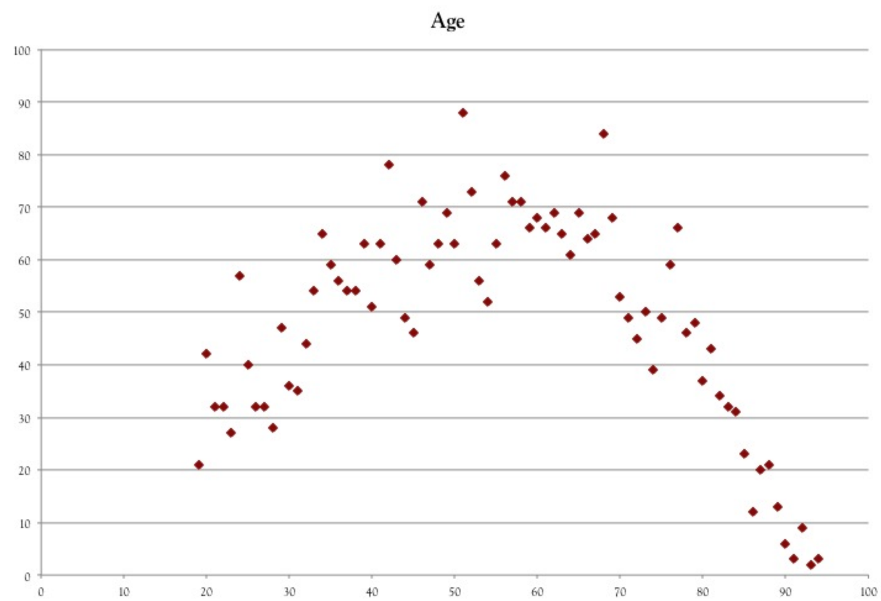

Age

Figure A.21:
The distribution of age is normal with some negative skew. The lowest age is 19 and the highest is 94, with an average of 54.

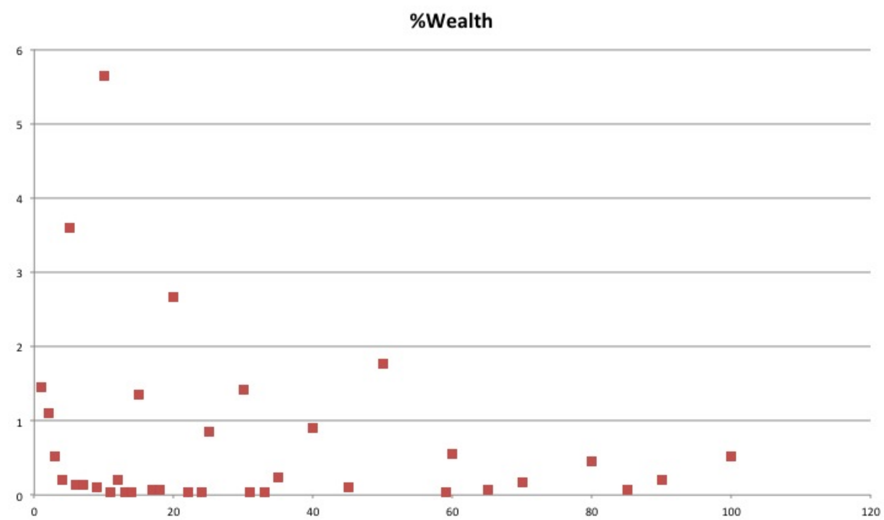

Wealth

Figure A.22:
This graph describes the percentage of wealth that households invest in the stock market. Non-participation (0%) is the most popular category at 75% but is omitted. Participation spikes at round numbers (5, 10, 20, 50, 100%) with the local peak being at 10%. Surprisingly, 0.5% of households invest all their wealth in the stock market, this is the same as the percentage that invest 60% of wealth.

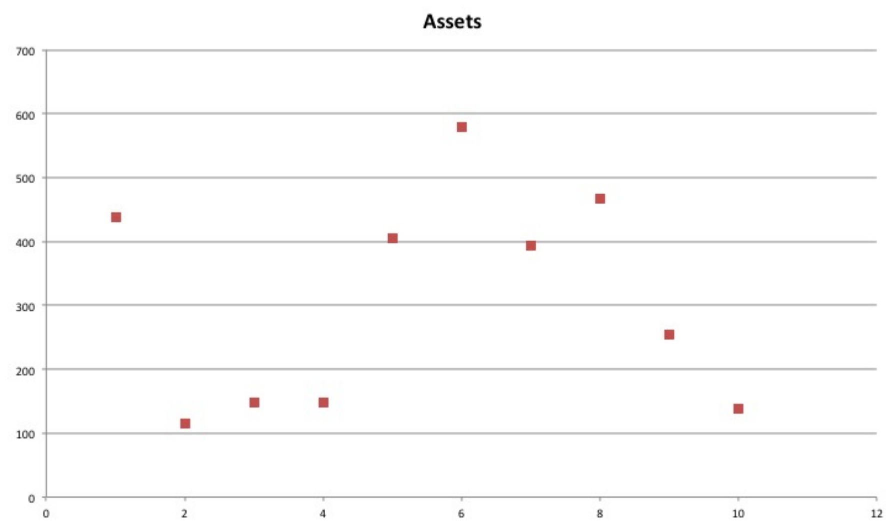

Figure A.23:
The most common assets categories are €75,000 to €449,999. There is also a spike at less than €8,000. Marriage and investing in the stock market are associated with higher asset holdings.

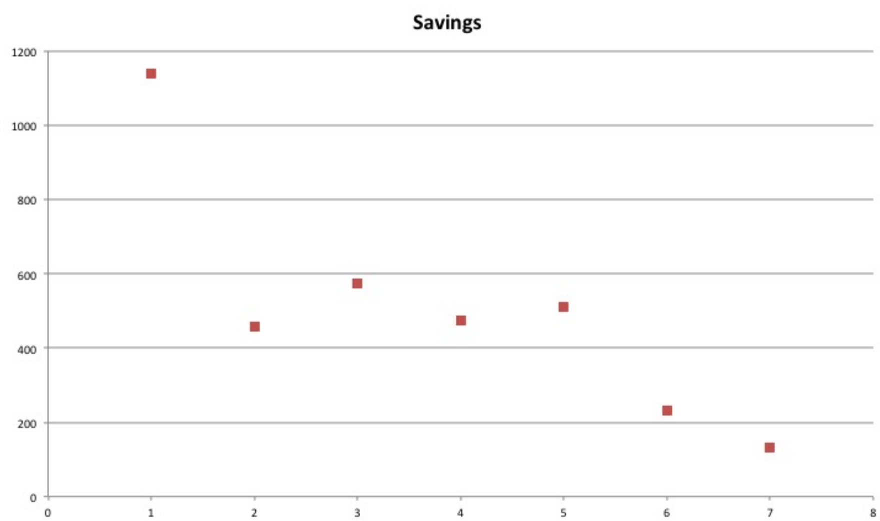

Savings

Figure A.24:
Higher savings categories are less popular. Savings decline with stock market participation and marriage, the inverse to asset holdings.

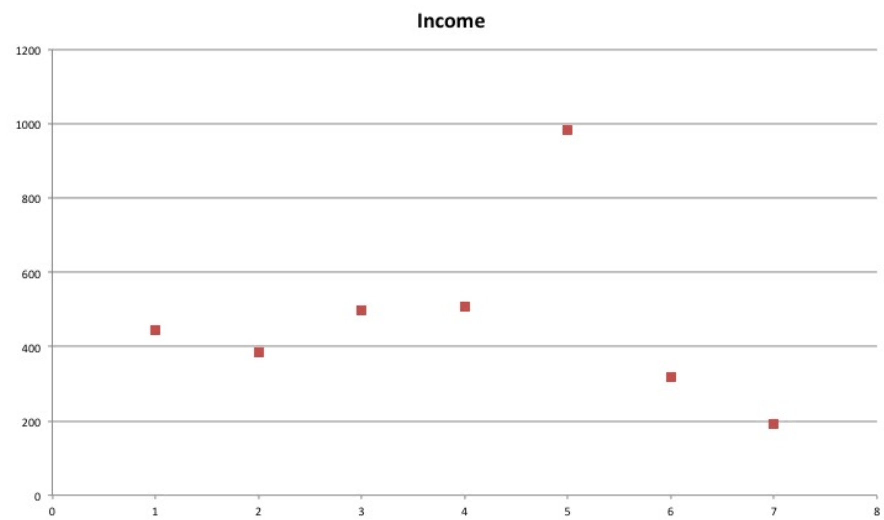

Income

Figure A.25:
Income is quite uniform but is lower for over €40,000 and much higher for €20,000 to €29,999.

Average participation of each subgroup and perceptions of others' participation & information.

Participation	
Subsample	Mean
All	21.69
Higher education	30.08
Top 3 income	31.59
Top income	58.95
Married	22.36

% Population participation	
Subsample	Mean
All	7.29
Higher education	10.23
Investors	12.28
Top 3 income	9.04
Top income	11.40
Married	7.50

% Population information	
Subsample	Mean
All	8.23
Higher education	11.35
Investors	13.05
Top 3 income	9.93
Top income	12.79
Married	8.44

Percentage of each subgroup that correctly identified a positive return and expects a positive future return.

Perceptions	
Subsample	(+)
All	52.49
Higher education	49.77
Investors	63.76
Top 3 income	58.94
Top income	70.91
Married	50.92

Expectations	
Subsample	(+)
All	44.85
Higher education	55.20
Investors	57.80
Top 3 income	51.67
Top income	60.47
Married	44.01

A.3.2 Regression analysis

List of Tables

VARIABLES	(1) BasePar Stocksi	(2) ParP Stocksi	(3) ParI Stocksi	(4) ParPI Stocksi	(5) ParPIInt Stocksi
Perceived % Pop. Part.		0.00474*		0.00930**	0.00731
		(0.00272)		(0.00371)	(0.00509)
Perceived % Pop. Informed			-0.00117	-0.00667*	-0.00848*
			(0.00251)	(0.00341)	(0.00455)
(% Pop. Part.)x(% Pop. Informed)					6.71e-05
					(0.000114)
Perceived Ret.	0.718***	0.709***	0.685***	0.683***	0.682***
	(0.219)	(0.218)	(0.218)	(0.218)	(0.218)
Sd. Perceived Ret.	-0.633***	-0.631***	-0.613***	-0.610***	-0.614***
	(0.226)	(0.227)	(0.226)	(0.227)	(0.227)
Age	0.00914	0.00930	0.00936	0.00918	0.00902
	(0.00948)	(0.00960)	(0.00957)	(0.00961)	(0.00962)
Age Sq.	1.28e-06	1.45e-05	1.42e-05	1.66e-05	1.78e-05
	(8.55e-05)	(8.60e-05)	(8.59e-05)	(8.62e-05)	(8.63e-05)
Male	0.131**	0.117**	0.113**	0.115**	0.115**
	(0.0540)	(0.0544)	(0.0544)	(0.0545)	(0.0545)
Married	-0.194***	-0.186***	-0.182***	-0.182***	-0.182***
	(0.0568)	(0.0571)	(0.0570)	(0.0571)	(0.0571)
CARA	-0.0172***	-0.0146**	-0.0148**	-0.0150**	-0.0151**
	(0.00595)	(0.00603)	(0.00599)	(0.00608)	(0.00610)
High school	0.00893	0.0528	0.0331	0.0480	0.0490
	(0.152)	(0.150)	(0.151)	(0.151)	(0.151)
Technical/Professional	0.0276	0.0488	0.0466	0.0494	0.0487
	(0.116)	(0.115)	(0.116)	(0.115)	(0.115)
College or more	0.252**	0.253**	0.247**	0.251**	0.250**
	(0.125)	(0.124)	(0.124)	(0.124)	(0.124)
75000<Assets<224999	0.140***	0.135***	0.130***	0.134***	0.134***
	(0.0413)	(0.0417)	(0.0416)	(0.0417)	(0.0417)
224500<Assets<449999	0.215***	0.210***	0.208***	0.211***	0.211***
	(0.0284)	(0.0286)	(0.0286)	(0.0286)	(0.0286)
450000<Assets	0.267***	0.260***	0.258***	0.258***	0.258***
	(0.0255)	(0.0256)	(0.0256)	(0.0256)	(0.0256)
0<Saving<999	0.111***	0.110***	0.113***	0.112***	0.112***
	(0.0357)	(0.0360)	(0.0359)	(0.0360)	(0.0360)
1000<Saving<4999	0.132***	0.129***	0.129***	0.129***	0.129***
	(0.0233)	(0.0234)	(0.0235)	(0.0235)	(0.0234)
Saving>5000	0.131***	0.126***	0.127***	0.128***	0.128***
	(0.0229)	(0.0231)	(0.0231)	(0.0232)	(0.0232)
12000<Income<19999	-0.00435	-0.00859	-0.0111	-0.00890	-0.00879
	(0.0372)	(0.0375)	(0.0374)	(0.0375)	(0.0375)
20000<Income<29999	0.00822	0.00528	0.00374	0.00405	0.00429
	(0.0249)	(0.0250)	(0.0249)	(0.0250)	(0.0250)
Income>30000	0.0639***	0.0590**	0.0569**	0.0571**	0.0574**
	(0.0230)	(0.0231)	(0.0231)	(0.0231)	(0.0231)
NR(Per. Ret.)	-0.315***	-0.267**	-0.264**	-0.261**	-0.265**
	(0.120)	(0.122)	(0.122)	(0.122)	(0.122)
NR(CARA)	-0.675***	-0.630**	-0.632**	-0.653**	-0.657**
	(0.260)	(0.264)	(0.263)	(0.266)	(0.267)
NR(Assets)	0.124	0.0985	0.101	0.0985	0.0969
	(0.106)	(0.109)	(0.109)	(0.110)	(0.110)
NR(Income)	0.306	0.316	0.307	0.322	0.318
	(0.231)	(0.233)	(0.234)	(0.232)	(0.234)
NR(Saving)	0.111	0.0423	0.0438	0.0456	0.0471
	(0.174)	(0.172)	(0.172)	(0.172)	(0.172)
NR(Per.% Pop. Part.)		0.189		0.534	0.364
		(0.202)		(0.551)	(0.673)
DK(Per.% Pop. Part.)		-0.191**		0.0796	0.0693
		(0.0869)		(0.150)	(0.153)
NR(Per.% Pop. Inf.)			-0.00586	-0.423	-0.654
			(0.198)	(0.539)	(0.720)
DK(Per.% Pop. Inf.)			-0.337***	-0.350**	-0.381**
			(0.0857)	(0.147)	(0.152)
NR(% Pop. Part.)x(% Pop. Informed)					0.347
					(0.944)
Constant	-1.305***	-1.372***	-1.216***	-1.290***	-1.240***
	(0.346)	(0.358)	(0.357)	(0.362)	(0.371)

Figure A.26: Regress participation in the stock market (discrete variable) on coordination motives and household characteristics under a probit technique.

Observations	3,606	3,606	3,606	3,606	3,606

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

VARIABLES	(1) ParPI Stocksi	(2) ParPIT Stocksi	(3) ParPIS Stocksi	(4) ParPII Stocksi	(5) ParPIK Stocksi	(6) ParPITS Stocksi	(7) ParPIIK Stocksi	(8) ParPIAll Stocksi
Perceived % Pop. Part.	0.00927**	0.00686	0.00773*	0.00727	0.00816**	0.00460	0.00663	0.00627
	(0.00394)	(0.00478)	(0.00408)	(0.00517)	(0.00392)	(0.00479)	(0.00508)	(0.00536)
Perceived % Pop. Informed	-0.00640*	-0.00491	-0.00548	-0.00481	-0.00657*	-0.00410	-0.00501	-0.00597
	(0.00358)	(0.00431)	(0.00373)	(0.00457)	(0.00358)	(0.00431)	(0.00455)	(0.00483)
Perceived Ret.	1.031***	0.662**	0.833***	0.745***	0.714***	0.644**	0.601**	0.731**
	(0.246)	(0.288)	(0.261)	(0.285)	(0.248)	(0.300)	(0.286)	(0.321)
Sd. Perceived Ret.	-0.698*	-0.496	-0.546	-0.556	-0.509	-0.368	-0.465	-0.302
	(0.412)	(0.491)	(0.457)	(0.478)	(0.422)	(0.527)	(0.479)	(0.556)
Age	0.0432***	0.0296**	0.0131	0.0158	0.0351***	0.0161	0.0134	0.00298
	(0.0114)	(0.0137)	(0.0127)	(0.0124)	(0.0116)	(0.0147)	(0.0125)	(0.0151)
Age Sq.	-0.000230**	-0.000154	1.96e-05	-6.09e-05	-0.000156	-3.77e-05	-3.54e-05	5.05e-05
	(0.000105)	(0.000128)	(0.000117)	(0.000117)	(0.000107)	(0.000137)	(0.000118)	(0.000141)
Male	0.172***	0.0363	0.122*	0.0375	0.0899	0.0124	0.00270	-0.0340
	(0.0600)	(0.0710)	(0.0644)	(0.0703)	(0.0618)	(0.0733)	(0.0712)	(0.0787)
Married	-0.0162	-0.0229	-0.0442	0.0332	-0.0334	-0.0318	0.0221	0.00321
	(0.0619)	(0.0732)	(0.0668)	(0.0736)	(0.0632)	(0.0765)	(0.0740)	(0.0823)
CARA	-0.0193***	-0.0176**	-0.0138**	-0.0127	-0.0127*	-0.0137*	-0.00961	-0.0116
	(0.00646)	(0.00722)	(0.00686)	(0.00786)	(0.00661)	(0.00754)	(0.00786)	(0.00828)
High school	0.0792	0.402*	0.337	0.611***	0.199	0.539**	0.666***	0.829***
	(0.200)	(0.231)	(0.215)	(0.230)	(0.204)	(0.244)	(0.232)	(0.255)
Technical/Professional	0.174	0.163	0.225	0.396**	0.248	0.205	0.434**	0.390*
	(0.152)	(0.182)	(0.163)	(0.192)	(0.156)	(0.194)	(0.195)	(0.213)
College or more	0.580***	0.410**	0.535***	0.630***	0.594***	0.427**	0.645***	0.580***
	(0.156)	(0.187)	(0.168)	(0.196)	(0.160)	(0.200)	(0.199)	(0.218)
Not investing: NSP		-7.099***				-7.381***		-6.728***
		(0.151)				(0.215)		(0.205)
Descending adviser trust		0.262***				0.152***		0.114***
		(0.0280)				(0.0340)		(0.0364)
Trust others		-0.124*				-0.135*		-0.131*
		(0.0713)				(0.0741)		(0.0791)
Friends & relations			0.0113			0.0681		0.103
			(0.0730)			(0.0808)		(0.0859)
Family			-0.0181			-0.105		-0.139*
			(0.0666)			(0.0738)		(0.0780)
Financial advisers			0.286***			0.311***		0.276***
			(0.0517)			(0.0598)		(0.0655)
General media			0.0251			0.0954		0.179
			(0.110)			(0.123)		(0.138)
Specialist media			0.229***			0.0124		-0.0968
			(0.0691)			(0.0826)		(0.0971)
Friends & relations: Finan decision - rely on...?			-0.152**			-0.109		-0.109
			(0.0676)			(0.0748)		(0.0762)
Family: Finan decision - rely on...?			0.0150			0.0772		0.0987*
			(0.0505)			(0.0569)		(0.0577)
Banker/financial adviser: Finan decision - rely on...?			0.0266			-0.0419		0.0146
			(0.0488)			(0.0594)		(0.0627)
Financial knowledge: Finan decision - rely on...?			0.123***			0.0792		0.0609
			(0.0448)			(0.0511)		(0.0563)
Media: Finan decision - rely on...?			0.0490			0.0126		-0.0927
			(0.0659)			(0.0754)		(0.0828)
Written press			0.112			0.104		0.141
			(0.0749)			(0.0862)		(0.0968)
Audiovisual			0.00541			0.0258		-0.0376
			(0.0705)			(0.0847)		(0.0984)
Online social websites			-0.146**			-0.118		-0.0999
			(0.0739)			(0.0823)		(0.0926)
Finance websites			0.0234			0.0503		-0.0968
			(0.0822)			(0.0954)		(0.114)
Financial authorities websites			0.0837			0.0198		-0.0291
			(0.103)			(0.128)		(0.150)
No. Operations(t-1)				0.273***			0.242***	0.279***
				(0.0806)			(0.0824)	(0.0965)
Descending finan awareness					0.342***		0.170***	0.0586

Figure A.27: Regress participation in the stock market (discrete variable) on coordination motives, household characteristics and information variables under a probit technique.

					(0.0318)		(0.0364)	(0.0464)
NR(CARA)	-0.963**	-	-0.879**	-0.698	-0.632	-	-0.564	-1.063*
		1.478***				1.355***		
	(0.406)	(0.547)	(0.413)	(0.439)	(0.399)	(0.519)	(0.442)	(0.564)
DK(Per.% Pop. Part.)	-0.0737	-0.0974	-0.0630	0.147	-0.0360	-0.118	0.149	0.0492
	(0.161)	(0.185)	(0.172)	(0.195)	(0.163)	(0.192)	(0.193)	(0.211)
DK(Per.% Pop. Inf.)	-0.305*	-0.127	-0.142	-0.332*	-0.253	-0.0626	-0.305	-0.196
	(0.158)	(0.182)	(0.168)	(0.190)	(0.160)	(0.187)	(0.189)	(0.205)
NR(Not investing: NSP)		7.905***				8.149***		7.215***
		(0.179)				(0.240)		(0.231)
NR(Descending adviser trust)		0.666				-0.0130		-0.189
		(0.610)				(0.551)		(0.507)
NR(Trust others)		-0.284				-0.325		-0.754**
		(0.407)				(0.432)		(0.332)
NR(Friends & relations)			0.154			0.243		0.218
			(0.362)			(0.494)		(0.457)
NR(Family)			-0.279			-0.566		-0.531
			(0.382)			(0.481)		(0.419)
NR(Financial advisers)			0.888***			0.959**		0.843**
			(0.327)			(0.403)		(0.424)
NR(General media)			0.0513			-0.255		-0.0398
			(0.449)			(0.528)		(0.476)
NR(Specialist media)			0.273			0.154		-0.218
			(0.444)			(0.476)		(0.441)
NR(Friend & relations: Finan decision - rely on...?)			-0.191			0.0609		0.133
			(0.394)			(0.463)		(0.438)
NR(Family: Finan decision - rely on...?)			0.155			-0.150		0.118
			(0.371)			(0.489)		(0.426)
NR(Banker/financial adviser: Finan decision - rely on...?)			-0.137			-0.0720		0.510
			(0.417)			(0.444)		(0.428)
NR(Financial knowledge: Finan decision - rely on...?)			0.260			0.403		0.281
			(0.388)			(0.473)		(0.520)
NR(Media: Finan decision - rely on...?)			0.0173			-0.172		-0.673
			(0.407)			(0.542)		(0.567)
NR(Written press)			0.228			0.0419		0.512
			(0.494)			(0.481)		(0.520)
NR(Audiovisual)			0.755			0.720		0.214
			(0.604)			(0.678)		(0.713)
NR(Online soc websites)			-0.674			-0.555		-0.489
			(0.590)			(0.807)		(0.835)
NR(Finance websites)			-0.486			-0.246		-0.760
			(0.598)			(0.624)		(0.607)
NR(Finan authorities websites)			0.465			0.249		0.0853
			(0.521)			(0.583)		(0.576)
NR(Online investment info)			-0.197			0.0113		0.319
			(0.672)			(0.765)		(0.711)
NR(Others)			0.138			0.126		-0.00175
			(0.166)			(0.186)		(0.213)
NR(No. Ops.(t-1))				-			-	-
				1.487***			1.493***	0.904***
				(0.227)			(0.230)	(0.268)
NR(Descending finan awareness)					0.644**		0.415	0.181
					(0.299)		(0.295)	(0.349)
Online investment information = o,			-			-		-
Others = o,			-			-		-
Constant	-	-	-	-0.00696	-	-	-0.574	-1.395*
	1.702***	2.230***	2.759***		2.769***	2.976***		
	(0.415)	(0.501)	(0.566)	(0.490)	(0.439)	(0.639)	(0.500)	(0.715)
Observations	2,318	2,318	2,318	2,318	2,318	2,318	2,318	2,318
ll	-1224	-844.5	-1081	-821.4	-1162	-799.3	-810.6	-682.6
r2_p	0.0944	0.375	0.200	0.392	0.140	0.408	0.400	0.495

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

VARIABLES	(1) sBaselineE Exp. Ret.	(2) sE1DK Exp. Ret.	(3) sE2DK Exp. Ret.	(4) sE3DK Exp. Ret.	(5) sE4DK Exp. Ret.
Perceived % Pop. Part.		0.000293**		0.000418**	0.000643**
		(0.000134)		(0.000206)	(0.000276)
Perceived % Pop. Informed			0.000190*	-0.000149	-1.65e-06
			(0.000108)	(0.000178)	(0.000205)
(% Pop. Part.)x(% Pop. Informed)					-8.60e-06
					(6.60e-06)
Age	0.000791	0.000728	0.000723	0.000718	0.000741
	(0.000618)	(0.000618)	(0.000616)	(0.000617)	(0.000617)
Age Sq.	-6.90e-06	-5.97e-06	-6.06e-06	-5.90e-06	-6.06e-06
	(5.93e-06)	(5.93e-06)	(5.91e-06)	(5.92e-06)	(5.92e-06)
Male	0.0133***	0.0129***	0.0131***	0.0128***	0.0124***
	(0.00355)	(0.00357)	(0.00356)	(0.00357)	(0.00359)
Married	-0.00936**	-0.00914**	-0.00925**	-0.00914**	-0.00924**
	(0.00407)	(0.00407)	(0.00407)	(0.00406)	(0.00407)
CARA	-0.000840**	-0.000745*	-0.000764*	-0.000758*	-0.000757*
	(0.000387)	(0.000390)	(0.000390)	(0.000389)	(0.000389)
High school	0.00587	0.00679	0.00679	0.00655	0.00655
	(0.0140)	(0.0141)	(0.0140)	(0.0141)	(0.0141)
Technical/Professional	-0.00532	-0.00510	-0.00466	-0.00537	-0.00536
	(0.0106)	(0.0107)	(0.0106)	(0.0107)	(0.0107)
College or more	0.00516	0.00421	0.00529	0.00406	0.00352
	(0.0108)	(0.0109)	(0.0108)	(0.0109)	(0.0109)
NR(CARA)	-0.0753***	-0.0707***	-0.0729***	-0.0726***	-0.0726***
	(0.0241)	(0.0243)	(0.0241)	(0.0241)	(0.0241)
NR(Per.% Pop. Part.)		-0.00331		-0.00387	0.0346
		(0.00410)		(0.00412)	(0.0280)
NR(Per.% Pop. Inf.)			0.0329	0.0354	0.0389
			(0.0242)	(0.0243)	(0.0264)
NR(% Pop. Part.)x(% Pop. Informed)					-0.0375
					(0.0283)
Children at Home>0	0.00362	0.00345	0.00352	0.00379	0.00388
	(0.00454)	(0.00454)	(0.00453)	(0.00453)	(0.00453)
Constant	0.0269	0.0224	0.0225	0.0239	0.0222
	(0.0229)	(0.0229)	(0.0229)	(0.0229)	(0.0229)
Observations	2,535	2,535	2,535	2,535	2,535
R-squared	0.019	0.022	0.020	0.023	0.024

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure A.28: Regress expectations of returns (continuous variable) on coordination motives and household characteristics under an OLS technique.

VARIABLES	(1) BaseInf Perceived Ret.	(2) InfP Perceived Ret.	(3) InfI Perceived Ret.	(4) InfPI Perceived Ret.	(5) InfPIInt Perceived Ret.
Perceived % Pop. Part.		-0.000238		0.000214	7.82e-06
		(0.000260)		(0.000365)	(0.000525)
Perceived % Pop. Informed			-0.000495**	-0.000624*	-0.000793**
			(0.000233)	(0.000327)	(0.000398)
(% Pop. Part.)x(% Pop. Informed)					6.59e-06
					(1.02e-05)
Age	0.000776	0.000712	0.000675	0.000663	0.000652
	(0.000848)	(0.000851)	(0.000851)	(0.000854)	(0.000854)
Age Sq.	-1.50e-06	-4.21e-07	-1.34e-07	3.79e-09	1.08e-07
	(8.09e-06)	(8.13e-06)	(8.13e-06)	(8.15e-06)	(8.15e-06)
Male	0.0314***	0.0299***	0.0295***	0.0294***	0.0294***
	(0.00495)	(0.00500)	(0.00494)	(0.00497)	(0.00497)
Married	-0.00946*	-0.00981*	-0.00945*	-0.00938*	-0.00943*
	(0.00517)	(0.00518)	(0.00516)	(0.00516)	(0.00517)
CARA	-0.00210***	-0.00200***	-0.00199***	-0.00201***	-0.00202***
	(0.000597)	(0.000596)	(0.000595)	(0.000596)	(0.000596)
High school	-0.0163	-0.0160	-0.0157	-0.0159	-0.0159
	(0.0179)	(0.0179)	(0.0179)	(0.0180)	(0.0179)
Technical/Professional	-0.00752	-0.00788	-0.00741	-0.00748	-0.00757
	(0.0148)	(0.0148)	(0.0148)	(0.0149)	(0.0149)
College or more	0.0156	0.0125	0.0125	0.0124	0.0124
	(0.0152)	(0.0152)	(0.0152)	(0.0154)	(0.0154)
NR(CARA)	-0.109***	-0.104***	-0.103***	-0.105***	-0.105***
	(0.0350)	(0.0337)	(0.0343)	(0.0335)	(0.0336)
NR(Per.% Pop. Part.)		0.00823		0.0547	0.0537
		(0.0809)		(0.0925)	(0.0925)
NR(Per.% Pop. Inf.)			-0.0222	-0.0514	-0.0546
			(0.0497)	(0.0455)	(0.0461)
DK(Per.% Pop. Part.)		-0.0207**		0.000760	-0.000252
		(0.00850)		(0.0166)	(0.0166)
DK(Per.% Pop. Inf.)			-0.0280***	-0.0275*	-0.0305*
			(0.00841)	(0.0165)	(0.0178)
NR(% Pop. Part.)x(% Pop. Informed) = o,					-
Constant	0.0676**	0.0811**	0.0883***	0.0879***	0.0928***
	(0.0327)	(0.0338)	(0.0335)	(0.0339)	(0.0349)
Observations	2,328	2,328	2,328	2,328	2,328
R-squared	0.040	0.044	0.045	0.046	0.046

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure A.29: Regress perceptions of returns (continuous variable) on coordination motives and household characteristics under an OLS technique.

VARIABLES	(1) InfPIInt Perceived Ret.	(2) InfPIIntTr Perceived Ret.	(3) InfPIIntSo Perceived Ret.	(4) InfPIIntIn Perceived Ret.	(5) InfPIIntKn Perceived Ret.	(6) InfPIIntTS Perceived Ret.	(7) InfPIIntIK Perceived Ret.	(8) InfPIIntAll Perceived Ret.
Perceived % Pop. Part.	7.82e-06	-3.32e-05	-3.37e-06	-3.88e-05	-9.91e-05	-3.69e-05	-0.000128	-9.40e-05
	(0.000525)	(0.000516)	(0.000518)	(0.000511)	(0.000524)	(0.000511)	(0.000516)	(0.000515)
Perceived % Pop. Informed	-0.000793**	-0.000756*	-0.000718*	-0.000708*	-0.000814**	-0.000699*	-0.000751*	-0.000708*
	(0.000398)	(0.000399)	(0.000401)	(0.000398)	(0.000395)	(0.000400)	(0.000395)	(0.000398)
(% Pop. Part.)x(% Pop. Informed)	6.59e-06	6.82e-06	7.10e-06	6.12e-06	8.54e-06	7.21e-06	8.06e-06	8.31e-06
	(1.02e-05)	(1.01e-05)	(1.03e-05)	(1.01e-05)	(1.02e-05)	(1.02e-05)	(1.02e-05)	(1.03e-05)
CARA	-0.00202***	-0.00186***	-0.00158***	-0.00179***	-0.00165***	-0.00152***	-0.00155***	-0.00139**
	(0.000596)	(0.000589)	(0.000582)	(0.000590)	(0.000588)	(0.000577)	(0.000587)	(0.000573)
Age	0.000652	0.000396	-5.06e-05	0.000297	0.000215	-1.56e-05	7.53e-05	-3.73e-05
	(0.000854)	(0.000872)	(0.000891)	(0.000851)	(0.000856)	(0.000898)	(0.000853)	(0.000897)
Age Sq.	1.08e-07	1.12e-06	5.38e-06	1.93e-06	3.05e-06	4.43e-06	3.70e-06	4.32e-06
	(8.15e-06)	(8.34e-06)	(8.49e-06)	(8.13e-06)	(8.21e-06)	(8.57e-06)	(8.18e-06)	(8.57e-06)
High school	-0.0159	-0.0141	-0.0156	-0.0134	-0.0119	-0.0148	-0.0107	-0.0127
	(0.0179)	(0.0181)	(0.0184)	(0.0180)	(0.0182)	(0.0185)	(0.0182)	(0.0185)
Technical/Professional	-0.00757	-0.0101	-0.00961	-0.00779	-0.00694	-0.0108	-0.00681	-0.00942
	(0.0149)	(0.0151)	(0.0153)	(0.0149)	(0.0149)	(0.0153)	(0.0149)	(0.0153)
College or more	0.0124	0.00508	0.00291	0.00861	0.00648	-0.000809	0.00543	-0.00103
	(0.0154)	(0.0156)	(0.0158)	(0.0154)	(0.0155)	(0.0159)	(0.0155)	(0.0158)
Male	0.0294***	0.0270***	0.0258***	0.0272***	0.0241***	0.0244***	0.0233***	0.0221***
	(0.00497)	(0.00496)	(0.00500)	(0.00498)	(0.00495)	(0.00498)	(0.00496)	(0.00500)
Married	-0.00943*	-0.0102**	-0.01000*	-0.00907*	-0.00942*	-0.00996*	-0.00891*	-0.00930*
	(0.00517)	(0.00518)	(0.00515)	(0.00516)	(0.00517)	(0.00516)	(0.00517)	(0.00519)
Not investing: NSP		-0.0306***				-0.0199**		-0.0131
		(0.00981)				(0.00973)		(0.0121)
Descending adviser trust		0.00350*				0.00340		0.00234
		(0.00194)				(0.00229)		(0.00230)
Trust others		-0.0104**				-0.00999**		-0.00949*
		(0.00501)				(0.00502)		(0.00503)
Friends & relations			-0.000875			-0.000760		-0.00155
			(0.00608)			(0.00604)		(0.00605)
Family			-0.000188			-0.000455		-0.000920
			(0.00557)			(0.00552)		(0.00548)
Financial advisers			-0.00472			-0.00592		-0.00685
			(0.00412)			(0.00417)		(0.00439)
General media			-0.0208**			-0.0189**		-0.0151
			(0.00938)			(0.00938)		(0.00955)
Specialist media			0.0252***			0.0224***		0.0179***
			(0.00624)			(0.00624)		(0.00625)
Friends & relations: Finan decision - rely on...?			0.00144			0.00191		0.00231
			(0.00475)			(0.00471)		(0.00474)
Family: Finan decision - rely on...?			-0.00266			-0.00223		-0.00296
			(0.00399)			(0.00397)		(0.00401)
Banker/financial adviser: Finan decision - rely on...?			-0.00374			-0.00559		-0.00430
			(0.00352)			(0.00366)		(0.00365)
Financial knowledge: Finan decision - rely on...?			0.00946***			0.00843**		0.00413
			(0.00344)			(0.00347)		(0.00369)
Media: Finan decision - rely on...?			0.0105**			0.0106**		0.0101*
			(0.00508)			(0.00510)		(0.00517)
Written press			0.00793			0.00728		0.00511
			(0.00655)			(0.00655)		(0.00668)
Audiovisual			0.00336			0.00367		0.00346
			(0.00699)			(0.00697)		(0.00702)
Online social websites			0.00357			0.00437		0.00494
			(0.00632)			(0.00634)		(0.00637)
Finance websites			0.00217			0.00177		0.00178

Figure A.30: Regress perceptions of returns (continuous variable) on coordination motives, household characteristics and information variables under an OLS technique.

Finance websites			0.00217			0.00177		0.00178
			(0.00696)			(0.00692)		(0.00707)
Financial authorities websites			-0.00686			-0.00747		-0.00843
			(0.00908)			(0.00919)		(0.00928)
No. Operations(t-1)				0.0114*			0.00741	0.00254
				(0.00622)			(0.00621)	(0.00637)
Descending finan awareness					0.00725**		0.00614**	0.00412
					(0.00306)		(0.00309)	(0.00316)
Understand financial press					-0.0174***		-0.0165***	-0.0130***
					(0.00342)		(0.00343)	(0.00357)
Discuss with adviser					0.00128		0.000921	-0.00200
					(0.00396)		(0.00398)	(0.00418)
Read commercial advert					0.00542		0.00548	0.00583
					(0.00383)		(0.00383)	(0.00390)
Evaluate opportunity profitability					-0.00148		-0.00108	-0.000901
					(0.00536)		(0.00538)	(0.00540)
Evaluate own investments risk					0.00134		0.000950	0.000354
					(0.00537)		(0.00536)	(0.00534)
Choose your investments					-0.000252		5.14e-05	0.000579
					(0.00441)		(0.00443)	(0.00454)
NR(CARA)	-0.105***	-0.104***	-0.0861***	-0.0955***	-0.0860**	-0.0856***	-0.0823**	-0.0740**
	(0.0336)	(0.0334)	(0.0333)	(0.0336)	(0.0356)	(0.0331)	(0.0356)	(0.0337)
NR(Assets)		-0.00766	-0.00178	-0.00727	-0.00340	-0.00208	-0.00322	-0.000564
		(0.0123)	(0.0120)	(0.0123)	(0.0125)	(0.0120)	(0.0125)	(0.0123)
NR(Income)		0.0209	0.0139	0.0174	0.0109	0.0147	0.0115	0.0104
		(0.0415)	(0.0417)	(0.0417)	(0.0413)	(0.0415)	(0.0416)	(0.0414)
NR(Saving)		0.0259	0.0239	0.0224	0.0192	0.0268	0.0191	0.0247
		(0.0263)	(0.0264)	(0.0263)	(0.0254)	(0.0263)	(0.0256)	(0.0261)
NR(Per.% Pop. Part.)	0.0537	0.0543	0.0389	0.0505	0.0373	0.0410	0.0364	0.0283
	(0.0925)	(0.0878)	(0.0877)	(0.0881)	(0.0906)	(0.0875)	(0.0905)	(0.0895)
NR(Per.% Pop. Inf.)	-0.0546	-0.0479	-0.0346	-0.0457	-0.0356	-0.0309	-0.0310	-0.0200
	(0.0461)	(0.0456)	(0.0473)	(0.0450)	(0.0438)	(0.0477)	(0.0432)	(0.0448)
DK(Per.% Pop. Part.)	-0.000252	0.00220	0.00329	0.00203	0.00294	0.00461	0.00338	0.00501
	(0.0166)	(0.0166)	(0.0163)	(0.0164)	(0.0164)	(0.0163)	(0.0164)	(0.0164)
DK(Per.% Pop. Inf.)	-0.0305*	-0.0285	-0.0213	-0.0278	-0.0253	-0.0213	-0.0239	-0.0198
	(0.0178)	(0.0178)	(0.0178)	(0.0178)	(0.0176)	(0.0178)	(0.0176)	(0.0178)
NR(Not investing: NSP)		0.0426***				0.0290**		0.0195
		(0.0124)				(0.0123)		(0.0149)
NR(Descending adviser trust)		-0.0495				-0.0398		-0.0401
		(0.0665)				(0.0681)		(0.0678)
NR(Trust others)		-0.00154				0.00591		0.00525
		(0.0290)				(0.0282)		(0.0277)
NR(Friends & relations)			-0.0179			-0.0191		-0.0236
			(0.0275)			(0.0274)		(0.0282)
NR(Family)			0.0306			0.0312		0.0339
			(0.0445)			(0.0445)		(0.0450)
NR(Financial advisers)			-0.0141			-0.0185		-0.0268
			(0.0350)			(0.0347)		(0.0359)
NR(General media)			-0.0911**			-0.0900**		-0.0809**
			(0.0398)			(0.0401)		(0.0408)
NR(Specialist media)			0.107***			0.104**		0.106**
			(0.0407)			(0.0413)		(0.0421)
NR(Friend & relations: Finan decision - rely on...?)			0.00323			0.00563		0.00280
			(0.0356)			(0.0364)		(0.0372)
NR(Family: Finan decision - rely on...?)			-0.0654*			-0.0659*		-0.0667*
			(0.0376)			(0.0382)		(0.0391)
NR(Banker/financial adviser: Finan decision - rely on...?)			0.0263			0.0280		0.0306
			(0.0361)			(0.0381)		(0.0386)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
NR(Financial knowledge: Finan decision - rely on...?)			0.0182			0.0172		0.00735
			(0.0351)			(0.0357)		(0.0359)
NR(Media: Finan decision - rely on...?)			0.0435			0.0412		0.0395
			(0.0344)			(0.0351)		(0.0349)
NR(Written press)			0.118***			0.107**		0.108**
			(0.0453)			(0.0458)		(0.0477)
NR(Audiovisual)			-0.0525			-0.0486		-0.0536
			(0.0466)			(0.0471)		(0.0492)
NR(Online soc websites)			0.0567			0.0614		0.0583
			(0.0653)			(0.0651)		(0.0657)
NR(Finance websites)			-0.0935			-0.0880		-0.0788
			(0.0755)			(0.0733)		(0.0764)
NR(Finan authorities websites)			-0.0529			-0.0569		-0.0582
			(0.0442)			(0.0449)		(0.0453)
NR(Online investment info)			0.0531			0.0544		0.0524
			(0.0549)			(0.0542)		(0.0547)
NR(Others)			-0.00238			-0.00252		-0.00353
			(0.0147)			(0.0147)		(0.0148)
NR(No. Ops.(t-1))				0.00610			0.00687	0.00333
				(0.0186)			(0.0186)	(0.0195)
NR(Descending finan awareness)					-0.00860		-0.0101	-0.0186
					(0.0234)		(0.0233)	(0.0238)
NR(Understand financial press)					0.0357**		0.0333*	0.0210
					(0.0181)		(0.0178)	(0.0216)
NR(Discuss with adviser)					0.0283		0.0307	0.00259
					(0.0224)		(0.0225)	(0.0301)
NR(Read commercial advert)					-0.0589		-0.0629*	-0.0386
					(0.0387)		(0.0379)	(0.0440)
NR(Evaluate opportunity profitability)					0.01000		0.0118	0.0273
					(0.0524)		(0.0515)	(0.0449)
NR(Evaluate own investments risk)					-0.0523		-0.0507	-0.0608*
					(0.0347)		(0.0343)	(0.0347)
NR(Choose your investments)					0.0647		0.0648	0.0640
					(0.0403)		(0.0400)	(0.0462)
Online investment information = o,			-			-		-
Others = o,			-			-		-
NR(% Pop. Part.)x(% Pop. Informed) = o,	-	-	-	-	-	-	-	-
Constant	0.0928***	0.101***	0.0514	0.0964***	0.113***	0.0630	0.111***	0.0972**
	(0.0349)	(0.0355)	(0.0415)	(0.0372)	(0.0403)	(0.0417)	(0.0424)	(0.0495)
Observations	2,328	2,328	2,328	2,328	2,328	2,328	2,328	2,328
R-squared	0.046	0.057	0.084	0.055	0.073	0.089	0.076	0.101

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

A.3.3 Questionnaire

This part of the appendix presents the questions posed to participants in the questionnaire and links them to the variables used in the econometric analysis. The text is translated into English from French.

1. Participation in the stock market, Question C3

   C3. Amongst the types of financial investments following, which ones do you or a member of your household own?

   Respondents are given a list of options that can be identified as: no investments, direct stock market investments or indirect stock market investments.

2. Perceptions of the stock market, Question C42

   C42. Over the last 3 years, do you think that the stock market... -For each category write down the probability of occurrence assigning a value between 0 and 100. The sum of all your answers must be equal to 100-:

   -... has increased by more than 25%

   -... has increased by 10 to 25%

   -... has increased by less than 10%

   -... has remained the same

   -... has decreased by less than 10%

   -... has decreased by 10 to 25%

   -... has decreased by more than 25%

3. Expectations of the stock market, Question C39

   C39. In 5 years' time, do you think that the stock market... -For each category write down the probability of occurrence assigning a value between 0 and 100. The sum of all your answers must be equal to 100-:

   -... will have increased by more than 25%

-... will have increased by 10 to 25%

-... will have increased by less than 10%

-... will have remained the same

-... will have decreased by less than 10%

-... will have decreased by 10 to 25%

-... will have decreased by more than 25%

4. Perceptions of others' participation, Question C35

   C35. In your opinion, out of 100, how many people in the French population invest in the stock market is (directly or via mutual funds)? If you do not know then select "I don't know"

5. Perceptions of others' information, Question C36

   C36. In your opinion, out of 100, how many people in the French population are informed regarding the evolution of the stock market? If you do not know then select "I don't know"

6. Risk aversion, Question C44

   C44. You can invest in a financial product for which there is a one in two chance of gaining €5,000 and a one in two chance of losing all the capital invested. What is the maximum amount you would pay for this financial product? Note the sum in euros.

Trust

7. Not investing, Question C9

   C9. From the following reasons which best describe why you have never invested (or no longer invest) in the stock market? **Response given:** I don't know.

8. Not investing, Question C53

   C53. When your banker or financial adviser recommends an investment to you:

   -... You understand how it fits with your needs.

-... You don't understand at all but you trust him.

-... You don't understand at all and you don't trust him.

-... Not applicable (no recommendations received).

9. Trust others, Question J20

   J20. In general, do you trust others? Yes/No.

Sources of information

10. Gather information from who or what, Question C45

    C45. Indicate for each of the following sources of information how often you consult it for your own information on financial investments.

    -... Friends & relations:
    often, sometimes, never or no financial investments.

    -... Family:
    often, sometimes, never or no financial investments.

    -... Financial advisers:
    often, sometimes, never or no financial investments.

    -... Generalist media:
    often, sometimes, never or no financial investments.

    -... Specialist media:
    often, sometimes, never or no financial investments.

11. Rely on who or what, Question C52

    C52. When you take a financial decision, who or what do you rely on to make your choice?

    -... Friends & relations:
    often, sometimes, never or not applicable.

-... Family:

often, sometimes, never or not applicable.

-... Banker/financial adviser:

often, sometimes, never or not applicable.

-... Your own financial knowledge:

often, sometimes, never or not applicable.

-... The media:

often, sometimes, never or not applicable.

12. Media consulted, Question C46

C46. What types of media do you consult?

-... Written press:

at least once per week, at least once per month, less often, never.

-... Audiovisual (radio, television):

at least once per week, at least once per month, less often, never.

-... Internet: social networks, forums, discussion groups:

at least once per week, at least once per month, less often, never.

-... Internet: financial establishment websites:

at least once per week, at least once per month, less often, never.

-... Internet: websites of institutions (central bank):

at least once per week, at least once per month, less often, never.

-... Internet: websites of investment opportunities:

at least once per week, at least once per month, less often, never.

-... Others:

at least once per week, at least once per month, less often, never.

Inertia

13. Recent stock market transactions, Question C10

    C10. In the course of the last 2 years, about how many buying or selling actions have you (or another member of your household or an intermediary on your behalf) carried out in the stock market?

    - No actions.

    - 1 or 2 actions.

    - 3 to 5 actions.

    - 6 or more actions.

Knowledge Level

14. Financial culture, Question C47

    C47. Would you say that you financial culture is...?

    - high.

    - average.

    - weak.

    - very weak.

    - inexistant.

15. Financial knowledge, Question C48

    C48. Do you have the impression that you financial knowledge is sufficient for...?

    -... Understanding the financial press:
    completely agree, somewhat agree, somewhat disagree, completely disagree, n/a.

    -... Discussing with a financial adviser:
    completely agree, somewhat agree, somewhat disagree, completely disagree, n/a.

-... Reading a commercial document of information about an investment opportunity:
completely agree, somewhat agree, somewhat disagree, completely disagree, n/a.

-... Evaluating the profitability of your investments:
completely agree, somewhat agree, somewhat disagree, completely disagree, n/a.

-... Evaluating the risk of your investments:
completely agree, somewhat agree, somewhat disagree, completely disagree, n/a.

-... Choosing your financial investments:
completely agree, somewhat agree, somewhat disagree, completely disagree, n/a.

APPENDIX B

References

ADAM, K. and MARCET, A., 2011. Internal Rationality, Imperfect Market Knowledge and Asset Prices. *Journal of Economic Theory* , 146(3), 1224-1252.

ADMATI, A., 1985. A Noisy Rational Expectations Equilibrium for Multi-Asset securities Markets. *Econometrica* , 53(3), 629-657.

ARMANTIER, O., NELSON, S., TOPA, G., VAN DER KLAAUW, W. and ZAFAR, B., 2014. The Price is Right: Updating Inflation Expectations in a Randomized Price Information Experiment. The Review of Economics and Statistics, forthcoming.

BLACK, F., JENSEN, M. and SCHOLES, M., 1972. The Capital Asset Pricing Model: Some Empirical Tests. *In:* JENSEN, M., eds. *Studies in the Theory of Capital Markets*. New York: Praeger Publishers, 79-121.

ARRONDEL, L., CALVO-PARDO, C. and TAS, D., 2015. Subjective Return Expectations, Information and Stock Market Participation: Evidence from France. Working Paper, Southampton DP in Economics and Econometrics No. 141.

ARRONDEL, L., CALVO-PARDO, H., and MASSON, A., 2015. Subjective Expectations and Perceptions of Stock Market Returns, *mimeo*.

ARRONDEL, L., and MASSON, A., 2011. French Savers in the Economic Crisis: What has Changed?, PSE-Jourdan WP.

ARTHUR, W. B., 2006, Out-of-Equilibrium Economics and Agent-Based Modeling, in JUDD, K. and TESFATSION, L., eds, *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, Elsevier/North-Holland Publishers.

BRANDT, M., 2010. Portfolio Choice Problems. *In:* AIT-SAHALIA, Y. and HANSEN, L., eds. *Handbook of Financial Econometrics, Volume 1*. Amsterdam: Elsevier Science, 269-330.

BRUGGISSER, D., 2011. Parameter Uncertainty and Learning in Dynamic Financial Decisions. [Online] Electronic Dissertations at the University of St. Gallen. Available from:
www1.unisg.ch/www/edis.nsf/SysLkpByIdentifier/3874/$FILE/dis3874.pdf. [Accessed 22nd November 2014].

CAMPBELL, J., Y., 1996. Understanding Risk and Return *Journal of Political Economy*, 104(2), 298-345.

CAMPBELL, J., Y. and VICEIRA, L., 1999. Consumption and Portfolio Decisions When Expected Returns are Time Varying *Quarterly Journal of Economics*, 114(2), 433-495.

CAMPBELL, J. Y., 2006. Household Finance, *The Journal of Finance*, 61: 1553-1604.

CHEW, S.H. and EPSTEIN, L., 1989. The structure of Preferences and Attitudes towards the Timing of the Resolution of Uncertainty. *International Economic Review* , 30(1), 103-117.

COCHRANE, J. and HANSEN, L., 1992. Asset Pricing Explorations for Macroeconomics *NBER Macroeconomics Annual*. Cambridge, MA: MIT Press, 1992.

COIBION, O. and GORODNICHENKO, Y., 2012. What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1), 116-159.

COIBION, O., GORODNICHENKO, Y. and KUMAR, S, 2012. How Do Firms Form Their Expectations? New Survey Evidence, Working paper, NBER.

CURCURU, S., HEATON, J., LUCAS, D. and MOORE, D., 2010. Portfolio Choice

Problems. *In:* AIT-SAHALIA, Y. and HANSEN, L., eds. *Handbook of Financial Econometrics, Volume 1*. Amsterdam: Elsevier Science, 269.

DAVIS, S. and WILLEN, P., March 2000. Using Financial Assets to Hedge Labor Income Risk: Estimting the Benefits, Working paper.

DESGRANGES, G. and GAUTHIER, S., 2013. Asymmetric Information and Rationalizability, *Economic Theory* 54 (3):789-804.

DOMINITZ, J. and MANSKI, C., 2007. Expected Equity Returns and Portfolio Choice: Evidence from the Health and Retirement Study. *Journal of the European Economic Association*, 5(2-3), 369-379.

DOMINITZ, J. and MANSKI, C., 2011. Measuring and Interpreting Expectations of Equity Returns, *Journal of Applied Econometrics*, 26: 352–370.

FELDMAN, D., 2007. Incomplete information equilibria: Separation theorems and other myths. *Annals of Operational Research*, 151(1), 119-149.

GABAIX, X., 2016. Behavioural Macroeconomics Via Sparse Dynamic Programming, Working Paper, New York University.

GUESNERIE, R., (2001), (2005), Assessing Rational Expectations: 1- Sunspot multiplicity and Economic Fluctuations, 317 p.; 2-Eductive Stability in economics, 452 p., Cambridge: MIT Press.

GUISO, L., JAPPELLI, T. and TERLIZZESE, D., 1996. Income Risky, Borrowing Constraints, and Portfolio Choice. *The American Economic Review* , 86(1), 158-172.

GUISO, L. and PAIELLA, M., 2008. Risk Aversion, Wealth and Background Risk, *Journal of the European Economic Association* 6(6): 1109-1150.

HALL, R. E., 2010. Why Does the Economy Fall to Pieces after a Financial Crisis?, *Journal of Economic Perspectives*, 24(4): 3-20.

HANSEN, L.P. and SARGENT, T., 2012. Three Types of Ambiguity, *Journal of Monetary Economics* 59: 422-445.

HELLWIG, C. and VELDKAMP, L., 2009. Knowing What Others Know: Coordination Motives in Information Acquisition. *Review of Economic Studies* , 76, 223-251.

HOMMES, C, 2007. Bounded rationality and learning in complex markets, Working paper, CeNDEF Working Paper Universiteit van Amsterdam, Amsterdam: Universiteit van Amsterdam.

HOMMES, C., 2013. Reflexivity, expectations feedback and almost self-fulfilling equilibria: economic theory, empirical evidence and laboratory experiments, *Journal of Economic Methodology*, 20(4), 406-419.

HOMMES, C. and WAGENER, F., 2008. Complex Evolutionary Systems *In:* HENS, T. and SCHENK-HOPPÉ, K.R. eds. Behavioral Finance, in the *Handbook of Financial Markets: Dynamics and Evolution*, Elsevier.

HUDOMIET, P., KÉZDI, G. and WILLIS, R. J., 2011. Stock Market Crash and Expectations of American Households, *Journal of Applied Econometrics* 26(3): 393-415.

HURD, M. D., VAN ROOIJ, M. and WINTER, J., 2011. Stock Market Expectations of Dutch Households, *Journal of Applied Econometrics* 26(3): 416-436.

KEARNS, M., 2015. Why do individuals hold concentrated portfolios?, Working paper, University of Southampton.

LEBARON, B., 2013. Heterogeneous Gain Learning and long Swings in Asset Prices, *In:* FRYDMAN, R. and PHELPS, E. S., eds. *Rethinking Expectations: The Way Forward for Macroeconomics*, Princeton University Press.

LJUNGQVIST, L. and SARGENT, T. (2004). *Recursive Macroeconomic Theory*. United States of America: The MIT Press.

MANKIW, G. and REIS, R., 2007. Sticky Information in General Equilibrium. *Journal of the European Economic Association*, 5(2-3), 603-613.

MANSKI, C., 2004. Measuring Expectations, *Econometrica*, 72: 1329-76.

MARWOWITZ, H., 1952. Portfolio Selection. *The Journal of Finance*, 7(1), 77-91.

MEINHOLD, R. and SINGPURWALLA, N., 1983. Understanding the Kalman Filter. *The American Statistician*, 37(2), 123-127.

MERTON, R., 1969. Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *The Review of Economics and Statstics*, 51(3), 247-257.

MERTON, R., 1971. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory*, 3, 373-413.

MERTON, R., 1972. An analytic derivation of the efficient portfolio frontier. *Journal of Financial and Quantitative Analysis*, 7(4), 1851-1872.

MERTON, R., 1973. An intertemporal capital asset pricing model. *Econometrica*, 41, 867-887.

MICHAELIDES, A. and ZHANG, Y., 2015. Stock Market Mean Reversion and Portfolio Choice over the Life Cycle. *Journal of Financial and Quantitative Analysis*, forthcoming.

MINIACI, R., and WEBER, G., 2002, Econometric Issues in the Estimation of Household Portfolio Choice Models, *In:* GUISO, L., HALIASSOS, M. and JAPPELLI, T. eds. Household Portfolios, Cambridge: MIT Press.

MORRIS, S. and SHIN, H. S., 1998, Unique Equilibrium in a Model of Self-Fullfilling Currency Attacks, *American Economic Review* 88(3): 587-97.

MORRIS, S. and SHIN, H. S., 2001. Rethinking Multiple Equilibria in Macroeconomic Modeling, *NBER Macroeconomics Annual 2000* 15: 139-182.

MORRIS, S. and SHIN, H. S., 2002, The Social Value of Public Information, *American Economic Review* 92: 1521-34.

ORPHANIDES, A. and WILLIAMS, J., 2005. Inflation scares and forecast-based monetary policy. *Review of Economic Dynamics*, 8(2), 498-527.

PASRICHA, G., 2006. Kalman Filter and its Economic Applications. MPRA Paper 22734, University of Munich, Germany.

PESARAN, M. H., and WEALE, M., 2006. Survey Expectations, *In:* ELLIOTT, G., GRANGER, C.W.J. and TIMMERMANN, A., eds. *Handbook of Economic Forecasting*. North-Holland.

PHELPS, E., 1962. The Accumulation of Risky Capital: A Sequential Utility Analysis. *Econometrica*, 30(4), 729-743.

SAMUELSON, P., 1969. Lifetime Portfolio Selection by Dynamic Stochastic Programming. *The Review of Economics and Statistics*, 51(3), 239-246.

SIMS, C., 2003. Implications of rational inattention. *Journal of Monetary Economics* , 50, 665-690.

STOKEY, N. and LUCAS, R. (1989). *Recursive Methods in Economic Dynamics*. United States of America: Harvard University Press.

VAN NIEUWERBURGH, S. and VELDKAMP, L., 2010. Information Acquisition and Under-Diversification. *Review of Economic Studies* , 77(2), 779-805.

VELDKAMP, L. (2011). *Information Choice in Macroeconomics and Finance*. Woodstock, Oxfordshire, United Kingdom: Princeton University Press.

VICEIRA, L., 2001. Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income. *The Journal of Finance* , 56(2), 433-470.

VISSING-JORGENSEN, A., 2004. Perspectives on Behavioural Finance: Does Irrationality Disappear with Wealth? Evidence from Expectations and Actions, *In:* GERTLER, M. and ROGOFF, K., eds. *The NBER Macroeconomics Annual 2003*. Cambridge: MIT Press.