SYRTO

SYstemic Risk TOmography
Signals, Measurements, Transmission Channels, and Policy Interventions

Rational learning for risk-averse investors by conditioning on behavioral choices

Michele Costola, Massimiliano Caporin

SYRTO WORKING PAPER SERIES
Working paper n. 15 | 2015

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement n° 320270.

Rational learning for risk-averse investors by conditioning on behavioral choices[☆]

Michele Costola[a,*], Massimiliano Caporin[b]

[a] *University Ca' Foscari of Venice, Department of Economics, Cannaregio 873, Venezia, Italy.*
[b] *University of Padova, Department of Economics and Management "Marco Fanno", Via del Santo 22, Padova, Italy.*

Abstract

We present a rational learner agent, which considers the information coming from a behavioral counterpart during the allocation process. The learner agent adopts a herding behaviour by conditioning her choice on the selection of the portfolio's constituents. The considered framework has therefore two types of agents with two different utility functions: the rational agent with a hyperbolic absolute risk aversion (HARA) utility function and the other one with a general behavioral utility function. We use the concept of performance measure related to utility functions to define agents' preferences: the higher the measure, the higher the expected utility of a given asset. The rational learner agent updates her information in a Bayesian manner similarly to the Black-Litterman model, which makes use of a weighting factor in blending the two components. We support our methodological framework with an empirical analysis including all the assets present in the NASDAQ and NYSE stock exchange from September 1977 to December 2014.
Keywords: learner agent, investment decision, behavioral agents, Bayesian updating.
JEL-Classification: G110, G140, G150, G170.

1. Introduction

The main goal of decision theory is to determine how individuals should decide and to explain how they actually decide. In particular, while the prescriptive approach indicates how a rational choice should be made, the descriptive one models how the decisions are effectively made. By focusing on the latter approach, it is possible to observe individuals that systematically deviate from what the prescriptive method defines as rational: this approach is called behavioral.
According to the efficient market hypothesis, if agents are rational and there are no frictions in the market, the security's price will reflect all the available information, and it will be equal to its fundamental value without allowing for arbitrage activities. In other words, if the market is

[☆] Both authors acknowledge financial support from the European Union, Seventh Framework Program FP7/2007-2013 under grant agreement SYRTO-SSH-2012-320270. Both authors thank for financial support the Global Risk Institute in Financial Services and the Louis Bachelier Institute. The first author acknowledges support from CREATES - Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. The second author acknowledges financial support from the University of Padova. The authors gratefully acknowledge the comments provided by participants at the Computational Statistics & Data Analysis Conference (2010), the Association of Southern European Economic Theorists Conference (2012), the Workshop of Quantitative Finance (2013), and the Italian Congress of Econometrics and Empirical Economics Conference (2013).

[*] Corresponding author: Michele Costola. University Ca' Foscari of Venice, Department of Economics, Cannaregio 873, Venezia, Italy. michele.costola@unipd.it

Email addresses: `michele.costola@unive.it` (Michele Costola), `massimiliano.caporin@unipd.it` (Massimiliano Caporin)

efficient, a profitable trading strategy that would allow obtaining risk-adjusted excess returns above the market returns would not exist. On the basis of this assumption, classical financial economists such as Friedman (1953) assert that pricing anomalies cannot exist,[1] because if there would be some mispricing, this would imply a *de facto* arbitrage that rational investors would immediately grasp. Consequently, the mismatch would instantaneously disappear.

However, Lamont and Thaler (2003), amongst many others, have shown several empirical violations of the law of one price, proving the existence of arbitrage opportunities in the stock market. On the other hand, Malkiel (2003) argues the possibility that some investors are less rational than others, and thus, pricing irregularities and predictable patterns could occur in the market. Nevertheless, these patterns of irrationalities in the pricing are unlikely to continue and, at the end, they will not reward a significant risk–adjusted excess return.[2] Moreover, as reported by Hommes (2006), in an efficient market, assuming that all agents are rational and have a perfect common knowledge of all the available information, there should be no trade.

Summing up, the classical theory asserts that the absence of an arbitrage opportunity ensures that the prices are correct, and then the market is efficient. Conversely, according to the behavioral approach, deviations from the fundamental value are due to the presence of some agents that do not act in a fully rational way. From a different point of view, we might think that a mispricing could be present in the market, but its search could be too complicated for a rational investor and unattractive because of implementation costs (D'Avolio, 2002).

Both types of agents, rational and behavioral, can be present in a market. This calls for appropriate modeling framework, allowing for heterogeneous agents, an area extensively discussed in the economic literature (see Hommes, 2006, for a complete survey). The different types of agents are usually distinguished on the basis of their expectations about the future asset returns. De Long et al. (1990) differentiate noise traders from sophisticated traders. The first ones (i.e., as technical analysts and stock brokers) incorrectly rely on their information. Sophisticated traders, instead, exploit these false perceptions by adopting a herding or contrarian behaviour.

Zeman (1974) introduces a fundamentalist versus a chartist model. Fundamentalists trade on the basis of the market fundamentals and economic factors, while chartists base their trades on observed historical patterns in past prices. In Grossman and Stiglitz (1980), agents are divided into informed and uninformed. Since information is costly, prices cannot perfectly reflect all the available information in the market. The purpose of heterogeneous agent models is to explain stylized facts observed in financial markets, such as the random walk evidences, the absence on autocorrelations of asset prices, the fat tailed distribution of returns, and the well-know long–range volatility clustering (i.e., slow decay of autocorrelation of squared returns).

In this paper, we consider a novel framework for heterogeneous agents: a risk-averse agent equipped with a hyperbolic absolute risk aversion (HARA) and a behavioral counterpart who is endowed with a piece-wise linear plus power utility function. The aim is to propose a rational learning model where the HARA investor, in making her investment choices, considers the information coming from the behavioral counterpart. We define rational learning as the process undertaken through Bayesian updating of the prior beliefs provided by HARA agent's utility function given the presence of the

[1] A well–know story on the market efficiency tells about a professor and her student walking on the street, where at some point they find \$20 on the ground. The professor stops the student from picking up the bill by telling him that if it was really a \$20 note, it wouldn't be there anymore because someone else would already picked it up by somebody else.

[2] Market bubbles have also been considered by many economists as proof of some market irrationality, e.g Shiller(2008), while other, such as Garber (1990) analyzed the market bubbles, providing a fundamental explanation.

behavioral counterpart.
The main goal is to investigate this component's effect in terms of utility function on asset evaluation during the asset selection process. In this respect, we use the concept of performance measure, as derived from a utility function, where the higher the measure, the higher the expected utility provided by a given asset. In order to maintain a coherence between the two agent's views over the assets, we consider the generalized Sharpe ratio of Zakamouline and Koekebakker (2009b), as the benchmark performance measure for a rational investor, while for the behavioral agent we use the Z-ratio developed by Zakamouline (2011), starting from a general behavioral utility function. The measures proposed in the mentioned papers have been both obtained by following and exploiting the *maximum principle* approach introduced by Pedersen and Satchell (2002). In that paper, the authors define the optimal allocation between a risky and a risk-free asset in a single-period horizon. The solution of this allocation, which provides the maximum expected utility, is an increasing function of a quantity that can be viewed as a performance measure. Therefore, the maximization of the performance measure is equivalent to the maximization of the utility function.
Following the Bayesian approach, the model we introduce, used by the rational investor to blend the two different evaluations (described by the performance measures), is analogue to the approach followed by Black and Litterman (1992). From the the rational investor's perspective, the prior evaluation represents her view, while the conditional part represents the *behavioral* component. Further, the posterior provides the aggregated expectation according to the relevance given to the behavioral information. In our model's application, the rational learner adopts a herding behaviour and test if conditioning her choice towards a behavioral direction improves the selection amongst the assets in terms of cumulative returns and other ex-post performance evaluation criteria.[3] In other words, if the rational investor's choice is influenced, up to a certain degree, by the views of other type of agents present in the market, we deduce the rational investor is acting in a sophisticated way. In fact, the rational agent implicitly considers the aggregated evaluation coming from the different utility functions as the best way to select amongst the assets for the next period. Beside introducing out heterogeneous agent's framework and the approach for blending rational and behavioral views, we also introduce a criteria to evaluate the relevance of the behavioral views. Further, assuming the presence of different behavioral agents in the market, each characterized by different designs of the utility function, we also provide a methodology to determine if one type of agent has prevailing views in the blending process. Finally, we set an optimizing criterion for the weighting function of the behavioral component in absolute (weight on rational) and relative (weight of each behavioral agent type) terms.
We support our methodological contribution with an empirical example. We focus on the weekly data for all the stocks present in the NASDAQ and NYSE stock exchange from September 1972 to December 2014 (including also dead assets to avoid survivorship biases). Our findings show that the heterogeneous agents model is able to promptly react at the market momentum, providing an improvement in the selection of the portfolio constituents, and therefore showing the potential benefits for a sophisticated investor.
The paper is organized as follows. In Section 2, we illustrate the two heterogeneous agents. In Section 3, we present the rational learning model following the Bayesian approach similar to that of Black-Litterman. In Section 4, we define the optimizing criterion for the weighting function of the behavioral component. Finally, in Section 5 we perform the empirical analysis.

[3] In some sense, the herding behaviour can be seen in the same way as the bandwagon effect.

2. Rational and behavioral agents

We consider two agents with different utility functions. The first decision maker is equipped with a HARA utility function, and the second with a behavioral utility function. Generally, recalling Zakamouline and Koekebakker (2009a), we define a behavioral agent as a decision maker who discriminates between an outcome above (gain) and below (loss) a reference point. Consequently, the investor's utility function behaves differently in the domain of gains and in the one of losses with a kink at the reference point.
The main difference between the two agents can be explained by their different risk attitudes: the rational investor is risk-averse in all the domains of the utility functions while the behavioral investor might show different risk preferences. Examples are the risk aversion in the gains and risk-seeking in the losses, as in the S-shaped utility function by Kahneman and Tversky (1979).

2.1. The HARA utility function

We consider expected utility theory as the rational investor's reference for the optimal decision making.[4] In this setting, the agent's risk-aversion is associated with the concavity property of her wealth function.
Let's consider a general class of utility functions, concave, everywhere differentiable,

$$U(W) = \frac{\rho}{1-\rho}\left(\frac{\lambda W}{\rho} + b\right)^{1-\rho}, \text{ where } b > 0, \tag{1}$$

where the absolute risk aversion is

$$ARA(W) = r(W) = -\frac{u''(W)}{u'(W)} = \lambda\left(\frac{\lambda W}{\rho} + b\right)^{1-\rho}. \tag{2}$$

The utility function reduces to the quadratic utility when $\rho = -1$, to the negative exponential utility function described by constant absolute risk aversion (CARA) when $b = 1$ and $\rho \to \infty$, and to the logarithmic described by constant relative risk aversion (CRRA) when $b = 0$ and $\rho > 0$.
As reported in Zakamouline and Koekebakker (2009b), the CRRA utility function provides a performance measure consistent with a market equilibrium. This utility function is defined as,

$$U(W) = \begin{cases} \frac{1}{\rho}W^{1-\rho}, & \text{if } \rho > 0, \quad \rho \neq 1 \\ \ln W, & \text{if } \rho = 1 \end{cases} \tag{3}$$

where ρ measures the degree of relative risk aversion.
Mehra and Prescott (1985) indicate a ρ around 30 to be consistent with the observed equity premium in the financial market. For high values of ρ, Zakamouline and Koekebakker (2009b) highlight that the relative preferences for the moments of the distributions are closed to ones related to the CARA utility function. Following the authors and for computational convenience, we consider the CARA instead of the CRRA utility function,

$$U(W) = -e^{-\lambda W}, \tag{4}$$

[4] A decision maker is defined rational according to the Von-Neumann-Morgenstern utility theorem, which defines a set of four axioms: *completeness*, *transitivity*, *independence* and *continuity*. The expected utility theory always satisfies this theorem.

where λ represents the coefficient of risk aversion and W the investor's wealth.
The two-fund separation theorem states that all the investors with the same prior beliefs, independently from their risk aversion, will invest in the same portfolio of risky assets. Sharpe (1964) and Lintner (1965) show that this portfolio is efficient one and represents the core of the formulation of the Capital Asset Pricing Model (CAPM) in a mean–variance world. Cass and Stiglitz (1970) demonstrate that if all investors in the market have a HARA utility function with the same exponent, the two-fund separation principle still holds. Shefrin and Statman (2000) developed a behavioral portfolio theory (BPT) consistent with the Friedman and Savage (1948) puzzle, and show that the mean-variance frontier and the BPT do not coincide. Moreover, the two-fund separation theorem does not hold in their portfolio theory.
In our model, we deviate from the traditional approach and we separately evaluate each asset in terms of its utility function contribution. Hence, we refer to a *single* risky asset instead analysis of a portfolio composed of risky assets. We motivate such a choice with the need of comparing the two agent's views over the assets. Therefore, the evaluation of each risky asset utility will allow creating an ordering, a ranking, of risky assets. However, this is not impacting on the rational utility framework, as this corresponds to the limiting case in which the investor is willing to include a risky component in its capital allocation but her choices are limited in that she can select only one risky asset. Therefore, the investor could evaluate the expected utility associated with the investment in each risky asset (separately considered) and then produce a rank allowing to identify the best asset. In this process, any time the investor evaluates the expected utility given by a risky asset, she consider the expected value of the future wealth, obtained by combining the risky asset with the risk-free asset. We stress that the fraction of wealth invested in the risky asset is unknown and must be determined. According to the maximum principle, the optimal fraction of wealth invested in the risky asset is proportional to a performance measure; the higher the performance measure, the higher the maximum expected utility for the investor. Note that such an evaluation is repeated for each risky asset. Therefore, the asset with the highest performance measure is the best asset. In other words, the ranking across risky asset can be based on performance measures without any loss of generality.
The mean–variance proposed by Markowitz (1952) can be considered a particular case of expected utility theory when the financial returns are normally distributed. In fact, the Sharpe ratio provides the optimal solution for the maximization of the expected, utility since the distribution of returns is completely described by the first two moments.
In this regard, let's consider an investor endowed by a wealth W at the beginning of a period t_0, where a is the amount of wealth allocated in a risky asset x and the remaining $w - a$, the part allocated in the riskless asset r_f.
At the end of the period t_1, the investor's wealth will be

$$\tilde{w} = a \times (1 + x) + (w - a) \times (1 + r_f) = a \times (x - r_f) + w \times (1 + r_f). \tag{5}$$

The investor's objective is to maximize the wealth according to the choice of a,

$$\max_a E[U(\tilde{w})] \tag{6}$$

and therefore the maximized expected utility will be

$$E[U^*(\tilde{w})] = E[-e^{-\lambda[a(x-r_f)+w(1+r_f)]}] = E[-e^{-\lambda[a(x-r_f)]} \times \underbrace{e^{-\lambda w(1+r_f)}}_{\overline{q}}]. \tag{7}$$

It is worth noting that a^* is independent from the investor's initial wealth, and we can treat $\overline{q}$ as a fixed quantity.

By setting $x_0 = w(1+r_f)$ as in Zakamouline (2011), we can approximate the expected utility using Taylor's series,

$$E[U(\tilde{w})] = -1 + a\lambda E(x - r_f) - \frac{\lambda^2}{2} a^2 E(x - r_f)^2 + O(\tilde{w}) \tag{8}$$

and by the first order condition (FOC),

$$\frac{\partial E[U(\tilde{w})]}{\partial a} = \lambda E(x - r_f) - \lambda^2 E(x - r_f)^2 a = 0, \tag{9}$$

we obtain that the *Sharpe ratio* is proportional to the quantity that maximizes the expected utility function,

$$a^* = \frac{1}{\lambda}\frac{\mu - r_f}{\sigma^2} = \frac{1}{\lambda}\frac{\mathrm{SR}}{\sigma}. \tag{10}$$

As shown by Gatfaoui (2009), when there is a departure of risky asset returns from Gaussianity, the ratio begins to be biased, both in the measurement and in the ranking amongst the assets. Therefore, several authors started to consider alternative performance measures, more robust to non-normality evidences. Cherny (2003) and Zakamouline and Koekebakker (2009b), amongst others, propose an improvement of the ratio with the inclusion of higher moments. In particular, the authors propose a parametric Sharpe ratio adjusted for skewness and kurtosis, assuming the normal inverse Gaussian (NIG) as the underlying probability distributions of the financial returns. This probability density function is particularly suitable for distributions with fat tails.
Alternatively, using a non-parametric methodology, both authors followed the Hodges (1998) conjecture by deriving a generalized Sharpe ratio (GSR).[5]
Recalling the maximization of the expected utility,

$$E[U(\tilde{w})] \propto E\big[-e^{-\lambda(x-r_f)}\big] = \max_a \int_{-\infty}^{\infty} -e^{-\lambda a(x-r_f)} \hat{f}_h(x)dx, \tag{11}$$

where $\hat{f}_h(x)$ is the estimated kernel density function of observed returns. The GSR is obtained by the numerical optimization of the expected utility, obtaining:

$$GSR = \sqrt{-2\log\left(-E[U^*(\tilde{w})]\right)}. \tag{12}$$

Notably, such an approach takes into account all the empirical moments of the returns probability distributions, being thus robut to deviations from Gaussianity. We consider the GSR ratio as the performance measure adopted by the rational investor to rank assets.
Finally, it worth noting that the GSR approaches the standard Sharpe ratio when the underlying distribution of the risky asset is close to the Gaussian distribution.

2.2. The behavioral utility function

As mentioned above, a *rational* investor should behave as described in the expected utility theory. Nevertheless, the presence of people who systematically deviate from this behaviour is commonly observed.[6] Empirical analyses on stock market data has shown that prices present an *excess volatility*, compared to what we might expect given the dynamics of their economic fundamentals (Cutler et al., 1989, See for example). Another important stylized fact is *clustered volatility*, where

[5]See Zakamouline and Koekebakker (2009b) for a detailed explanation.

[6]The paradox of Allais (1953) was the first and the best empirical example of a systematic violation of rationality in the expected utility, particularly, the independence axiom.

asset price movements are characterized by periods of high volatility and by periods of low volatility.[7] These examples enforce the hypothesis of the presence of non-rational and heterogeneous agents.
In a behavioral framework, the investor's utility function behaves differently in the domain of gains and in the one of losses, with a kink at the reference point; as an example a behavioral utility function might have the following form

$$U(W)\begin{cases} U_+(W) \text{ if } W \geq W_0, \\ U_-(W) \text{ if } W < W_0, \end{cases} \tag{13}$$

where $U_+(W)$ is the utility function when the investor's wealth W is above the reference point W_0, and $U_-(W)$ is the utility function when the investor's wealth W is below the reference point W_0.
Zakamouline (2011) introduced a generalized *behavioral utility function* characterized by a piece–wise linear plus power utility function,

$$U(W) = \begin{cases} 1_+(W - W_0) - (\gamma_+/\alpha)(W - W_0)^\alpha, & \text{if } W \geq W_0, \\ -\lambda(1_-(W_0 - W) + (\gamma_-/\beta)(W_0 - W)^\beta), & \text{if } W < W_0. \end{cases} \tag{14}$$

where W_0 is the reference point, 1_+ and 1_- are the indicator functions in $\{0, 1\}$ which define the linear part of the utility, γ_+ and γ_- are real numbers that model for the shapes of the utility, and the parameters $\lambda > 0$, $\alpha > 0$ and $\beta > 0$ are real numbers. This utility function is continuous and increasing in wealth with the existence of the first and second derivatives, with respect to the investor's wealth.
Under some conditions and by using the maximum principle, Zakamouline (2011) derives the Z-ratio, the performance measure which maximizes the utility function; thus having the same role of the GSR for the representative rational investor. The expected generalized behavioral utility function can be thus replaced by a function of mean and partial moments of the returns,

$$Z_{\gamma_-,\gamma_+,\lambda,\beta,1_-} = \frac{E(x) - r - (1_-\lambda - 1)LPM_1(x, r)}{\sqrt[\beta]{\gamma_+ UPM_\beta(x, r) + \lambda\gamma_- LPM_\beta(x, r)}}, \tag{15}$$

where x is the returns series of the asset and r is set to the risk–free rate, LPM and UPM are respectively the lower and upper partial moments as defined by Fishburn(1977),

$$\begin{aligned} LPM_n(x, r) &= \int_{-\infty}^{r} (r - x)^n dF_x(x), \\ UPM_n(x, r) &= \int_{r}^{\infty} (x - r)^n dF_x(x), \end{aligned} \tag{16}$$

where n is the order of the partial moment of x at a given threshold r, usually the risk-free asset r_f, and $F_x(\cdot)$ is the cumulative distribution function of x.
The behavioral utility function of Zakamouline (2011), see equation (15), allows modelling several different preferences of a behavioral decision maker. In fact, we can obtain several behavioral types of utility through the calibration of the function parameters. Consequently, we can easily shape the concavity and convexity in the domain of gains or losses.
We consider here four different cases which recall some well-known behavioral utility functions. We label them in the same way as Zakamouline and Koekebakker (2009a).

[7] GARCH models introduced by Engle (1982) and Bollerslev (1986) explain these type of phenomena.

- **behavioral I** is the Fishburn's utility function.
The agent equipped with this utility is risk averse in the domain of losses and risk neutral in the domain of gains.
The parameters are set as follows:

$$\gamma_+ = 0, \gamma_- = .1, 1_+ = 1, 1_- = 1, \lambda = 1.5, \beta = \alpha = 2.$$

Notably, the Sortino ratio is the measure that maximizes the utility function when the risk aversion λ is equal to 1; see Sortino and Price (1994).

- **behavioral II** is analogous to the the utility function used in the prospect and cumulative prospect theories. In this utility function, the decision maker exhibits loss aversion, which is defined in a local sense around the reference point (Köbberling and Wakker, 2005).[8] The risk aversion coefficients equals

$$\lambda = \frac{U'(W_0-)}{U'(W_0+)},$$

where we have the left derivative in the numerator and the right derivative in the denominator . If λ is greater than 1 the individual exhibits loss aversion. The parameters are set as follows:

$$\gamma_+ = .1, \gamma_- = -.1, 1_+ = 1, 1_- = 1, \lambda = 1.5, \beta = \alpha = 2.$$

- **behavioral III** relates to the disappointment theory (DT) introduced by Bell (1985). The decision maker experiences disappointment when an outcome is worse than expected (the reference point). Conversely, when an outcome is better than the expected one, a magnification is generated. The utility function is concave below the reference point and it can be convex above. The parameters are set to:

$$\gamma_+ = .1, \gamma_- = .2, 1_+ = 1, 1_- = 1, \lambda = 1.5, \beta = \alpha = 2.$$

- **behavioral IV** is the utility where the decision maker is equipped with piece-wise power utility function with non-linear parts.
The parameters are set to:

$$\gamma_+ = -\alpha, \gamma_- = \beta, 1_+ = 0, 1_- = 0, \lambda = 1.5, \alpha = 1.5, \beta = 2.$$

The existence of a solution to the optimal capital allocation requires that $\beta > \alpha$; therefore, the investor does not show loss aversion. Zakamouline (2011) shows that the performance measure that maximizes their utility function is given by the Farinelli-Tibiletti ratio, see Tibiletti and Farinelli (2003).

It is worth noting that the general behavioral utility reduces to a quadratic utility when $\lambda = 1$, $\alpha = \beta = 2$ and $\gamma_+ = \gamma_- > 0$. If returns are normally distributed, the $CRRA$, the $CARA$ and the quadratic utility are maximized in the function of the Sharpe ratio measure. Therefore, when returns converge to normality, we can relate the rational investor to a peculiar case of the general behavioral utility function. This fact confirms the appropriateness of the two types of utilities. In Figure 6 we provide examples of the shape of the rational and behavioral utilities we consider.

[8] In contrast, Kahneman and Tversky (1979) define the loss aversion in a global sense,

$$-U(W_0 - \Delta W) > U(W_0 + \Delta W), \quad \forall \Delta W > 0.$$

3. A Rational Learning Model

The aim of the model is to blend the assets selection by the rational investor with the choice of the behavioral counterpart. This combination is done by conditioning the rational choice on a behavioral ordering where the evaluation is performed in terms of expected utility. If the rational investor modifies her choice by taking into account a behavioral selection in the market, then, up to a certain degree, she acts in a more sophisticated way.
In fact, this agent is adopting a herding behaviour in order to improve her investment. Herding behaviour is the tendency of an investor to abandon her own information in order to mimic the behaviour of other investors. This agent considers implicitly the aggregated evaluation coming from the different utility functions as the best way to perform the optimal selection amongst the assets, with a "one step ahead" horizon.
In practice, the more weight is given to the behavioral component, the greater is that component's relevance for the rational investor. The two extremes are the limiting cases where the mixed selection collapses into one of the components.
As reported in Forbes (2009), if investors are not irrational and are learning to invest better, their learning process takes place in accordance with the Bayes' rule,

$$P(A|B) = \frac{P(A)P(B|A)}{P(B)}. \tag{17}$$

Therefore, the most appropriated way to obtain the aggregated measure is using the Bayesian approach. We conjugate these two components analogously to the model proposed by Black and Litterman (1992).
We place our perspective on an investor with the HARA utility function, which is considered the benchmark for the rational investor in the expected utility framework. Hence, the generalized Sharpe ratio is the measure used to evaluate the assets in terms of this utility function. This is our prior distribution leading to a rational ranking of assets. We assume for simplicity that the rational investor is myopic by ignoring that other agents are also engaged in a dynamic learning process.
The general behavioral utility function represents the additional information used along with the prior distribution to infer the posterior one. The Z ratio from Zakamouline (2011) is the measure coming from this utility function. We structure the model in a similar way as that of He and Litterman (2002) where the main difference is due to the functional transformation of the returns through the performance measures. Consequently, the aggregated performance measure is defined by the posterior distribution.

3.1. The model

Generally, we consider the performance measure of an asset as a random variable independently and identically distributed,

$$PM_i \sim iid(\mu, \sigma^2). \tag{18}$$

We are interested in identifying the assets i with the highest values of the expected PM_i as those representing the best opportunity for the investor in terms of utility function. In performing this choice, we start from a prior distribution for μ, which is assumed to be normally distributed when centred to the generalized Sharpe ratio obtained from the optimized expected utility function,

$$\mu_{GSR} = GSR(E(U^*(W_i)))) + \varepsilon, \tag{19}$$

where ε is a normally distributed error with mean equal to zero and variance, $\tau\sigma^2$. As in the Black–Litterman model, τ represents the uncertainty on the prior density. The higher the τ, the

higher the uncertainty given to the prior density. Conversely, the closer τ is to zero, the lower is the variance of the prior density, and therefore the lower the relevance given to the conditioning information represented by the behavioral component. The parameter τ is defined in $[0, \infty]$.
If we assume a coexistence between a behavioral and a rational component in the market, it is reasonable to expect an improvement in a ranking that merges the points of view, either from a return perspective, from a risk view, or from a performance evaluation analysis.
The conditioning information coming of the general behavioral utility function can be of different form, according to the shape given to the utility in terms of gains and losses preferences. If we consider the behavioral types of utility as described above, we have four different types of investors. We assume thus that different behavioral investors are jointly active in the market and the rational investor accounts for their joint presence.
Generally speaking, we can have k different behavioral views (as in the Black-Litterman model). Those, similarly to the B&L model, are linearly combined with a selection vector. The latter, in our case, collapses to a k-dimensional vector of ones, named P. Moreover, we denote by $\mathbb{Z}_{\gamma_-,\gamma_+,\lambda,\beta,1\_}$ the k-vector of the behavioral measures that we declined from the Z ratio. We thus have the following quantities:

$$\begin{aligned} &P' = \vec{1}, \\ &\mathbb{Z}' = (Z_1, Z_2, Z_3, Z_4)\,, \end{aligned} \tag{20}$$

where we have the four different Z ratios associated with the four declinations of the behavioral utility function defined in Section 2.2. The mean of the behavioral measures is centred to the Z ratios plus an error term normally distributed with zero mean and variance matrix Ω:

$$\mu_{\mathbb{Z}} = \mathbb{Z}_{\gamma_-,\gamma_+,\lambda,\beta,1\_}(E(U^*(W_i))) + \eta \tag{21}$$

We assume that ϵ and η are independent,

$$\begin{pmatrix} \varepsilon \\ \eta \end{pmatrix} \sim N\left(0, \begin{bmatrix} \tau\sigma^2 & 0 \\ 0 & \Omega \end{bmatrix}\right) \tag{22}$$

and with an application of the Bayes theorem, similarly to Black and Litterman (1992), the aggregated expectations (behavioral and rational) are distributed as a normal with mean μ_p and covariance matrix M_p, with the following analytical expressions

$$\mu_p = \left[(\tau\sigma^2)^{-1} + P'\Omega^{-1}P\right]^{-1}\left[(\tau\sigma^2)^{-1}GSR + P'\Omega^{-1}\mathbb{Z}_{\gamma_-\gamma_+,\lambda,\beta,1\_}\right] \tag{23}$$

and

$$M_p = \left[(\tau\sigma^2)^{-1} + P'\Omega^{-1}P\right]^{-1},$$

where, μ_p represents the aggregated expected performance measure coming from a mixture of the two components: the first from the HARA type utility function and the second from the behavioral utility functions I to IV. Thus, μ_p can be used as a performance measure to rank the investment universe blending the rational and behavioral perspective.

4. Calibrating τ and the weight of the behavioral components

Since the (rational) agent is learning looking at the other (behavioral) agent types according to the Bayesian framework, we need to calibrate the weight of the behavioral component (the conditional) by tuning the uncertainty on her prior beliefs. Moreover, in order to act in a more

sophisticated way, the rational learner agent should also calibrate the relative weight of each component, the various behavioral views. In fact, in a particular market momentum, one of the considered behavioral agents could be more (less) relevant with respect to the others. More importantly, this prevent us to impose a specific structure on the model (for example attributing the same weight to each agent).
In this regard, we reformulate equation (20) as

$$\begin{aligned} P' &= \vec{\Theta}, \\ \mathbb{Z}' &= (Z_1, Z_2, Z_3, Z_4) . \end{aligned} \tag{24}$$

where $\Theta = (\theta_1, \theta_2, \theta_3, \theta_4)$. Finally, the rational learner agent should optimize τ and Θ according to a specified criteria. The chosen criterion is to weight the conditional component (and the relative weight of each behavioral component) using the past performances in terms of cumulated returns,

$$r_p = \frac{1}{m} \sum_{l=t-m+1}^{t} r_{p,l}, \tag{25}$$

where $r_{p,l}$ is the time l return of an equally weighted portfolio, and m represents the time range for the portfolio evaluation (from time $t - m + 1$ to time t). The portfolio is composed of the best-performing equities, according to equation (23). This criterion function can be interpreted as risk neutral, since we are focusing only on the first moment.
Let $\mathcal{A}_t(\tau, \mathbf{P}(\boldsymbol{\Theta}))$ be the set containing the M best assets selected across the K assets included in the market (with $M << K$) at time t. This set depends on τ and Θ, as a change in these parameters modifies the rankings produced by the agents. The set is also a function of time, given that the impact of the behavioral choices might change over time. Therefore, we represent the portfolio returns as

$$r_{p,l} = \frac{1}{M} \sum_{j \in \mathcal{A}_t(\tau P(\Theta))} r_{j,l}, \tag{26}$$

where $r_{j,l}$ is the return of asset j at time l; we stress that the index j varies from 1 to K, although only M values are included in the set $\mathcal{A}_t(\tau, P(\Theta))$. Given the dependence on τ and Θ of the best-performing asset set, the portfolio's cumulative return in (25) is also a function of both parameters. The optimal choice of τ and Θ is determined by maximizing the portfolio returns, that is,

$$\begin{aligned} &\max_{\tau, \boldsymbol{\Theta}} f(\tau) = \frac{1}{m} \sum_{l=t-m+1}^{t} r_{p,l} \\ &s.t. \begin{cases} r_{p,l} = \frac{1}{k} \sum_{j \in \mathcal{A}_t(\tau, P(\Theta))} r_{j,l}, \\ \sum_{i=1}^{4} \theta_i = 1, \quad \theta_i > 0. \end{cases} \end{aligned} \tag{27}$$

As previously noted, we are considering an equally weighted allocation scheme across the best assets. This choice is clearly restrictive, but it allows limiting the impact of the portfolio weights estimation error and it has been shown to be preferred over optimal weighting schemes by DeMiguel et al. (2009).
The optimal τ^* and Θ^* provide the maximum cumulative return obtained by investing in a subset of risky assets traded in the market and making decisions by blending the rational and behavioral rankings. Consequently, the estimated τ^* represents the relevance of behavioral choices or, conversely, the reliability of the rational rankings. The methodology for the optimal criteria investing in a subset of risky assets from an investment universe is similar to the one adopted by Billio et al.

(2012). Looking at the estimated optimal parameters, a high value of τ^* implies that, to obtain the optimal return, the rational investor should have corrected her choice towards a behavioral direction. On the other hand, a low value of τ^* means that the investor should make decision according to her prior rational rankings. The parameter Θ^* represents the optimal weight given to each behavioral agent and it lies in the domain $[0,1]$. The criterion function enables us in some sense to weight the components, rational versus behavioral, through τ. Moreover, by solving (27) over multiple samples, we obtain a sequence τ_t^* that gives further insight into the fluctuation and evolution of this factor. We see the choice of this criterion function as natural since it focuses only on the expected return of the given portfolio, that is, a risk-neutral evaluation. Caporin et al. (2014) in an analogous framework considered the rational and the behavioral agent equipped with an S-shaped utility function. Their analysis is focused on the financial interpretation of τ and it is based on the S&P 500 market.

5. Blending rational and behavioral ranking in the US market

Our purpose is to apply the model to an investment universe composed by a large number of assets. We want to compare the different strategies, behavioral, rational and the Bayesian blending, by building portfolios based on the $k < n$ selected constituents. The portfolios we compare stem from the ranking provided by the performance measure in terms of the utility of the rational investor (the generalized Sharpe ratio), the one based on the behavioral agent (the Z ratio) and finally, the one from the rational investor conditional to the behavioral component. We stress the portfolios are equally weighted in order to focus on the agents' selection, to avoid possible corner solutions obtained by utility maximization, and to avoid the impact of weights estimation error.

5.1. The dataset

Our investment universe is given by the full list of quoted stocks on the NASDAQ and NYSE from September 1972 to December 2014 for a total of 15,790 assets. Dead series are also included in the dataset. We handle the delisting and merger and acquisition of a stock by assuming that if the investor has selected this stock, she would have disinvested the asset during the last period of its quotation in the market. The series have been downloaded from Datastream at a weekly frequency. We also recover a proxy of the risk-free asset, the US Treasury 3 Month Bill rate. In a first evaluation, we refine the investment universe a priori by excluding the asset with a lower market value in order to mitigate the liquidity risk. The level of exclusion is fixed at 50%; Figure 2 shows the final number of assets for each period. We build the benchmark for market comparison as a value weighted (equally weighted) index based on the investment universe considered at a given time t. Since the number of assets in the investment universe is time varying, it follows that the number of selected constituents k changes according to n_t.

5.2. Model settings and estimation

We apply the model on rolling windows of 240 weekly returns to take into account the time-varying structure of the series (Zivot and Wang, 2007). The rolling window is moved forward by one week at each iteration. Hence, an asset enters the valuation process when its time series is longer than the dimension of the bandwidth. The dimension of the invested portfolio k represents a fraction α of the investment universe n_t in order to take into account its expansion through time as showed by Figure 2. In this regard, we set α equal to 25%.[9]

[9]As robustness check, we considered also $\alpha = \{1\%, 5\%, 10\%, 50\%, 75\%\}$. Results are reported in complementary material available upon request.

We report in Figure (4) the number of portfolio constituents for the sample. On each rolling window we estimate the performance measures (GSR and Z ratios), then we estimate their variances in equation (22) using a block bootstrap procedure (Lahiri, 1999) with a block of dimension 16 in order to consider the potential time dependence among the returns. The bootstrap procedure has been applied to the returns, measures have been computed for each iteration and then the variances have been obtained. Once the variances are obtained, we proceed with the estimation of τ and Θ through the maximization in equation (27). Finally, after calculating μ_p, we rank the investment universe and we select the k assets with the highest values as portfolio constituents. The sequential estimation have been implemented in Matlab and takes approximately 138 hours for 532,593,120 iterations on a cluster multiprocessor system which consists of 4 nodes; each comprises four Xeon E5-4610 v2 2.3GHz CPUs, with 8 cores, 256GB ECC PC3-12800R RAM.

5.3. In-sample results

In this section, we discuss the estimation of the optimal τ_t^* and the weights of each behavioral agent type, Θ^*. Figure 3 shows the dynamic of the τ_t^* from September 2, 1977 to December 19, 2014. The estimated factor is stationary in mean until March 2013, after that we have a change since the τ_t^* approaches zero, which means that the rational learner agent should remain completely on her prior beliefs in evaluating the market's assets. This can be observed also in Figure 6, which shows the estimated optimal weight Θ^* of each behavioral component through the sample. Until 2003, the predominant weights are constituted by the Fishburn (Z1) and Farinelli-Tibiletti (Z4) agent. After that period, the dynamic on the behavioral components changes. In fact, we can observe a similar dynamic of the τ_t^* on the weight assigned to the agent equipped with a Fishburn utility function (Z1), which approaches to one after March 2013. Clearly, the dynamic of the estimated parameters show that the model is "working" and therefore the learner process if profitable for the rational agent in terms of cumulative returns if she is conditioning her beliefs toward the direction of other agents. If there are no gains in the conditioning, τ converges to zero, as shown in Figure 3 from March 2013. To investigate the degree of similarity among the rational agent with and without the learning process, we compute a concordance index (CI) to highlight the common selection in terms of assets. That is, we want to compare the rankings and the optimal selection made by the rational agent when she does/does not account for the presence of behavioral choices in the market. We thus compute the following index

$$CI = \frac{\mid I_{RL} \cap I_{GSR} \mid}{M(\alpha)} \in [0,1], \tag{28}$$

where I_{RL}, I_{GSR} are the set of selected assets for two strategies GSR and RL, $\mid \cdot \mid$ indicates the cardinality of the set, and $M(\alpha)$ is the number of selected assets as a function of the chosen fraction of the investment universe. Figure 5 shows the concordance index, which approaches one when τ_t^* move toward zero at the end of the sample period. This confirms the coherence of the model. In fact, at the end of the sample, the rational agent's posterior collapses on the prior and therefore the two selection coincides. From the rational learner agent perspective, the change in τ^* has a clear meaning: the agent is not conditioning her choice since she obtains a higher cumulative return without conditioning. However, the economic and financial interpretation of this variation in τ is not immediate. One possible explanation could be represented by the adaptive market hypothesis (AMH) proposed by Lo (2004). If we consider the market as an evolutionary system where participants interact and evolve dynamically according to intrinsic rules of economic selection, changes in τ_t^* and Θ may explain part of this evolution. In fact, profit maximization, utility maximization, and general equilibrium are relevant aspects of market ecology (Lo, 2004). Hoffmann et al. (2013) combining monthly survey data with trading records document significant

swings in trading and risk-taking behavior during the 2007-2009 crisis, which are driven by changes in investor perceptions.
As robustness check, we test if the volatility of the τ parameter is driven by some factors such as the market returns, the volatility of the market or simply by the dimension of the investment universe. In this regard, we estimate the volatility of τ_t^* using an ARFIMA(1,d,1)-FIGARCH(1,d,1) model and the market volatility using a FIGARCH(1,d,1) on the value-weighted index.[10] Table 1 reports the estimation results and Figure 7 shows the estimated volatility of τ_t^*. Clearly, market returns and market volatility do not explain the volatility of the τ factor (first and second column), adjusted-R-squared are close to zero although market volatility is significant and negative related, which means that in turbulent periods the weight of the behavioral component decreases in absolute terms, while in relative terms it can be seen in Figure 6 that the weights are distributed across type II (S-shaped utility function) and type III (disappointment theory), which are related both to the concept of loss aversion. It is worth noting that these two types of behavioral agents arise by the mid-2000s until 2003, while before that date the weights are concentrated mainly on Type I (Fishburn's utility function) and Type IV (Farinelli and Tibiletti's agent). We control also for the dimension of the investment universe n_t, which could be a source of volatility in the factor since it directly affects the cardinality of the selection. Estimation results show a significant positive relationship, but with a coefficient very close to zero; the adjusted-R-squared is close to 12% (third column). Finally, the last column of Table 1 reports the estimation with all the covariates, which confirms the previous results.

5.4. Out of sample results

We conduct an out-of-sample application of the model by considering the time t the values τ_t^* and Θ^* as the best predictor for $t+1$ (a week ahead). In this regard, we perform the allocation from September 1977 to December 2014 and we initialize the portfolios with a value of 100. Table (2) reports the descriptive analyses of the portfolios, while Figure 8 plots the cumulative returns among different agents. Over the full sample the best portfolio in terms of cumulative returns and Sharpe Ratio is the value-weighted market index. The second-best is the rational agent, without learning mechanism. Figure 9 shows the comparison among the rational learner agent, the rational agent (whom remains on her prior), the value-weighted and equally weighted market index, while table 3 reports the annual portfolios' statistics. The results suggest that, when focusing on a long investment period, about 37 years, rational choices seems the most profitable and the less risky. This is somewhat coherent with the perception that behavioral views might be subject to larger volatility that could thus induce a loss in terms of returns. On the contrary, the poor results with respect to the benchmarks are not surprising nor lead to a negative evaluation of our approach as the allocation strategies have not been created for the purpose of beating the benchmark, but rather with the purpose of highlighting the effect of the rational learning approach. In order to shed further light on the comparison among the various agents, we move to a yearly evaluation of the strategy performances. The Sharpe Ratio (last column), which is a performance measures related with the rational agent, shows that the rational learner mechanism was profitable during the years 1979, 1980, 1988, 1999, 2003, 2004, 2005, 2007, 2010, 2012, 2013 and 2014 (13 out 38) using Θ_t^* and τ_t^* as predictor for $t+1$. Therefore, there are periods where the blending of rational views with behavioral views lead to an improvement of performances. Notably, the evidence is much stronger in the last 15 years, corresponding to the increase of direct trading by private/small investors, those who are more exposed to behavioral choices (we could think that a professional investor

[10]Estimation results of these regression are not reported but are available upon request.

should make choices in a more rational way). As stated in Barber and Odean (2001), the internet changed how information is delivered and the way agents can invest on this information. From the mid-90s the number of brokerage accounts began rapidly to expand opening the gate to access directly to the financial market even to the small investors. Meneu and Pardo (2004) investigated the existence of a pre-holiday effect in different assets due to the reluctance of small investors to invest on pre-holiday days. Their results show that institutional investors could have exploited this anomaly. In an analogous way, if we think about institutional investors as the rational ones, they should take into account the behavioral component, which is reasonably related to small and private investors. To further compare the approaches, we focus on the turnover reported in Figure 10. The Rational Learner agent shows an higher turnover, which is due to the learning process that cause a larger volatility in the portfolio composition. Notably, at the end of the period, the turnovers are equal since the selection of the rational learner converges to that of the rational agent. Figure 11 shows the portfolio's composition according to the industrial composition benchmark. In the Appendix, we also report the figures for each behavioral agent. There is a degree of similarity among the behavioral agents as shown in Figure A.14-A.17, for instance on the weight attributed to the Utilities industry around 1997. On the other side, the rational agent in the same period gives more weight to Financial and Technology industry. Figure 11 shows the composition for the rational learner agent, which is clearly more volatile due to the composite selection. In fact, the portfolios' composition reflect the different preferences among the agents. Figure 12 shows the diversification index for the rational agent and the rational learner agent obtained as,

$$DI_t = \sum_{j=1}^{s} \left| \omega_i - \frac{1}{s} \right|, \tag{29}$$

where s is the number of industries and ω_i is the weight of each sector in the given portfolio. The average of the DI for the rational agent (GSR) in the period is equal to 0.6662, while for the rational learner agent (RL) is equal to 0.7012. The values are similar for both agents. It is worth noting that at the end of the period the portfolio's composition of the rational learner agent collapses into the composition of her prior (GSR). Clearly, in the out of sample case the portfolios' returns are better for the rational agent without the learning process. However, it is worth noting that these results strongly depend on the forecast made on τ_{t+1} and Θ_{t+1}.

6. Conclusion

In this paper, we present a heterogeneous agent model, which considers two decision makers; the classical risk-averse agent equipped with a HARA utility function and an agent with a general behavioral utility function. The HARA agent adopts a learning process by updating her beliefs according to the presence of a behavioral component. The learning process takes place in a Bayesian manner, where the model conjugates the choice of a HARA investor towards a behavioral counterpart with weighing factor τ. In our investigation, the rational investor adopts a herding behaviour in the selection of the portfolio constituents. This effect has been checked by estimating the weighting factor τ and the weight of each behavioral component Θ with an optimizing criteria.

We performed an empirical analysis over all the assets listed in the NASDAQ and NYSE stock exchange from September 1972 to December 2014. In sample results show an improvement for the rational investor who adopts a learning process, modifying her choice in the evaluation of the assets with the behavioral counterpart. When the conditioning mechanism is no longer convenient τ goes very close to zero (as in last part of the sample, March 2013) showing its ability as learning factor since it quickly reacts to the market momentum. Moreover, the volatility of the τ_t^* factor is not driven by market returns, market volatility or changes in the investment universe. The relative weights of the behavioral component Θ_t^* suggests that different agents could capture different market momentum. Out-of-sample results are obtained using as best predictor for $t+1$, τ_t^* and Θ_t^*. In this case, we show that in some period the learning process is convenient for the learner agent. Clearly, the results in the out of sample application strongly depend on the forecast methodology used for the τ and θ factors to obtain the aggregated performance measure in $t+1$. This can be further refined as future extension given the persistence showed in the factors' dynamic. In our empirical analysis, the Bayesian learner investor with a HARA utility function has adopted a herding strategy; further research should also analyze a contrarian strategy allowing for negative weights on θ_i, or a conditioning switching strategy.

Allais, M., 1953. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. Econometrica 21 (4), 503–546.

Barber, B. M., Odean, T., 2001. The internet and the investor. Journal of Economic Perspectives, 41–54.

Bell, D. E., 1985. Disappointment in decision making under uncertainty. Operations Research 33 (1), 1–27.

Billio, M., Caporin, M., Costola, M., 2012. Backward/forward optimal combination of performance measures for equity screening. University Ca'Foscari of Venice, Dept. of Economics Research Paper Series.

Black, F., Litterman, R., 1992. Financial Analysts Journal 48 (5), 28–43.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31 (3), 307 – 327.

Caporin, M., Corazzini, L., Costola, M., 2014. Measuring the behavioral component of financial fluctuations: An analysis based on the s&p 500. Available at SSRN 2504381.

Cass, D., Stiglitz, J. E., 1970. The structure of investor preferences and asset returns, and separability in portfolio allocation: A contribution to the pure theory of mutual funds. Journal of Economic Theory 2 (2), 122–160.

Cherny, A., 2003. Generalized sharpe ratios and asset pricing in incomplete markets. European Finance Review 7 (2), 191–233.

Cutler, D. M., Poterba, J. M., Summers, L. H., 1989. What Moves Stock Prices? Journal of Portfolio Management 15 (2).

D'Avolio, G., 2002. The market for borrowing stock. Journal of Financial Economics 66, 271 – 306.

De Long, J. B., Shleifer, A., Summers, L. H., Waldmann, R. J., 1990. Noise trader risk in financial markets. Journal of political Economy, 703–738.

DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? Review of Financial Studies 22 (5), 1915–1953.

Engle, R. F., 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. Econometrica 50 (4), 987–1007.

Forbes, W., 2009. Behavioural finance. Wiley.

Friedman, M., 1953. The Case for Flexible Exchange Rates, 157–203.

Friedman, M., Savage, L. J., 1948. The utility analysis of choices involving risk. Journal of Political Economy 56 (4), 279–304.

Garber, P. M., 1990. Famous first bubbles. The Journal of Economic Perspectives 4 (2), 35–54.

Gatfaoui, H., 2009. Deviation from normality and sharpe ratio behavior: A brief simulation study. Investment Management and Financial Innovations 7 (4), 95–107.

Grossman, S. J., Stiglitz, J. E., 1980. On the impossibility of informationally efficient markets. The American economic review, 393–408.

He, G., Litterman, R., 2002. The intuition behind black-litterman model portfolios.

Hodges, S., 1998. A generalization of the sharpe ratio and its applications to valuation bounds and risk measures.

Hoffmann, A. O., Post, T., Pennings, J. M., 2013. Individual investor perceptions and behavior during the financial crisis. Journal of Banking & Finance 37 (1), 60–74.

Hommes, C. H., 2006. Chapter 23 heterogeneous agent models in economics and finance 2, 1109 – 1186.

Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica 47 (2), 263–292.

Köbberling, V., Wakker, P. P., 2005. An index of loss aversion. Journal of Economic Theory 122 (1), 119 – 131.

Lahiri, S. N., 1999. Theoretical comparisons of block bootstrap methods. Annals of Statistics, 386–404.

Lamont, O. A., Thaler, R. H., 2003. Anomalies: The law of one price in financial markets. The Journal of Economic Perspectives 17 (4), 191–202.

Lintner, J., 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. The Review of Economics and Statistics 47 (1), 13–37.

Lo, A. W., 2004. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. Journal of Portfolio Management, 15–29.

Malkiel, B. G., 2003. The efficient market hypothesis and its critics. The Journal of Economic Perspectives 17 (1), 59–82.

Markowitz, H., 1952. Portfolio selection. The journal of finance 7 (1), 77–91.

Mehra, R., Prescott, E. C., 1985. The equity premium: A puzzle. Journal of Monetary Economics 15 (2), 145 – 161.

Meneu, V., Pardo, A., 2004. Pre-holiday effect, large trades and small investor behaviour. Journal of Empirical Finance 11 (2), 231–246.

Pedersen, C. S., Satchell, S. E., 2002. On the foundation of performance measures under asymmetric returns. Quantitative Finance 2 (3), 217–223.

Sharpe, W. F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. The Journal of Finance 19 (3), 425–442.

Shefrin, H., Statman, M., 2000. Behavioral portfolio theory. Journal of Financial and Quantitative Analysis 35 (02), 127–151.

Sortino, F. A., Price, L. N., 1994. Performance measurement in a downside risk framework. The Journal of Investing 3 (3), 59–64.

Tibiletti, L., Farinelli, S., 2003. Upside and downside risk with a benchmark. Atlantic Economic Journal 31, 387–387.

Zakamouline, V., 2011. Portfolio performance evaluation with loss aversion. Quantitative Finance, 1–12.

Zakamouline, V., Koekebakker, S., 2009a. A generalisation of the mean-variance analysis. European Financial Management 15 (5), 934–970.

Zakamouline, V., Koekebakker, S., 2009b. Portfolio performance evaluation with generalized sharpe ratios: Beyond the mean and variance. Journal of Banking & Finance 33 (7), 1242 – 1254.

Zeman, E., 1974. The unstable behavior of stock exchange. Journal of Mathematical Economics (1), 39–49.

Zivot, E., Wang, J., 2007. Modeling Financial Time Series with S-PLUS®. Vol. 191. Springer Science & Business Media.

List of Figures

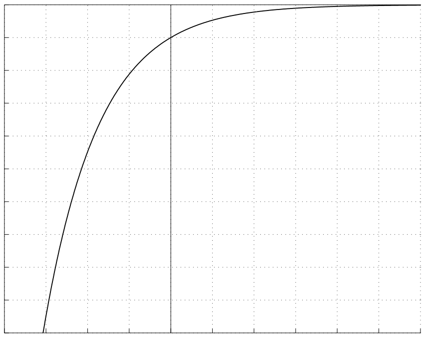

(a) Rational Type

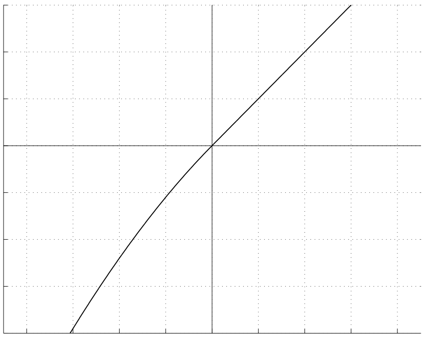

(b) Behavioural Type I

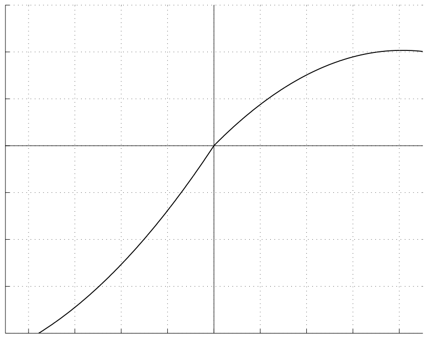

(c) Behavioural Type II

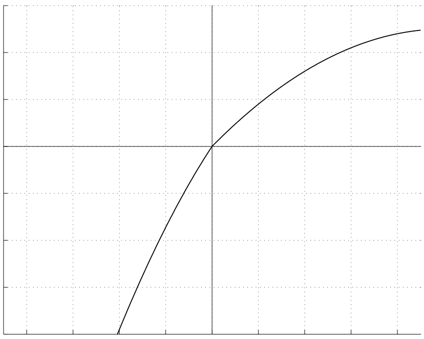

(d) Behavioural Type III

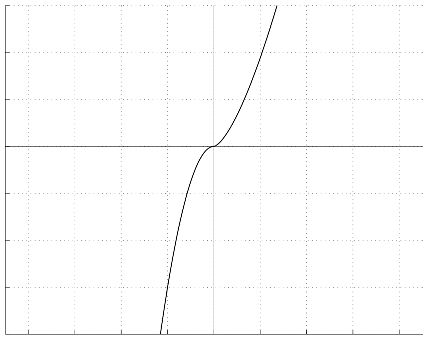

(e) Behavioural Type IV

Figure 1: Negative exponential utility function with constant absolute risk aversion (CARA) and the four specified utility types from Zakamouline's (2011) general behavioural utility function.

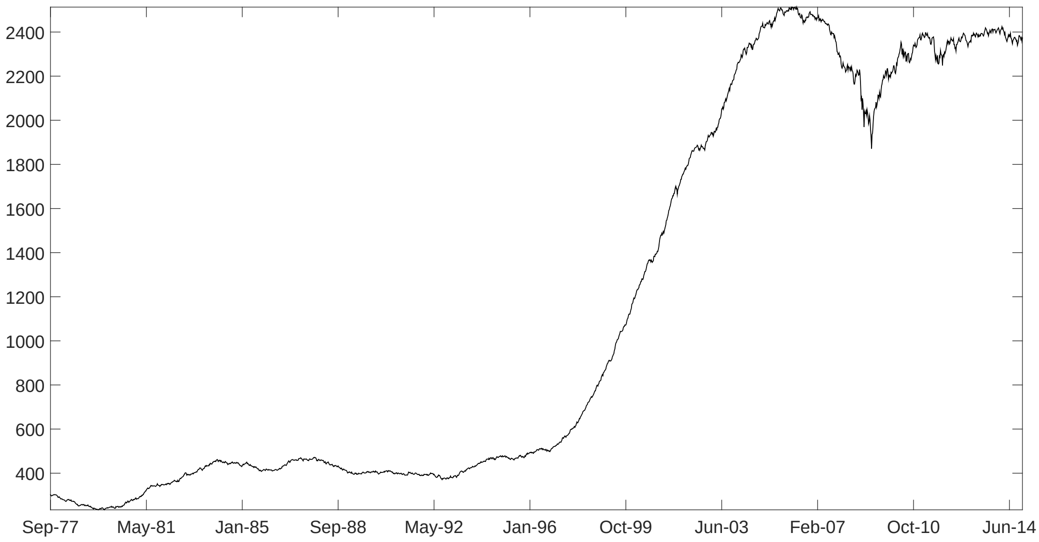

Figure 2: Number of assets in the investment universe across time.

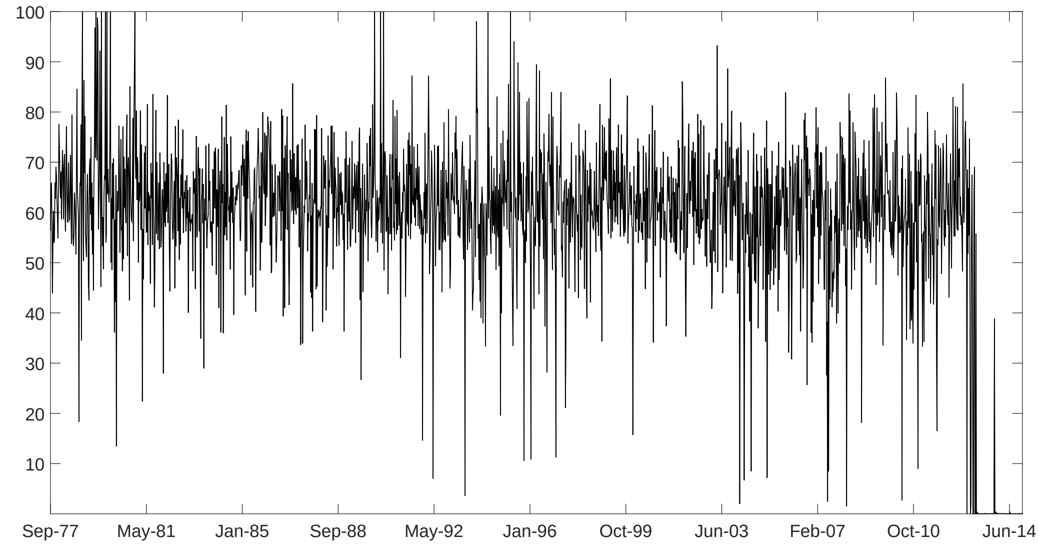

Figure 3: The estimated optimal τ in the period 1977-2014 (weekly frequency). The τ represents the uncertainty in the rational agent (prior). The higher the its value, the greater is the weight assign to the conditional part.

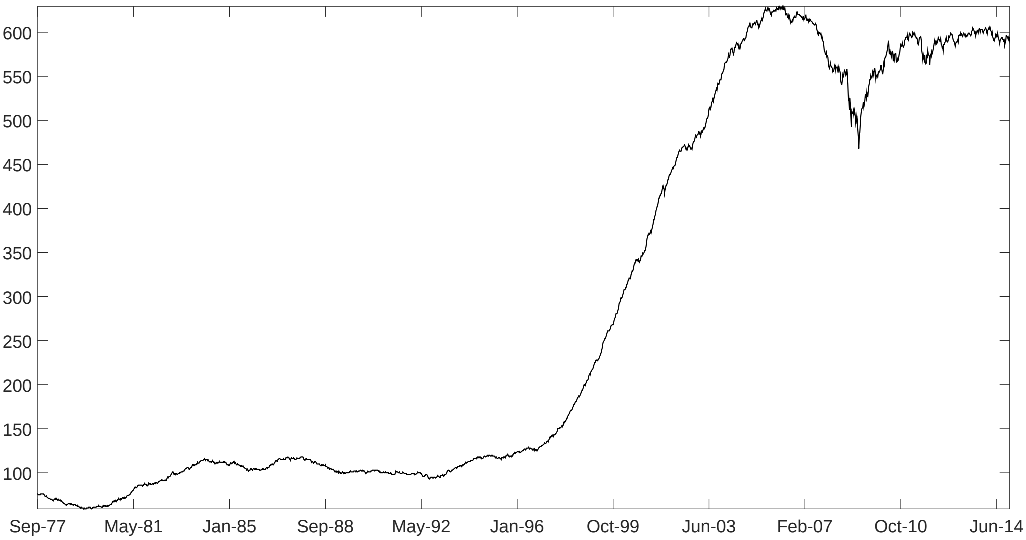

Figure 4: The number of portfolio's constituents in each period according to the invested fraction of the investment universe in the period 1977-2014 (weekly frequency).

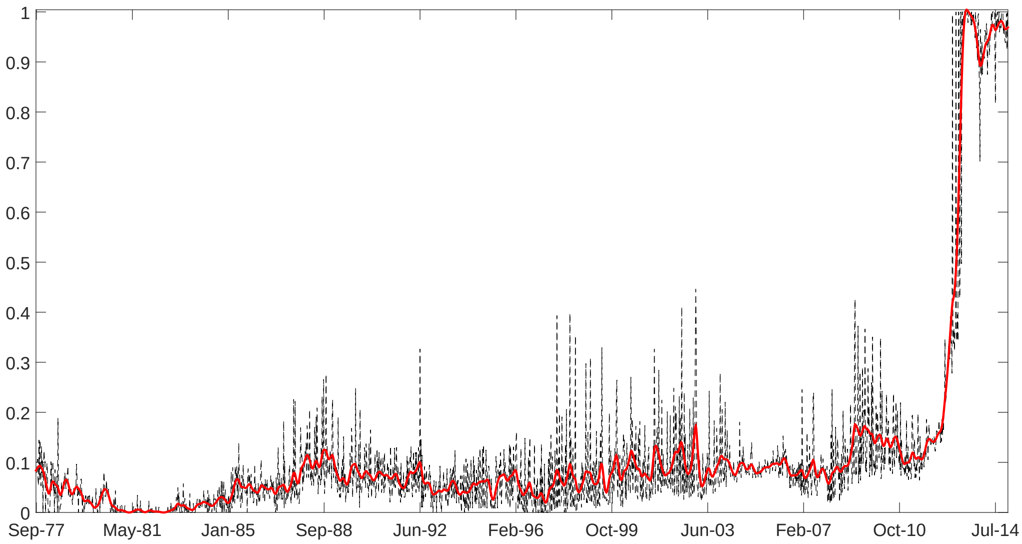

Figure 5: The concordance index among the selection of the rational learner agent (RL) and the rational agent (GSR). The solid line represents the smoothed series. The considered period is 1977-2014 (weekly frequency).

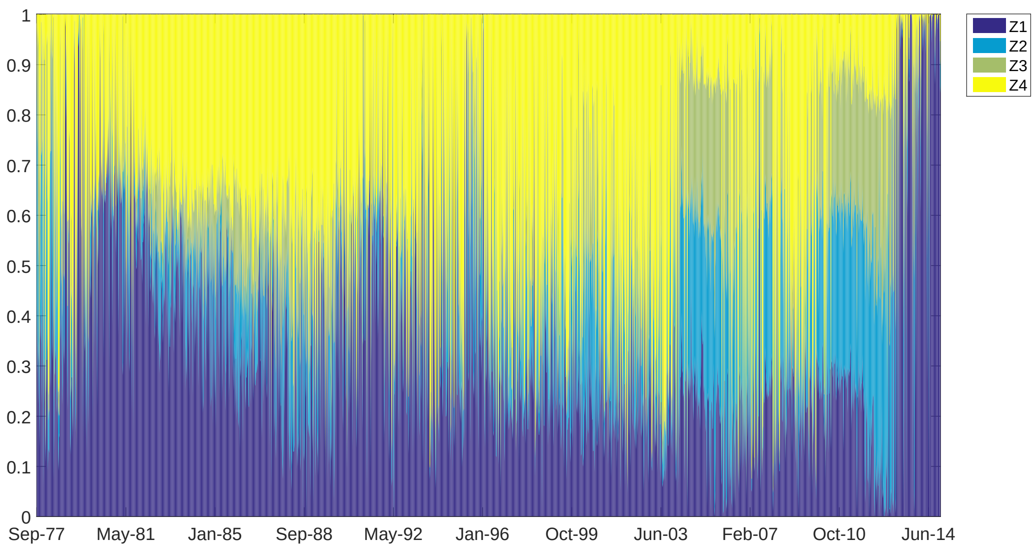

Figure 6: The estimated optimal weights Θ for each behavioral utility function in the period 1977-2014 (weekly frequency). The behavioral utility functions are: Fishburn (Z1), S-shaped (Z2), Disappointment theory (Z3), and Farinelli-Tibiletti (Z4).

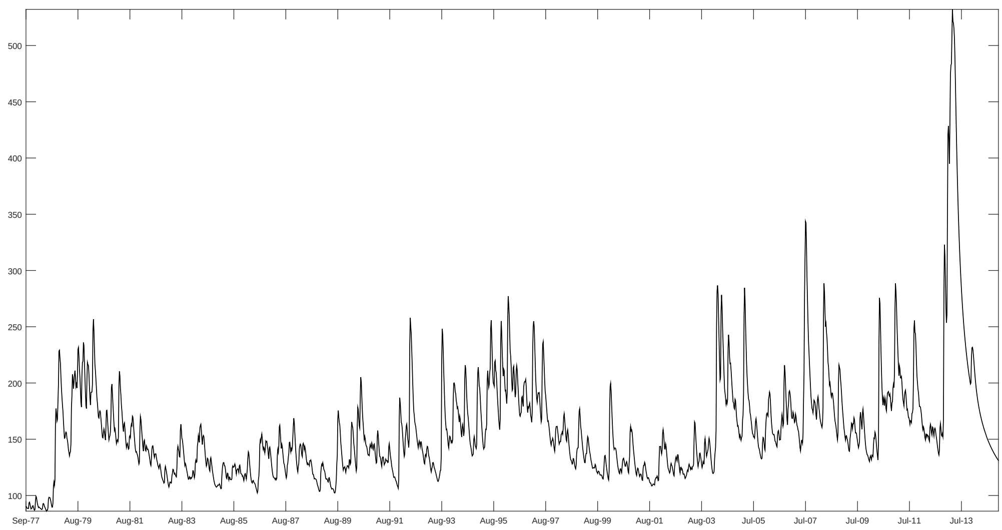

Figure 7: The estimated volatility for τ using the ARFIMA(1,d,0)-FIGARCH(1,d,1) model.

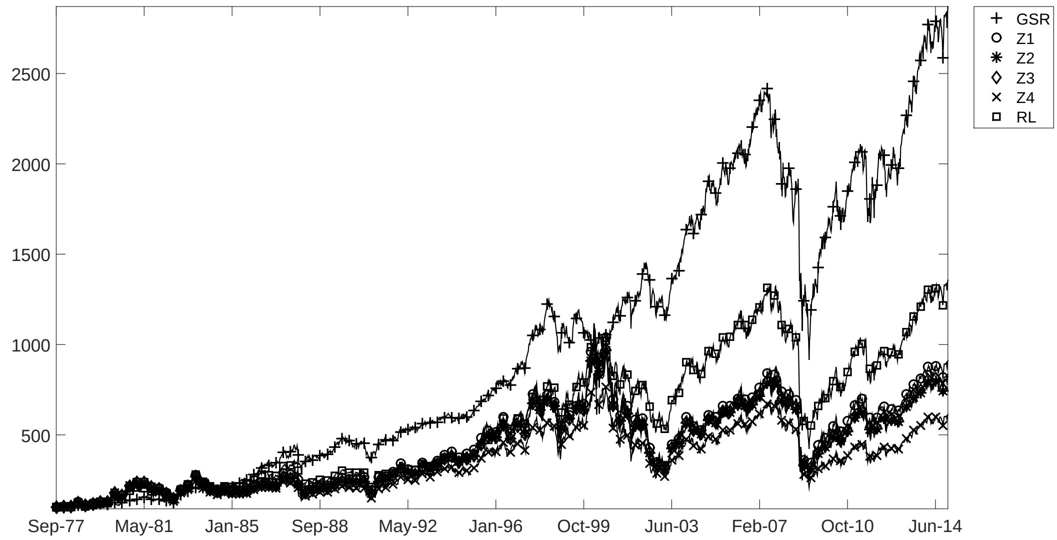

Figure 8: The cumulative returns of the agents' portfolio: the rational learner (RL) agent, the rational agent (GSR), the Fishburn agent (Z1), the S-shaped agent (Z2), the Disappointment theory (Z3), and the Farinelli-Tibiletti (Z4). The considered period is 1977-2014 (weekly frequency).

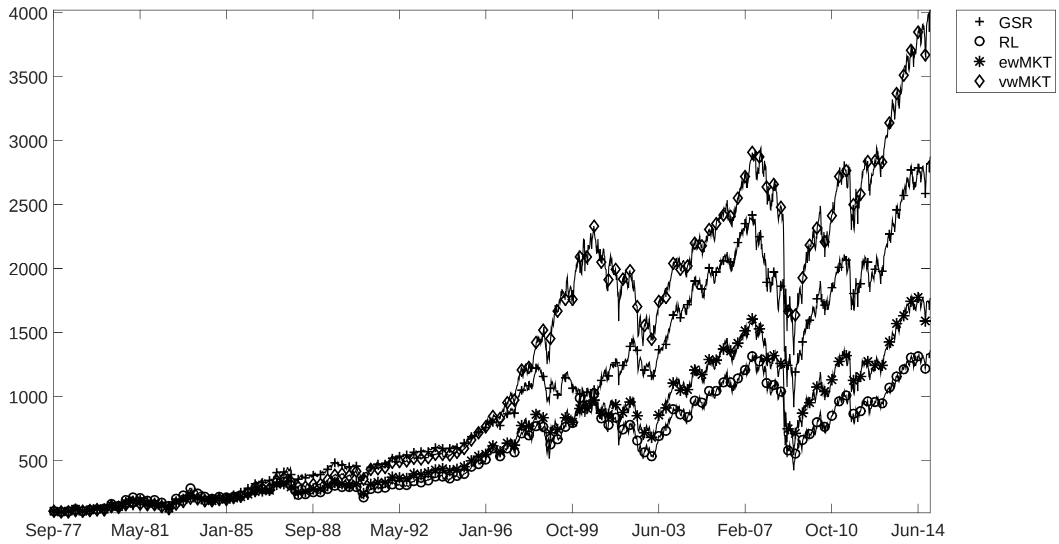

Figure 9: The cumulative returns of the value weighted market index (vw) and equally weighted market index (ew), the rational learner (RL) agent and the rational agent (GSR). The considered period is 1977-2014 (weekly frequency).

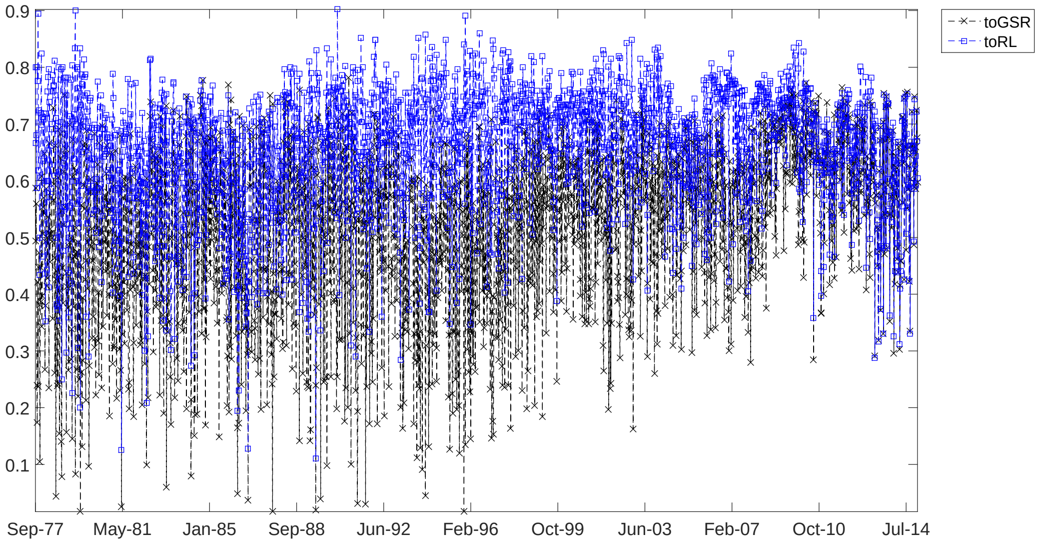

Figure 10: Turn-over of the agents' portfolio: the rational learner (RL) agent and the rational agent (GSR). The considered period is 1977-2014 (weekly frequency).

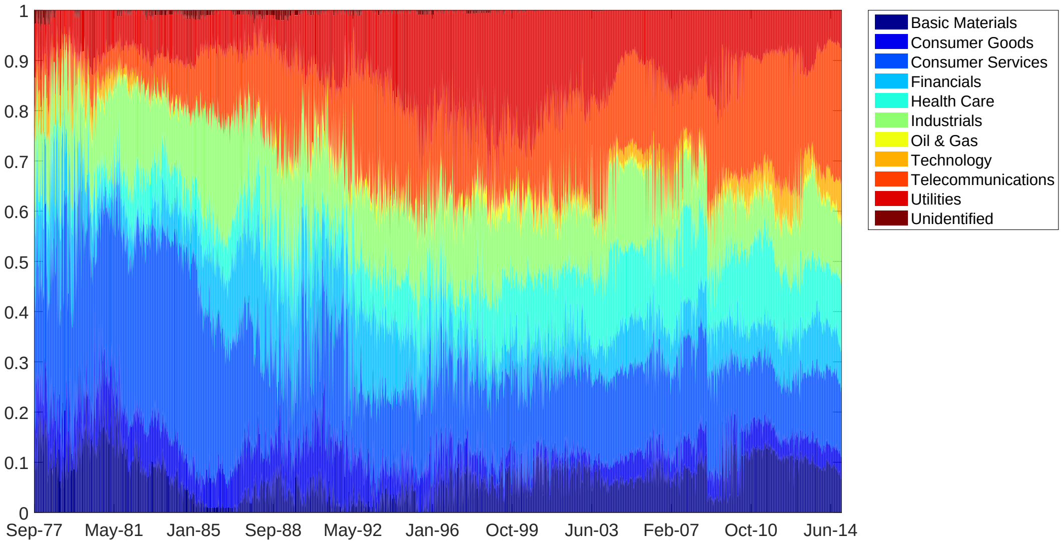

Figure 11: Portfolio's composition according to the Industrial composition benchmark (ICB) Supersectors for the rational learner agent (RL). The considered period is 1977-2014 (weekly frequency).

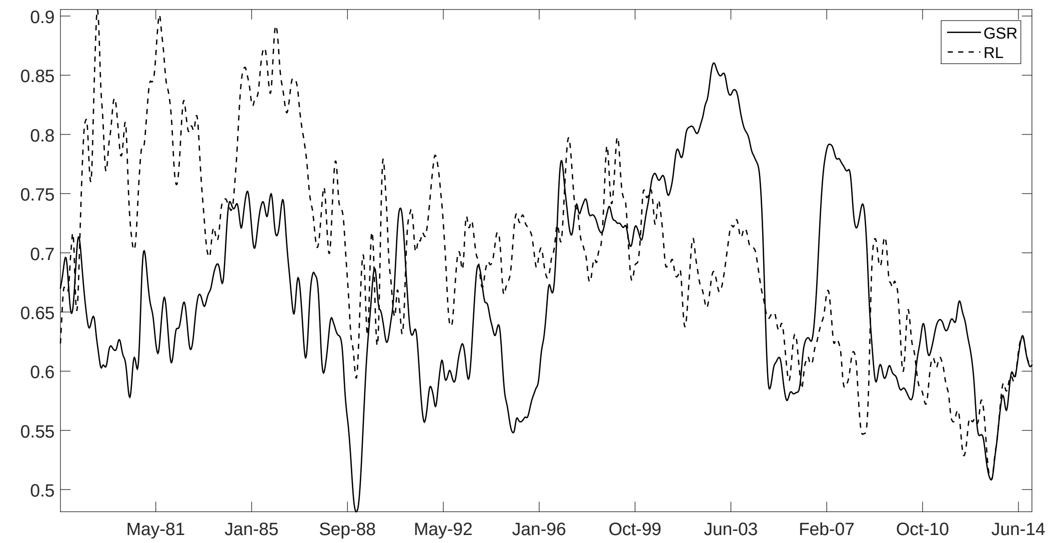

Figure 12: Diversification index on Portfolio's composition according to the Industrial composition benchmark for the rational agent (GSR) and the rational learner agent (RL). The considered period is 1977-2014 (weekly frequency).

Table 1: Results from regressing the estimated volatility of the factor τ on market returns, volatility and the investment universe dimension. The period is from September 1972 to December 2014 (1947 observations, weekly frequency).

	σ_τ			
(Intercept)	12.5800***	12.8952***	11.8015***	12.3765***
	(0.0394)	(0.0932)	(0.0602)	(0.0907)
r_{mkt}	1.2597	-	-	1.7392
	(1.4969)			(1.3790)
σ_{mkt}	-	-11.7593***	-	-25.5001***
		(3.1758)		(3.0300)
n_t	-	-	0.0007***	0.0008***
			(0.0000)	(0.0000)
R-squared	0.0004	0.0070	0.1212	0.1527
adj-R-squared	-0.0001	0.0065	0.1207	0.1514
LL	-3833.09	-3826.61	-3707.68	-3672.14
AIC	7670.18	7657.21	7419.36	7352.27
BIC	7681.33	7668.36	7430.51	7374.57

Table 2: Descriptive analysis of realized portfolio returns. The columns report the cumulated returns obtained during the period 1977-2014, the annualized average monthly return and annualized variance, the minimum and maximum monthly returns, the skewness and kurtosis indices computed on weekly returns, the 5% value-at-risk and expected shortfall, and the Sharpe ratio. We report the rational agent (GSR), the rational learner agent (RL), the Fishburn agent (Z1), the S-shaped agent (Z2), the agent according to the Disappointment theory (Z3), and the Farinelli and Tibiletti agent (Z4). For the market we include the equally weight market index (ewMKT) and the value weighted market index.

	cum.rets	std.ann	Min	Max	Skew	Kurtosis	VaR	ES	Sharpe
GSR	27.7058	0.1564	-0.1894	0.1543	-0.5425	10.5897	-0.0309	-0.0493	0.0905
Z1	8.1110	0.2436	-0.2289	0.2107	-0.5622	8.0903	-0.0528	-0.0812	0.0508
Z2	7.2380	0.2384	-0.2248	0.1994	-0.5742	8.0579	-0.0520	-0.0796	0.0496
Z3	7.3709	0.2386	-0.2248	0.1994	-0.5749	8.0512	-0.0516	-0.0797	0.0498
Z4	5.1328	0.2244	-0.2214	0.1674	-0.6356	7.8175	-0.0490	-0.0750	0.0457
RL	12.5581	0.2061	-0.1980	0.1499	-0.5567	7.5436	-0.0449	-0.0681	0.0613
ewMKT	16.6575	0.1878	-0.1801	0.1774	-0.5690	8.9472	-0.0391	-0.0616	0.0698
vwMKT	39.2019	0.1686	-0.1785	0.1244	-0.4466	7.4296	-0.0347	-0.0531	0.0930

Table 3: Descriptive analysis of realized portfolio returns. The columns report the cumulated returns obtained in each year during the period 1977-2014, the annualized average monthly return and annualized variance, the minimum and maximum monthly returns, the skewness and kurtosis indices computed on weekly returns, the 5% value-at-risk and expected shortfall, and the Sharpe ratio. We report the rational agent (GSR) and the rational learner agent (RL). For the market we include the equally weight market index (ewMKT) and the value weighted market index.

		cum.rets	**std.ann**	**VaR**	**ES**	**Sharpe**
1977	*GSR*	-0.0160	0.1329	-0.0241	-0.0276	-0.0429
	RL	-0.0076	0.1572	-0.0291	-0.0339	-0.0106
	ewMKT	0.0009	0.1547	-0.0281	-0.0336	0.0124
	vwMKT	-0.0179	0.1459	-0.0268	-0.0286	-0.0431
1978	*GSR*	0.0606	0.1485	-0.0330	-0.0459	0.0651
	RL	0.0311	0.1896	-0.0421	-0.0749	0.0357
	ewMKT	0.0525	0.1721	-0.0336	-0.0602	0.0532
	vwMKT	0.0628	0.1595	-0.0350	-0.0468	0.0639
1979	*GSR*	0.1661	0.1229	-0.0208	-0.0386	0.1821
	RL	0.3746	0.1729	-0.0325	-0.0520	0.2679
	ewMKT	0.2775	0.1457	-0.0259	-0.0452	0.2437
	vwMKT	0.1939	0.1296	-0.0242	-0.0366	0.1988
1980	*GSR*	0.1401	0.1576	-0.0350	-0.0413	0.1263
	RL	0.5085	0.2288	-0.0542	-0.0588	0.2657
	ewMKT	0.3430	0.1942	-0.0409	-0.0492	0.2244
	vwMKT	0.3922	0.2019	-0.0383	-0.0537	0.2417
1981	*GSR*	0.0090	0.1345	-0.0316	-0.0329	0.0184
	RL	-0.1227	0.2175	-0.0544	-0.0646	-0.0686
	ewMKT	-0.0993	0.1711	-0.0475	-0.0522	-0.0730
	vwMKT	-0.1404	0.1669	-0.0410	-0.0474	-0.1142
1982	*GSR*	0.2119	0.1960	-0.0309	-0.0341	0.1467
	RL	0.0706	0.2779	-0.0512	-0.0574	0.0521
	ewMKT	0.0783	0.2337	-0.0411	-0.0466	0.0596
	vwMKT	0.1076	0.2294	-0.0396	-0.0470	0.0760
1983	*GSR*	0.2338	0.1105	-0.0201	-0.0240	0.2717
	RL	0.2249	0.1805	-0.0356	-0.0526	0.1684
	ewMKT	0.1861	0.1339	-0.0286	-0.0361	0.1861
	vwMKT	0.1941	0.1327	-0.0237	-0.0326	0.1947
1984	*GSR*	-0.0041	0.1447	-0.0284	-0.0324	0.0058
	RL	-0.1670	0.2331	-0.0466	-0.0484	-0.0930
	ewMKT	-0.0768	0.1705	-0.0355	-0.0377	-0.0535
	vwMKT	-0.0170	0.1528	-0.0291	-0.0349	-0.0053

		cum.rets	**std.ann**	**VaR**	**ES**	**Sharpe**
1985	*GSR*	0.3493	0.0986	-0.0173	-0.0221	0.4291
	RL	0.2157	0.1505	-0.0298	-0.0376	0.1905
	ewMKT	0.2195	0.1175	-0.0225	-0.0275	0.2427
	vwMKT	0.2657	0.1086	-0.0189	-0.0245	0.3089
1986	*GSR*	0.2247	0.1555	-0.0365	-0.0502	0.1918
	RL	0.0823	0.1757	-0.0420	-0.0583	0.0747
	ewMKT	0.1179	0.1570	-0.0375	-0.0521	0.1094
	vwMKT	0.1643	0.1641	-0.0409	-0.0545	0.1400
1987	*GSR*	-0.0009	0.2300	-0.0634	-0.0919	0.0157
	RL	-0.1230	0.3042	-0.0809	-0.1254	-0.0376
	ewMKT	-0.0628	0.2794	-0.0761	-0.1175	-0.0120
	vwMKT	0.0234	0.2511	-0.0708	-0.1002	0.0305
1988	*GSR*	0.1299	0.1253	-0.0284	-0.0322	0.1412
	RL	0.0514	0.1249	-0.0297	-0.0352	0.0631
	ewMKT	0.1094	0.1316	-0.0314	-0.0343	0.1164
	vwMKT	0.0690	0.1475	-0.0320	-0.0371	0.0716
1989	*GSR*	0.2121	0.1096	-0.0275	-0.0354	0.2513
	RL	0.1949	0.1203	-0.0281	-0.0435	0.2139
	ewMKT	0.1693	0.1189	-0.0291	-0.0437	0.1909
	vwMKT	0.2672	0.1299	-0.0309	-0.0420	0.2623
1990	*GSR*	-0.1417	0.1505	-0.0401	-0.0452	-0.1303
	RL	-0.2109	0.1927	-0.0544	-0.0592	-0.1568
	ewMKT	-0.1704	0.1791	-0.0453	-0.0534	-0.1321
	vwMKT	-0.0517	0.1630	-0.0405	-0.0449	-0.0340
1991	*GSR*	0.2352	0.1212	-0.0191	-0.0211	0.2503
	RL	0.2963	0.1765	-0.0301	-0.0306	0.2163
	ewMKT	0.3167	0.1568	-0.0227	-0.0248	0.2545
	vwMKT	0.2842	0.1393	-0.0235	-0.0256	0.2589
1992	*GSR*	0.1269	0.0868	-0.0153	-0.0183	0.1969
	RL	0.1050	0.1178	-0.0226	-0.0231	0.1256
	ewMKT	0.1187	0.1070	-0.0198	-0.0213	0.1528
	vwMKT	0.0850	0.0940	-0.0206	-0.0215	0.1269

		cum.rets	std.ann	VaR	ES	Sharpe
1993	*GSR*	0.0735	0.0675	-0.0119	-0.0173	0.1476
	RL	0.1248	0.1131	-0.0213	-0.0292	0.1494
	ewMKT	0.1019	0.0848	-0.0164	-0.0225	0.1617
	vwMKT	0.0512	0.0802	-0.0178	-0.0228	0.0901
1994	*GSR*	-0.0113	0.0837	-0.0252	-0.0283	-0.0131
	RL	-0.0152	0.1393	-0.0357	-0.0429	-0.0057
	ewMKT	-0.0137	0.1164	-0.0297	-0.0391	-0.0085
	vwMKT	0.0126	0.1011	-0.0239	-0.0307	0.0241
1995	*GSR*	0.2590	0.0622	-0.0111	-0.0154	0.5186
	RL	0.3423	0.1116	-0.0177	-0.0239	0.3743
	ewMKT	0.2980	0.0866	-0.0191	-0.0195	0.4244
	vwMKT	0.3574	0.0716	-0.0106	-0.0120	0.5985
1996	*GSR*	0.1502	0.1033	-0.0205	-0.0233	0.1951
	RL	0.1944	0.1334	-0.0284	-0.0344	0.1940
	ewMKT	0.1593	0.1190	-0.0199	-0.0284	0.1805
	vwMKT	0.2605	0.1306	-0.0241	-0.0255	0.2552
1997	*GSR*	0.2380	0.1162	-0.0190	-0.0244	0.2632
	RL	0.1254	0.1563	-0.0303	-0.0420	0.1155
	ewMKT	0.1599	0.1432	-0.0271	-0.0393	0.1535
	vwMKT	0.2632	0.1567	-0.0301	-0.0326	0.2178
1998	*GSR*	-0.0107	0.1645	-0.0385	-0.0464	0.0021
	RL	0.0080	0.2704	-0.0747	-0.0891	0.0227
	ewMKT	0.0234	0.2376	-0.0567	-0.0684	0.0297
	vwMKT	0.3595	0.1960	-0.0486	-0.0535	0.2312
1999	*GSR*	0.0000	0.1430	-0.0294	-0.0377	0.0097
	RL	0.4048	0.1727	-0.0379	-0.0436	0.2803
	ewMKT	0.1831	0.1504	-0.0341	-0.0407	0.1627
	vwMKT	0.2795	0.1807	-0.0402	-0.0460	0.1982
2000	*GSR*	0.1012	0.1656	-0.0319	-0.0374	0.0920
	RL	-0.0980	0.3147	-0.0600	-0.0978	-0.0235
	ewMKT	0.0106	0.2210	-0.0432	-0.0628	0.0217
	vwMKT	-0.0085	0.2334	-0.0456	-0.0717	0.0109

		cum.rets	std.ann	VaR	ES	Sharpe
2001	*GSR*	0.0989	0.1731	-0.0247	-0.0629	0.0878
	RL	-0.0700	0.2974	-0.0576	-0.0998	-0.0129
	ewMKT	0.0283	0.2508	-0.0439	-0.0890	0.0332
	vwMKT	-0.0590	0.2272	-0.0413	-0.0763	-0.0213
2002	*GSR*	-0.0364	0.1449	-0.0392	-0.0534	-0.0254
	RL	-0.2707	0.2237	-0.0609	-0.0643	-0.1798
	ewMKT	-0.2130	0.1950	-0.0571	-0.0633	-0.1565
	vwMKT	-0.2203	0.1934	-0.0481	-0.0635	-0.1648
2003	*GSR*	0.2961	0.1262	-0.0272	-0.0311	0.2942
	RL	0.4651	0.1880	-0.0344	-0.0440	0.2955
	ewMKT	0.4304	0.1707	-0.0337	-0.0415	0.3033
	vwMKT	0.2882	0.1526	-0.0340	-0.0389	0.2410
2004	*GSR*	0.2073	0.1077	-0.0185	-0.0216	0.2458
	RL	0.1619	0.1618	-0.0323	-0.0460	0.1374
	ewMKT	0.1752	0.1525	-0.0281	-0.0422	0.1547
	vwMKT	0.1371	0.1121	-0.0278	-0.0321	0.1637
2005	*GSR*	0.0159	0.1101	-0.0271	-0.0337	0.0273
	RL	0.0618	0.1385	-0.0289	-0.0412	0.0695
	ewMKT	0.0472	0.1360	-0.0309	-0.0425	0.0563
	vwMKT	0.0548	0.0999	-0.0238	-0.0276	0.0809
2006	*GSR*	0.1677	0.1075	-0.0217	-0.0261	0.2076
	RL	0.1357	0.1473	-0.0362	-0.0423	0.1300
	ewMKT	0.1432	0.1422	-0.0360	-0.0399	0.1403
	vwMKT	0.1323	0.1047	-0.0250	-0.0285	0.1718
2007	*GSR*	-0.1148	0.1492	-0.0453	-0.0516	-0.1029
	RL	0.0213	0.1714	-0.0467	-0.0565	0.0288
	ewMKT	-0.0422	0.1630	-0.0464	-0.0548	-0.0255
	vwMKT	0.0487	0.1369	-0.0377	-0.0446	0.0575
2008	*GSR*	-0.3845	0.3838	-0.0923	-0.1322	-0.1477
	RL	-0.5169	0.3918	-0.1140	-0.1386	-0.2285
	ewMKT	-0.4669	0.4194	-0.1229	-0.1445	-0.1778
	vwMKT	-0.4047	0.3348	-0.0857	-0.1217	-0.1903

		cum.rets	std.ann	VaR	ES	Sharpe
2009	*GSR*	0.3507	0.2878	-0.0539	-0.0823	0.1649
	RL	0.2835	0.3380	-0.0663	-0.0808	0.1254
	ewMKT	0.3556	0.3388	-0.0725	-0.0933	0.1480
	vwMKT	0.3611	0.2644	-0.0494	-0.0634	0.1800
2010	*GSR*	0.1599	0.1652	-0.0440	-0.0512	0.1337
	RL	0.2455	0.2277	-0.0633	-0.0735	0.1472
	ewMKT	0.2216	0.2218	-0.0600	-0.0701	0.1383
	vwMKT	0.1497	0.1783	-0.0452	-0.0557	0.1189
2011	*GSR*	-0.0489	0.2278	-0.0612	-0.0702	-0.0150
	RL	-0.0651	0.2518	-0.0561	-0.0801	-0.0197
	ewMKT	-0.0837	0.2696	-0.0670	-0.0836	-0.0265
	vwMKT	-0.0234	0.2237	-0.0502	-0.0665	0.0006
2012	*GSR*	0.0653	0.1503	-0.0365	-0.0425	0.0687
	RL	0.0791	0.1368	-0.0326	-0.0400	0.0866
	ewMKT	0.0965	0.1503	-0.0361	-0.0438	0.0953
	vwMKT	0.1039	0.1218	-0.0270	-0.0354	0.1209
2013	*GSR*	0.3599	0.1202	-0.0239	-0.0252	0.3639
	RL	0.3497	0.1171	-0.0242	-0.0253	0.3641
	ewMKT	0.3743	0.1204	-0.0240	-0.0265	0.3756
	vwMKT	0.3022	0.1017	-0.0208	-0.0217	0.3677
2014	*GSR*	0.0637	0.1214	-0.0302	-0.0327	0.0789
	RL	0.0659	0.1217	-0.0288	-0.0326	0.0811
	ewMKT	0.0279	0.1339	-0.0323	-0.0378	0.0376
	vwMKT	0.0997	0.1133	-0.0276	-0.0325	0.1241

Appendix A. Portfolios' Composition.

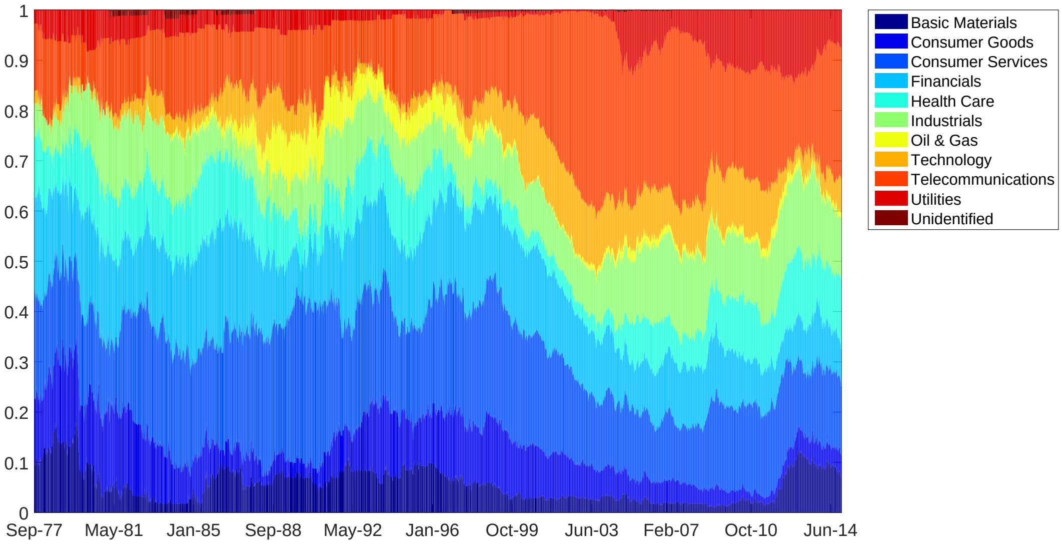

Figure A.13: Portfolio's composition according to the Industrial composition benchmark (ICB) Supersectors for the rational agent (GSR). The considered period is 1977-2014 (weekly frequency).

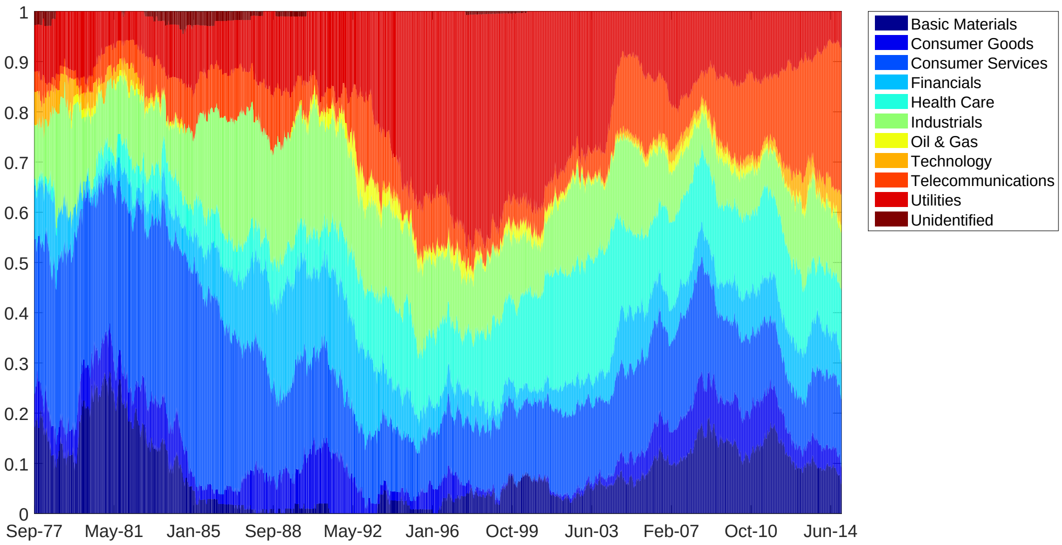

Figure A.14: Portfolio's composition according to the Industrial composition benchmark (ICB) Supersectors for the Fishburn agent (Z1). The considered period is 1977-2014 (weekly frequency).

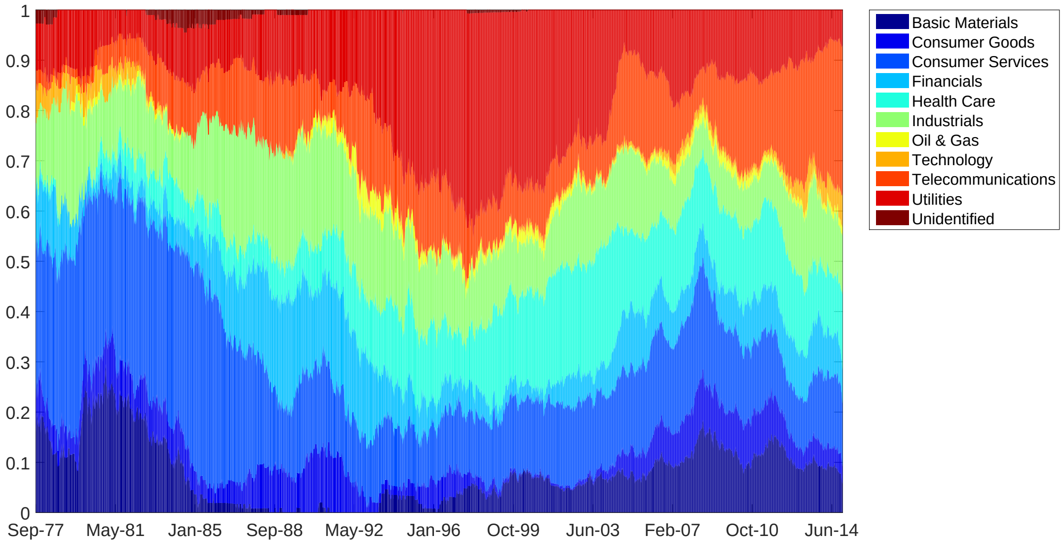

Figure A.15: Portfolio's composition according to the Industrial composition benchmark (ICB) Supersectors for the S-shaped agent (Z2). The considered period is 1977-2014 (weekly frequency).

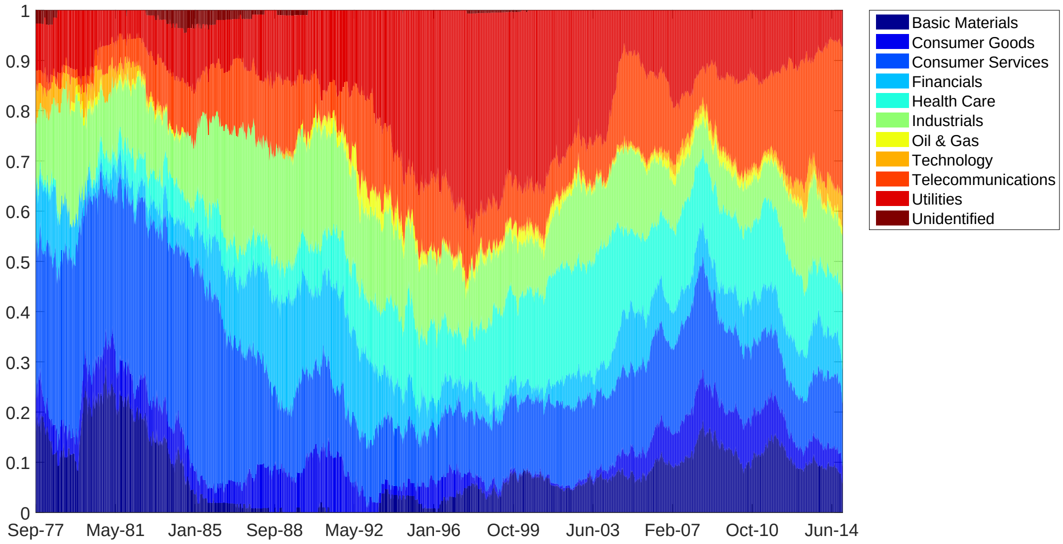

Figure A.16: Portfolio's composition according to the Industrial composition benchmark (ICB) Supersectors for the agent endowed with the utility for disappointment theory (Z3). The considered period is 1977-2014 (weekly frequency).

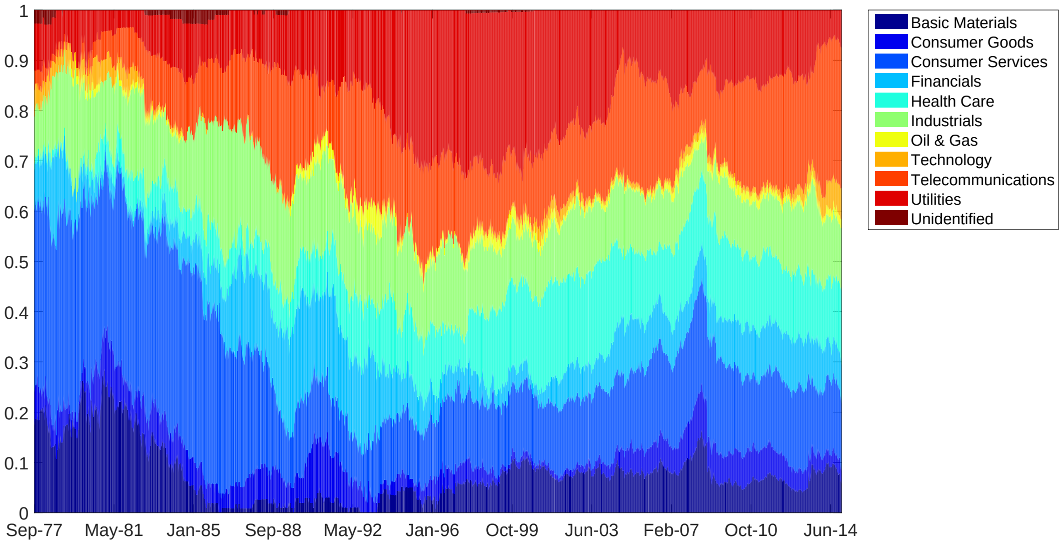

Figure A.17: Portfolio's composition according to the Industrial composition benchmark (ICB) Supersectors for the Farinelli-Tibiletti agent (Z4). The considered period is 1977-2014 (weekly frequency).