

Relative risk aversion and wealth dynamics

Shu-Heng Chen ^{a,*}, Ya-Chi Huang ^b

^a *AI-ECON Research Center, Department of Economics, National Chengchi University, Taipei 116, Taiwan*

^b *Department of International Trade, Lunghwa University of Science and Technology, Taoyuan 333, Taiwan*

Abstract

As a follow-up to the work of Chen and Huang [S.-H. Chen, Y.-C. Huang, Risk preference, forecasting accuracy and survival dynamics: simulations based on a multi-asset agent-based artificial stock market, Working Paper Series 2004-1, AI-ECON Research Center, National Chengchi University, 2004; S.-H. Chen, Y.-C. Huang, Risk preference and survival dynamics, in: T. Terano, H. Kita, T. Kaneda, K. Arai, H. Deghchi (Eds.), *Agent-Based Simulation: From Modeling Methodologies to Real-World Applications*, Springer Series on Agent-Based Social Systems, vol. 1, 2005, pp. 135–143], this paper continues to explore the relationship between wealth share dynamics and risk preferences in the context of an agent-based multi-asset artificial stock market. We simulate a multi-asset agent-based artificial stock market composed of heterogeneous agents with different degrees of relative risk aversion. As before, we find that the difference in risk aversion and the resultant saving behavior are the primary forces in determining the survivability of agents. In addition to the stability of the saving behavior, the level of the saving rate also plays a crucial role. The agents with stable saving behavior, e.g., the log-utility agents, may still become extinct because of their low saving rates, whereas the agents with unstable saving behavior may survive because of their high saving rates, implied by their highly risk-averse preferences.

© 2006 Elsevier Inc. All rights reserved.

Keywords: Risk preferences; CRRA (constant relative risk aversion); Blume–Easley theorem; Agent-based artificial stock markets; Genetic algorithms

1. Motivation and introduction

The contribution of *risk preference* to the survivability (wealth share) of investors has recently received a series of theoretic and simulation studies (e.g., [1,18,17,2,4,5]). The results are mixed, depending on how we approach this issue. While the standard analytic approach proves the irrelevance of risk preference to survivability [17,2], the agent-based computational approach indicates the opposite [4,5]. This kind of inconsistency, as quite often seen in the agent-based computational economics literature, simply reflects the sensitivity of the classical (analytical) results to the interacting heterogeneous boundedly-rational behavior.

* Corresponding author. Tel.: +886 2 29387308; fax: +886 2 29387308/27386874.

E-mail addresses: chchen@nccu.edu.tw (S.-H. Chen), yachi@aiecon.org (Y.-C. Huang).

URL: <http://www.aiecon.org> (S.-H. Chen).

Ref. [4] actually supports an earlier result obtained in [1], which is also known as the *Kelly criterion* in financial economics. This result basically points out the optimal type of risk preference, namely, the *CRRRA* (*constant relative risk aversion coefficient*) agent with an *RRA* coefficient of one. Equivalently, it is the log utility function. [5] reestablishes this result, while in an agent-based computational setting. They examine the long-run wealth share dynamics of eleven different types of CRRRA agents, with RRA coefficients ranging from 0 to 1 with increments of 0.1. They find that in finite time the wealth share is positively related to the CRRRA coefficient, and in the long run, only the agents with high CRRRA coefficients can survive. All others become extinct.

This paper is an extension of [5] in the sense that we wish to extend the earlier testing domain of the CRRRA coefficient from [0, 1] to an even larger positive domain. In doing so, we are inquiring whether a higher degree of risk aversion can actually enhance the survivability of agents. Notice that the degree of risk aversion is not the original concern of either the Kelly criterion or the Blume–Easley theorem [1], both of which are only concerned with the dominance of the log-utility type agents. Risk aversion is involved because the log-utility agent is also known as a CRRRA type of agent with an RRA coefficient of one. Now, is this the optimal degree of risk aversion? Will more risk-averse agents (i.e., those with RRA coefficients greater than one) be also driven out of the market when they are competing with the log-utility agents? Or, would higher risk aversion help them survive? These are the questions that we try to answer in this paper.

We consider these questions to be particularly relevant because the empirical literature actually suggests a large range of relative risk aversion coefficients. Some of them are exactly one or less than one, but many more are greater than one. Of course, it is doubtful whether one can directly compare our results with those empirical values, since they refer to quite different stories. However, given the prevailing empirical results on high risk aversion, it is definitely useful to know what makes them so, and our agent-based computational setting can serve as a good starting point.

The remainder of this paper is organized as follows. Section 2 gives a brief introduction to a simple multi-asset model, which is originally used in [1] and later extended and modified by [17]. Section 3 presents the artificial multi-asset market, which is an agent-based version of the analytical model presented in Section 2. Section 4 gives the experimental design. To justify the range of the relative risk aversion coefficient considered in this paper, it starts with a brief review of the literature on the empirical estimation of the RRA coefficient in Section 4.1, followed by the setting of other control parameters in Section 4.2. The simulation results are presented and discussed in Section 5, and are followed by the concluding remarks in Section 6.

2. A simple multi-asset model

The simulations presented in this paper are based on an agent-based version of the multi-asset market as [1,17] have studied. The market is complete in the sense that the number of states is equal to the number of assets, say M . At each date t , the outstanding volume of each asset is exogenously fixed at one unit. There are I investors in the market, each indexed by i . At time t asset m will pay dividends w_m if the corresponding state m occurs, and 0 otherwise. The behavior of states follows a finite-state stochastic process, which does not have to be stationary. The dividends w_m will be distributed among the I investors proportionately according to their owned share of the respective asset. The dividends can only be either re-invested or consumed. Hoarding is prohibited. If agent i chooses to consume c , her satisfaction is measured by her utility function $u(c)$. This simple multi-asset market clearly defines an optimization problem for each individual.

$$\max_{\{\{\delta_{t+r}^i\}_{r=0}^\infty, \{\alpha_{t+r}^i\}_{r=0}^\infty\}} E \left\{ \sum_{r=0}^\infty (\beta^i)^r u^i(c_{t+r}^i) | B_t^i \right\} \tag{1}$$

$$\text{subject to } c_{t+r}^i + \sum_{m=1}^M \alpha_{m,t+r}^{i,*} \cdot \delta_{t+r}^{i,*} \cdot W_{t+r-1}^i \leq W_{t+r-1}^i, \quad \forall r \geq 0, \tag{2}$$

$$\sum_{m=1}^M \alpha_{m,t+r}^i = 1, \quad \alpha_{m,t+r}^i \geq 0, \quad \forall r \geq 0. \tag{3}$$

In Eq. (1), u^i is agent i 's temporal utility function, and β^i , also called the discount factor, reveals agent i 's time preference. The expectation $E(\cdot)$ is taken with respect to the most recent belief B_t^i , which is a probabilistic

model used to represent agent i 's subjective belief regarding the stochastic nature of the state. The maximization problem asks for two sequences of decisions, one on saving, and one on portfolios, as denoted by

$$\{\{\delta_{t+r}^i\}_{r=0}^{\infty}, \{\alpha_{t+r}^i\}_{r=0}^{\infty}\},$$

where δ_t^i is the saving rate at time t , and

$$\alpha_t^i = (\alpha_{1,t}^i, \alpha_{2,t}^i, \dots, \alpha_{M,t}^i)$$

is the portfolio comprising the M assets.

Eqs. (2) and (3) are the budget constraints. W_t^i is the wealth of agent i at time t , which is earned from the dividends paid at time t . Notice that these budget constraints do not allow agents to consume or invest by borrowing.

The equilibrium price $\rho_{m,t}$ is determined by equating the demand for asset m with the supply of asset m , i.e.,

$$\sum_{i=1}^I \frac{\alpha_{m,t}^{i,*} \cdot \delta_t^{i,*} \cdot W_{t-1}^i}{\rho_{m,t}} = 1, \quad m = 1, 2, \dots, M. \quad (4)$$

Rearranging Eq. (4), we obtain the market equilibrium price of asset m :

$$\rho_{m,t} = \sum_{i=1}^I \alpha_{m,t}^{i,*} \cdot \delta_t^{i,*} \cdot W_{t-1}^i. \quad (5)$$

3. The agent-based multi-asset artificial stock market

An *agent-based computational version* of the Blume–Easley–Sandroni standard multi-asset model is developed in [4]. One of the mainstays of the agent-based computational economics is *autonomous agents* [19]. The idea of autonomous agents was initially presented in [11]. Briefly, these agents are able to learn and to adapt to the changing environment without too much external intervention, say, from the model designer. Their behavior is very much endogenously determined by the environment which they are interacting with. Accordingly, it can sometimes be very difficult to trace and to predict the resulting outcome, known as the *emergent behavior*.

In this paper, we follow what was initiated in [11], and equip our agents with the genetic algorithm to cope with the finite-horizon stochastic dynamic optimization problem, (1)–(3). The GA is applied here at two different levels, a high level (learning level) and a low level (optimization level). First, at the high level, it is applied as a *belief-updating scheme*. This is about the B_t^i appearing in (1). Agents start with some initial beliefs regarding state uncertainty which are technically characterized by parametric models, say, Markov processes. However, agents do not necessarily confine themselves with just stationary Markov processes. Actually, they can never be sure whether the underlying process will remain unchanged over time. So, they stay alert to that possibility, and keep on trying different Markov processes with different time frames (time horizons). Specifically, each belief can be described as “a k th order Markov process appearing over the last d days and may continue”. These two parameters can be represented by a binary string, and a canonical GA is applied to evolve a population of these two parameters with a set of standard genetic operators.¹

Once the belief is determined, the low-level GA is applied to solve the stochastic dynamic optimization problem defined in (1)–(3). Basically, we use Monte Carlo simulation to generate many possible ensembles consistent with the given belief and use them to evaluate a population of investment plans composed of a saving rate and a portfolio. GA is then applied to evolve this population of candidates.²

In sum, the high-level GA finds an appropriate belief, and under that belief the low-level GA searches for the best decisions regarding savings and portfolios. This style of adaptive design combines *learning how to forecast* with *learning how to optimize*, a distinction made in [3]. These two levels of GA, however, do not repeat themselves with the same frequency. As a matter of fact, the belief-updating scheme is somewhat slow, whereas the numerical optimization scheme is more frequent. Intuitively, changing our belief in the meta-level of the world

¹ Details can be found in [4], Appendix A.2.

² Details can be found in [4], Appendix A.1.

tends to be slower and less frequent than just fine-tuning or updating some parameters associated with a given structure. In this sense, the idea of *incremental learning* is also applied to our design of autonomous agents.

4. Experimental design

4.1. Empirical RRA coefficients

This paper and the coming experimental design are both very much motivated by the existing intensive empirical studies on the RRA coefficient. Therefore, we would like to give a brief survey here to present a general flavor on the wide dispersion of various estimates. For convenience we shall denote the RRA coefficient by ρ .

Let us look at this from the left extreme, i.e., $\rho \leq 1$. While [5] already questions the survivability of the CRRA agents with ρ being less than 1, empirical studies supporting small ρ still exist. The lowest value of ρ is found in [15], which is only 0.3. A little above this is [10]'s estimate, that ranges between 0.3502 and 0.9903. The one that seems to best fit our early simulation result of [5] is [7], who find values of RRA of around unity. Instead of using real economic data, [9] derived their estimate from a laboratory experiment with human subjects, which lay between 0.6 and 1.4. At the other extreme, there are cases of $\rho \geq 10$, such as 18 in [16] and 30 in [13]. These estimates are so large and are puzzling to many economists, having become part of the very famous *equity premium puzzle*. Next to those big ones, [12] also find a double-digit RRA in their estimates, which is close to 12. Despite these two extremes, most estimates suggest a moderate range for ρ , and include the [1, 3] interval [8], the [2, 3] interval [14, 20], the [2, 5] interval [6], etc.

4.2. Control parameters

According to the brief review above, we have set ρ to lie in the [0.5, 5] interval. It starts from 0.5 with increments of 0.5 and continues up to 5. In this setting, we have a total of 10 types of CRRA agents, and assign five market participants to each type of agent. Hence, there are 50 agents in the market in total. There are also five assets available in the market ($M = 5$), corresponding to five states. Asset m pays dividends $6 - m$ ($m = 1, 2, \dots, 5$). Two stochastic processes are considered in the experiments, namely, *iid* and the *first-order Markov*. Each is employed for one half of the total number of runs. The parameters of these two stochastic processes are also randomly generated in such a way that the axioms of the probability function are satisfied. Parameter values for the low-level and high-level GA can be found in [4], 4.1. We have 100 independent runs of the same experiment, and each run lasts for 100 market periods ($T = 100$). The simulation is conducted using the software *AIE-ASM Version 5.0*. The software is written with *Delphi, Version 6.0*. It usually takes 3 h for a single run on a Pentium III 1000 computer with 256 MB of RAM.

5. Simulation results

5.1. Wealth share dynamics

Fig. 1 shows the time series of the wealth share among the 10 types of agents. Notice that each line is based on an average of 100 simulations. The diagram seems to clearly indicate that the higher the risk aversion coefficient, the higher the wealth share. For example, the wealth share of the agents with the highest risk aversion coefficient ($\rho = 5$) shows strong growth from 10% in the initial period to 35% in the final period. What is particularly striking is that, unlike what the Blume–Easley theorem [1] predicts, the log-utility agents ($\rho = 1$) do not survive well in this case. Instead, their wealth share keeps on declining toward zero. Agents with lower values of ρ share the same destiny. This result combined with our earlier result [5, Fig. 3] well articulates the contribution of risk aversion to survivability.

5.2. Saving rates

Following the analysis in [4], we raise the same question: what may cause the agents with high values of ρ to survive and those with low values to vanish? In [4], the difference lies in the saving behavior. It is mainly the

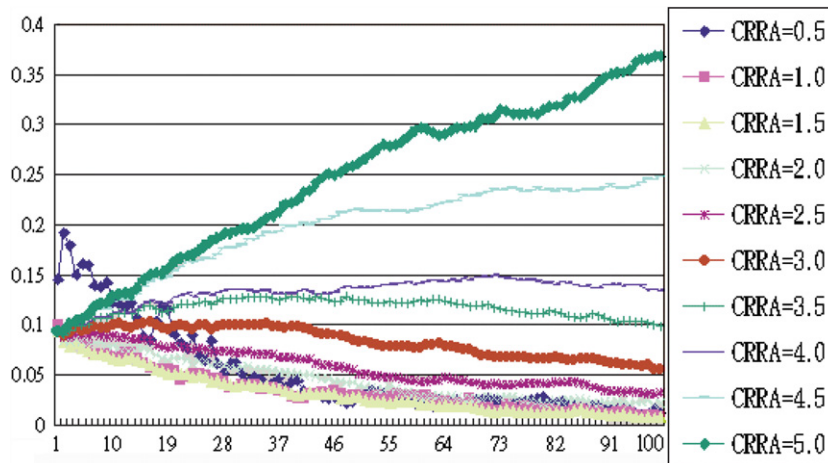


Fig. 1. Wealth share dynamics.

stable saving behavior that distinguishes log-utility agents from other types of agents. However, this argument is no longer applicable to our case. Fig. 2 is the box-whisker plot of the saving rates. Each plot shows the lifetime distribution of the saving rate δ_t associated with a specific RRA coefficient. To generate each plot, we first take an average of the saving rate of the five agents for the same ϱ . This is done period by period. A single history of $\delta_t (t = 1, 2, \dots, 100)$ is then derived by further taking an average over the entire 100 simulation runs. So, in the end, we have a single time history of δ_t for each ϱ . The ten boxes are drawn accordingly.

We first notice that the log-utility agent (the case where $\varrho = 1$) still has the most stable saving behavior with a very narrow niche around 0.45, which is exactly the discount rate (β) set in this paper. [1] shows analytically that $\delta_t = \beta$ for $\varrho = 1$. In other words, the optimal saving rate for agents with $\varrho = 1$ is just a constant and is the discount rate. Therefore, their saving behavior is independent of other economic variables, such as beliefs, wealth, asset prices or expected asset returns. Our genetic algorithm (the low-level GA) just confirms this property numerically. The property, however, is not shared by other types of agents. The most successful agents, agents with $\varrho = 5$, obviously have a box and whisker wider and longer than the log-utility agents, indicating a not very stable saving behavior. Hence, the dominance of a stable saving behavior as evidenced in [4] fails to carry over to this case.

However, why does the unstable saving behavior no longer cause a threat to the survival of the agents with a high RRA coefficient? And why can these agents drive out the log-utility agents? The answer probably lies in their high level of saving rates and the avoidance of extremely low saving rates. [1] already addressed the significance of the level of the saving rates, while their saving behavior is exogenously given. In our context, saving behavior is endogenously generated, and from the median, i.e., the line appearing in the middle of the




Fig. 2. Distribution of saving rates.

box in Fig. 2, it is found that the level of the saving rates increases with the RRA coefficient.³ It starts from 0.45 or so and increases up to 0.75. Therefore, although the agents of high risk aversion have very unstable saving behavior, their high long-run saving rates compensate for this weakness. Besides, what deserves to be mentioned further is that the high saving rate does not provide the entire compensation. To survive, agents must also avoid extremely low saving rates, i.e., the low whisker cannot come down to anything almost zero. Fortunately, an analysis using Eq. (6) may show that our CRRA-type agents will avoid this fatal point. As shown in Fig. 2, almost all of our agents have the minimum saving rate above 0.20. However, many other types of agents studied in [4] are not so lucky to step away from this dangerous zone.

Putting it together, we find that both the *first moment* and the *second moment* of saving behavior matter for the survivability issue. [1] does not consider the unstable saving behavior, so the second moment is neglected. [4] does not simulate the cases in which there exist some agents whose level of saving rates is sufficiently high and simultaneously without choosing an extremely low down-side saving rate, so the importance of the first moment is not revealed. This paper supplements [4] by integrating the effect of the first moment and the second moment.

Having said that, we compare the wealth share dynamics of two types of agents, i.e., agents with $\varrho = 1.5$ and $\varrho = 1$. From Fig. 2, we can see that agents with $\varrho = 1.5$ have a higher medium saving rate than agents with $\varrho = 1$, but they have a weaker stability. On the other hand, from Fig. 1, we actually see that the wealth share of the former ($\varrho = 1.5$) declines at almost the same speed as the latter ($\varrho = 1$), which suggests the substitution effect between the two moments: a sufficiently high level of saving rates can compensate a lower degree of instability so as to keep survivability unaffected.⁴

5.3. Portfolio performance

In addition to the saving rate, portfolio performance may be another contributing factor to survivability. However, this possibility has already been excluded in [4], and is excluded here again. Table 1 gives the three basic performance measures: the mean return, the risk (variance), and the Sharpe ratio.⁵ These statistics are averaged over the five agents of the identical type and are further averaged over the entire 100 simulation runs.

As we have seen in [4], the surviving agents do not have the highest rates of return. For example, surviving agents, agents with a ϱ of 4–5, have earned rates of return ranging from 1.22 to 1.16, which is almost one fourth or one third of those of the extinct agents ($\varrho = 0.5, 1$). Nonetheless, the column “variance of return” indicates that these agents are under different exposures to risk. Agents with higher returns are exposed to a higher risk and vice versa. This result is exactly what we expect. Agents with higher relative risk aversion coefficients choose to behave more prudently: low risk but also low return. On the other hand, agents with lower risk aversion coefficients choose to behave more aggressively: high return but also high risk.

Motivated by this finding, we go further to examine the *risk-adjusted return*, also known as the *Sharpe ratio*, and we find that, despite their low mean rate of return, the precautionary behavior of highly risk-averse agents actually helps them to earn a higher Sharpe ratio. However, there is no simple intuition to tell us why the agents with higher Sharpe ratios would survive. In particular, if we draw the risk-return plot in Fig. 3, we see that all these ten type of agents are standing on the *efficient frontier*. Strictly speaking, there is no one type of agent being strongly dominated by any other types of agent.⁶ Therefore, we see no particular reason to attribute their survivability to portfolio performance.

³ Analytically, via the *Euler equation* (6), it can be shown that in a standard continuous-time dynamic optimization context, the RRA coefficient is positively related to the saving rate.

$$r = \beta + \varrho \frac{\dot{c}(t)}{c(t)} \quad (6)$$

While our situation is much more complicated, that context may still provide us some useful insights to work with.

⁴ The substitution effect between the first and second moments obviously needs more careful analysis in the future. What is needed for supporting the existing empirical evidence is whether one can have a range of ϱ , e.g., $3 \leq \varrho \leq 5$, within this range, all agents survive, characterized by a large positive wealth share.

⁵ For the definition or calculation of these statistics, please see [4], Eqs. (21) and (22), for details.

⁶ Since all agents are equipped with the same soft computing technique, namely, GA, to solve their utility maximization problems, this result simply shows that GA is at least doing equally well for all types of agents. In fact, not only for the low-level GA, [4] shows that, in terms of forecasting accuracy, the high-level GA also performs equally well for all types of agents.

Table 1
Performance measurements

	RRA									
	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
Mean	4.479	3.303	2.182	1.713	1.484	1.360	1.281	1.222	1.181	1.162
Variance	109.923	53.816	18.923	10.779	7.666	6.261	5.349	4.688	4.382	4.351
Sharpe ratio	0.427	0.450	0.502	0.522	0.536	0.543	0.554	0.564	0.564	0.557

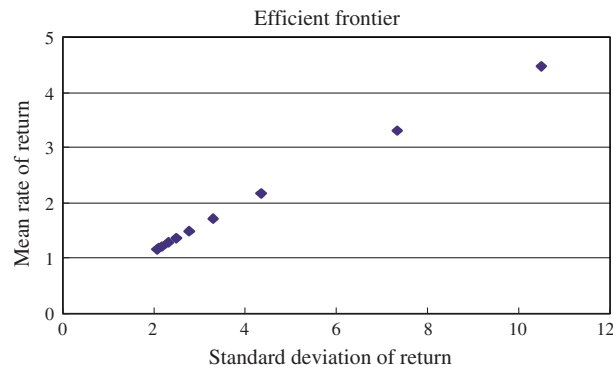


Fig. 3. Efficient frontier.

6. Concluding remarks

The basic conclusion that *risk preference matters* as established in [4,5] remains unchanged in this paper. The only thing modified here is *what causes it to matter*. The early analysis singles out the stable saving behavior as the distinguishing feature of the log-utility agents, and considers it to be the contributing factor to survivability. In this analysis, we add *the level of the saving rate* to it, and suggest that the level (the first moment) and the stability (the second moment) can both be contributing factors and in effect work together. When the first-moment effect dominates the second-moment effect, the log-utility agents become extinct and those who survive are the very risk-averse agents.

This study can be possibly related to the extensive literature on the empirical estimation of the RRA coefficients. A wide range of RRA coefficients has been found in the literature, and so far no unanimous conclusion has been reached. The agent-based computational approach as demonstrated in this paper proposes the possibility that in reality there may be such a wide *survival range* of the relative risk aversion coefficient. Within this range, the first-moment effect counter balances the second-moment effect, and all agents can survive with some positive and large wealth share.

Acknowledgements

The authors are grateful for the useful comments received from two referees. Research grant for this project in the form of NSC 91-2415-H-004-005 is also greatly acknowledged.

References

- [1] L. Blume, D. Easley, Evolution and market behavior, *Journal of Economic Theory* 58 (1992) 9–40.
- [2] L. Blume, D. Easley, If you're so smart, why aren't you rich? belief selection in complete and incomplete markets, Working paper, Department of Economics, Cornell University, 2001.
- [3] J. Bullard, J. Duffy, Using genetic algorithms to model the evolution of heterogeneous beliefs, *Computational Economics* 13 (1) (1999) 41–60.
- [4] S.-H. Chen, Y.-C. Huang, Risk preference, forecasting accuracy and survival dynamics: simulations based on a multi-asset agent-based artificial stock market, Working Paper Series 2004-1, AI-ECON Research Center, National Chengchi University, 2004.

- [5] S.-H. Chen, Y.-C. Huang, Risk preference and survival dynamics, in: T. Terano, H. Kita, T. Kaneda, K. Arai, H. Deghchi (Eds.), *Agent-Based Simulation: From Modeling Methodologies to Real-World Applications*, Springer Series on Agent-Based Social Systems, vol. 1, Springer, 2005, pp. 135–143.
- [6] G.M. Constantinides, J.B. Donaldson, R. Mehra, Junior can't borrow: a new perspective on the equity premium puzzle, *Quarterly Journal of Economics* 117 (2002) 269–297.
- [7] L.G. Epstein, S.E. Zin, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: an empirical analysis, *Journal of Political Economy* 99 (1991) 263–286.
- [8] I. Friend, M.E. Blume, The demand for risky assets, *American Economic Review* 65 (1975) 900–922.
- [9] M.J. Gordon, G.E. Paradis, C.H. Rorke, Experimental evidence on alternative portfolio decision rules, *American Economic Review* 62 (1972) 107–118.
- [10] L.P. Hansen, K. Singleton, Generalized instrumental variables estimation of nonlinear rational expectation models, *Econometrica* 50 (1982) 1269–1286.
- [11] J. Holland, J. Miller, Artificial adaptive agents in economic theory, *American Economic Review* 81 (2) (1991) 365–370.
- [12] P. Jorion, A. Giovannini, Time series tests of a non-expected-utility model of asset pricing, *European Economic Review* 37 (1993) 1083–1100.
- [13] S. Kandel, R.F. Stambaugh, Asset returns and intertemporal preferences, *Journal of Monetary Economics* 27 (1991) 39–71.
- [14] D. Lucas, Asset pricing with undiversifiable risk and short sales constraints: deepening the equity premium puzzle, *Journal of Monetary Economics* 34 (1994) 325–341.
- [15] G.N. Mankiw, J.J. Rotemberg, L.H. Summers, Intertemporal substitution in macroeconomics, *Quarterly Journal of Economics* 100 (1985) 225–251.
- [16] M. Obstfeld, Risk taking, global diversification, and growth, *American Economic Review* 84 (1994) 1310–1329.
- [17] A. Sandroni, Do markets favor agents able to make accurate predictions? *Econometrica* 68 (6) (2000) 1303–1341.
- [18] E. Sciuabba, *The evolution of portfolio rules and the capital asset pricing model*, DAE Working Paper No. 9909, University of Cambridge, 1999.
- [19] L. Tesfatsion, Introduction to the special issue on agent-based computational economics, *Journal of Economic Dynamics and Control* 25 (2001) 281–293.
- [20] S.P. Zeldes, Consumption and liquidity constraints: an empirical investigation, *Journal of Political Economy* 97 (2) (1989) 305–346.