Journal of Experimental Psychology: Applied
2006, Vol. 12, No. 4, 233–250

1076-898X/06/$12.00 DOI: 10.1037/1076-898X.12.4.233

Positive and Negative Recency Effects in Retirement Savings Decisions

Jörg Rieskamp
Max Planck Institute for Human Development and Indiana University Bloomington

Retirement savings decisions can be influenced by the fund composition of the retirement savings plan. In 2 experiments, strong composition effects were observed, with a larger percentage of resources being invested in stock funds when more stock than bond funds were offered. Although participants changed their allocations repeatedly, the opportunity to learn did not alter the composition effects. Learning processes led to positive and negative recency effects as well, providing evidence that allocations were strongly influenced by the recent performance of the different allocation options. Two learning models were tested to explain these learning processes. The first, a local adaptation learning model, assumes that people change their behavior on the basis of recent experience, whereas the second, a reinforcement learning model, assumes that decisions are made on the basis of the totality of accumulated experience. The local adaptation model was more accurate in predicting allocation decisions, in explaining positive and negative recency effects, and in showing why composition effects are not overcome by learning.

Keywords: retirement savings, resource allocation, reinforcement learning, local adaptation learning, 1/*n* strategy

Many retirement saving plans allow participants to make their own investment decisions (Huberman & Jiang, 2006). Yet the question of *how* participants decide—and whether the decisions they make are *good*—has been the subject of much debate. Benartzi and Thaler (2001) have argued that people often follow a naive diversification strategy of dividing their resources equally among the financial assets offered in a given savings plan. The first goal of the present article is to examine the extent to which such diversification strategies influence repeated savings decisions. The second goal is to test the ability of two learning models to predict these decisions.

Many retirement savings plans allow participants to decide how to allocate their resources among different financial assets, such as equity, bond, and property funds. Huberman and Jiang (2006) analyzed 650 different defined-contribution retirement savings plans with more than half a million participants. These diverse plans offered their participants the possibility to allocate resources between 4 and 59 funds, with at least 1 equity fund offered in all 650 plans and at least 1 bond fund offered in 620 of the plans. Because of the different types of funds offered in a retirement savings plan, participants face the problem of allocating their resources wisely. Past psychological research on retirement savings decisions generally examined how well people solve different types of resource allocation decision problems (Ball, Langholtz, Auble, & Sopchak, 1998; Busemeyer & Myung, 1987; Busemeyer, Swenson, & Lazarte, 1986; Fox, Ratner, & Lieb, 2005; Gingrich & Soli, 1984; Langholtz, Ball, Sopchak, & Auble, 1997; Langholtz, Gettys, & Foote, 1993, 1994, 1995; Northcraft & Neale; 1986; Rieskamp, Busemeyer, & Laine, 2003). The main finding from these studies was that when individuals make only a single resource-allocation decision with no opportunity to learn from experience, their decisions are often not as good as they could have been (e.g., Fox et al., 2005; Gingrich & Soli, 1984; Northcraft & Neale, 1986). This is not surprising given the complexity of most allocation problems. Individuals have a tendency to allocate an equal share of their resources to the different assets offered, which, depending on the situation, can lead to negative outcomes. However, in studies in which individuals were given the opportunity to improve their allocations through feedback, substantial learning effects were found, and individuals often approached optimal allocations (e.g., Busemeyer et al., 1986; Langholtz et al., 1993, 1994, 1995). Retirement savings decisions differ in some respects from the allocation decisions examined in the aforementioned studies, in which there were often constraints on the percentage of the resources that could be allocated to a single option (e.g., Langholtz et al., 1993, 1994, 1995), or the different options were interdependent in their performance outcomes (e.g., Rieskamp et al., 2003). These constraints often do not apply to retirement savings decisions (e.g., Benartzi & Thaler, 2001; Huberman & Jiang, 2006).

The main difficulty for retirement savings decisions consists in finding an allocation that corresponds to the investor's risk attitude (Huberman & Jiang, 2006). Often, investors have to decide how much of their resources to allocate to stock funds as opposed to property or bond funds. Investment options are commonly described by their expected rate of return and the expected variance

Jörg Rieskamp, Max Planck Institute for Human Development, Berlin, Germany, and Department of Psychology, Indiana University Bloomington.

Special thanks are due to the Max Planck Institute for Human Development and to the Center for the Study of Institutions, Population, and Environmental Change (Indiana University Bloomington) for financial support provided through National Science Foundation Grants SBR9521918 and SES0083511. I thank Tei Laine for programming the experimental software for Experiment 1 and Uwe Czienskowski for programming the software for Experiment 2. I also thank Donna Alexander for editing a draft of this article.

Correspondence concerning this article should be addressed to Jörg Rieskamp, Max Planck Institute for Human Development, Lentzeallee 94, 14195, Berlin, Germany. E-mail: rieskamp@mpib-berlin.mpg.de

in returns, the latter representing an option's *risk* (e.g., Markowitz, 1952; Sharpe, 1981). The problem is that the rate of return is typically positively correlated with the variance. Therefore, an investor needs to trade off the advantage of high expected returns against the disadvantage of high return variability (e.g., Sharpe, 1981). Mehra and Prescott (1985) analyzed the annual rate of return for equities and securities for the period 1889–1978. In this period, the average real return of Standard & Poor's 500 composite stock price index was 7%, with a standard deviation of 17%, compared with the average real return for relatively risk-free short-term securities of 1%, with a standard deviation of 6%. Thus, when considering the average past performance of stocks, they appear attractive but, because of their wider variability, bear the risk of large losses. The standard economic answer to this allocation problem is captured by the mean-variance model (see Markowitz, 1952), which is optimal when normal distributed returns and risk-averse investors are assumed (see Kroll, Levy, & Markowitz, 1984). An allocation is mean-variance efficient if it maximizes the expected rate of return for a given variance and minimizes the variance for a given expected return. Applying the mean-variance model eliminates inefficient allocations. Nevertheless, the efficient allocation ultimately selected by an individual should depend on the investor's risk attitude. A risk-averse investor could prefer investments in bonds that offer lower variability. Thus, the *risk attitude hypothesis* asserts that the percentage invested in stocks is correlated with the investor's risk attitude. This hypothesis was tested in the two present experiments.

In contrast to the risk attitude hypothesis, it has been argued that people's financial investments are influenced to a great extent by the context surrounding their decisions (for an overview, see Hirshleifer, 2001). Benartzi and Thaler (2001) asked employees of the University of California how they would allocate their retirement funds among five investment funds on offer. In the first condition, the five funds consisted of one stock and four fixed-income funds; in the second, they consisted of four stock and one fixed-income fund. The composition of funds offered had a strong effect on participants' decisions: When only one stock fund was offered, participants allocated an average of 43% of their resources to stocks, but when four stock funds were offered, they allocated an average of 68% to stocks. Benartzi and Thaler suggested that these individuals may have applied a naive diversification strategy to the investment problem—namely, the 1/*n* strategy. An investor applying this strategy allocates resources evenly across the *n* investment options offered, regardless of any particular option's expected performance or risk. When following this strategy, even very risk-averse investors will allocate a large percentage of their resources to stocks when a disproportionately large number of stock funds are offered in a retirement plan. The 1/*n* strategy may seem sensible particularly to unsophisticated investors, who may interpret the offer of retirement plans composed mainly of stock funds to mean that stock funds are the more appropriate form of retirement savings.

Benartzi and Thaler (2001) also examined composition effects in 170 real retirement plans with 1.56 million participants. They found that in plans with a relatively low percentage (only 37%) of equity options, the average allocation to stocks was 49%. In the case of plans with a relatively large percentage (81%) of equity options, however, participants allocated an average of 64% to stocks. These results suggest that the composition of investment options offered in a retirement plan influences retirement savings decisions. In contrast, Huberman and Jiang (2006), who also analyzed retirement savings plans, found less support for composition effects: First, participants did not allocate their resources evenly across all of the funds offered by the retirement savings plan. Instead, they often only selected a subset of funds for their investments (with a median of three funds), and the number of funds selected did not correlate with the number of funds offered by the plans. However, many participants did allocate their resources equally across the selected funds, and thus, they applied a "conditional 1/*n* strategy." However, because participants could select funds according to their risk attitude, the conditional 1/*n* strategy does not directly contradict the risk attitude hypothesis.

In the experiments reported by Benartzi and Thaler (2001), participants were not given the opportunity to revise their decisions. This might be explained by the fact that participants in retirement savings plans tend not to change their allocations frequently (see Agnew, Balduzzi, & Sunden, 2003; Samuelson & Zeckhauser, 1988). Nevertheless, it is necessary to examine how robust the composition effect is when people have the opportunity to revise their allocations. If participants in Benartzi and Thaler's experiments had received feedback about a fund's performance, they might have detected the performance advantage of stocks and changed their allocations accordingly. Similarly, Cutler, Poterba, and Summers (1990) have argued that investors change their investments on the basis of the results of past investments.

Kroll, Levy, and Rapoport (1988a; see also Kroll, Levy, & Rapoport, 1988b) conducted experiments studying whether people changed their allocation of resources when making repeated investment decisions and whether their allocations were consistent with the mean-variance model. Participants were provided with a fixed amount of resources that they repeatedly allocated to one of three assets: two different risky assets and one risk-free asset with a fixed rate of return. Participants were told that the returns of the risky assets were randomly drawn from a normal distribution, with a higher mean and variance for one of the two. Although the majority of allocations observed were mean-variance efficient, a substantial percentage (27%) were not. Moreover, for about 30% of their decisions, participants changed their allocations depending on the investment option's past returns, demonstrating a recency effect. Thus, when one asset had brought in higher returns in the previous and current period, many participants changed their allocations by switching to the opposite risky asset, exhibiting a *countercyclical investment strategy* leading to a negative recency effect. Other participants exhibited a *cyclical investment strategy*, switching to the asset that had achieved higher returns in the previous and current period, leading to a positive recency effect. Consistent with these results, Kroll et al. (1988a, 1988b) reported that in postexperiment questioning, many participants said that they had been looking for "trends" and "patterns." This finding is surprising from an economic perspective, because participants had been told that the rates of return were determined randomly and independently of each other. However, it is less surprising when considered in light of the substantial body of psychological research demonstrating that people perceive dependencies in random events (for early studies, see, e.g., Hake & Hyman, 1953; Jarvik, 1951; for an overview, see Myers, 1976). In probability learning studies, participants (in the simplest case) have to repeatedly predict which of two events will occur in the next trial. Even when the events are independent of each other, people make use of sequential information. This typically results in a positive recency effect: People tend to predict that the same event that occurred in the last trial will occur again in the next (e.g., Edwards, 1961; Jarvik, 1951). Yet negative recency effects have been observed as well (e.g., Estes, 1964; Winefiel, 1966; for an overview, see Ayton & Fischer, 2004). Overall, the above results indicate that people change their behavior on the basis of the success of their past behavior. However, particularly in the case of repeated investment decisions such as those studied by Kroll et al. (1988a, 1988b), no specific model has yet been proposed that can describe these changes. This is the second goal of the present article.

In the following section, two learning models are proposed to describe how people change their retirement savings allocations on the basis of experience. These models were tested in the two experiments reported below. In addition, the experiments examined whether the opportunity to learn will counteract any composition effects. If composition effects can be overcome by learning, this will help in developing training tools to assist people in making their retirement savings decisions. In contrast, if learning does not rule out composition effects and thereby shows their robustness, it is even more important to take these effects seriously.

Learning Models

According to the 1/*n* strategy, people allocate a resource equally among the available options. But how do they change their allocations over time? A strict view of the 1/*n* strategy implies that the allocation is kept constant for all following investment periods. However, it appears plausible that people might change their allocations on the basis of experience. Here, two approaches to predicting learning in resource allocation decisions are described.

The first approach, *reinforcement learning*, assumes that individuals form expectancies for the alternatives they choose by keeping track of the history of all previous decisions and by repeating decisions that worked well in the past. The general reinforcement learning idea can be traced back to early work by Bush and Mosteller (1955), Estes (1950), and Luce (1959; for more recent learning models, see Börgers & Sarin, 1997; Camerer & Ho, 1999a, 1999b; Erev, 1998; Erev & Barron, 2005; Erev & Roth, 1998; Harley, 1981; Rieskamp & Otto, 2006; Roth & Erev, 1995; Stahl, 1996; or Sutton & Barto, 1998).

The second approach, *local adaptation learning*, assumes that people compare the consequences of a current decision with a reference point (i.e., the previous decision or the most successful decision up to that point) and adjust their decisions in the direction of successful decisions. The main learning models that take this approach are the learning direction theory of Selten and Stöcker (1986), the error-correction learning models (e.g., Dorfman, Saslow, & Simpson, 1975; Thomas, 1973), and the hill-climbing learning model of Busemeyer and Myung (1987; see also Busemeyer & Myung, 1992). Rieskamp et al. (2003) have applied both approaches to an allocation problem with interdependent options. They were able to demonstrate that local adaptation learning described people's allocation decisions more accurately than reinforcement learning. In particular, local adaptation learning explained why many participants, even after extensive learning, converged at a local rather than the global payoff maximum. Thus, the local adaptation model supports the theory (see Selten, 1991) that learning does not necessarily lead people to find the optimal solution to a problem. The question of which learning approach better describes changes in retirement saving decisions, however, still remains open.

Reinforcement Learning

The basic idea of the reinforcement learning (RL) model is that decisions are made relative to an option's *expectancies*, which increase whenever the option produces a positive payoff. When applied to the retirement savings problem, the RL model assumes the following: Each investment fund is assigned an expectancy. First, allocations are made relative to the funds' expectancies. Second, the funds' returns are used to determine the reinforcement for updating the expectancies of each fund, and the process returns to the first step. Mathematically, RL is defined as follows: The preferences for each fund i in period t are expressed by expectancies $q_t(i)$. The percentage of resources invested in each fund in a certain period is defined by

$$PR_t(i) = \frac{\exp(\theta \cdot q_t[i])}{\sum_{i=1}^{N} \exp(\theta \cdot q_t[i])}, \tag{1}$$

where θ is a free *sensitivity parameter* restricted to positive values that determine the extent to which a fund that has a higher expectancy than other funds will receive a larger percentage of the available resources. For instance, with a very high value for the sensitivity parameter, total resources are invested in the fund with the highest expectancy. For the first period, all expectancies are assumed to be equal across all available funds and determined by the average return that can be expected from random choice, multiplied by w, which is the second free parameter —an *initial strength parameter*—restricted to positive values. After a decision, the expectancies are updated by the relative reinforcement of a decision—that is, the fund's rate of return, $r_{t-1}(i)$:

$$q_t(i) = (1 - \phi)q_{t-1}(i) + r_{t-1}(i), \tag{2}$$

where $\phi \in [0,1]$ is the third free parameter—a *forgetting rate parameter*. The forgetting rate determines how strongly previous expectancies affect new ones. With a high forgetting rate, recent reinforcements have a strong effect on the new expectancies, thereby potentially explaining positive recency effects. After the updating process, the percentage of resources invested in the available assets is determined again. In sum, the RL model has three free parameters: (a) the sensitivity parameter, θ; (b) the initial strength parameter, w; and (c) the forgetting rate parameter, ϕ. In addition, the RL model includes the 1/*n* strategy as a special case. The RL model makes identical predictions to the 1/*n* strategy in the case of a sensitivity parameter value of 0, so that, regardless of a fund's returns, the model will always predict that an equal percentage is invested in each fund.

Local Adaptation Learning

The basic idea of the local adaptation learning (LA) model is to start with a particular decision as a temporary solution and to change the allocation based on the outcome of the current decision compared with the previous decision, ignoring all other past decisions. Each investment fund is evaluated comparing its current return with its previous return. On the basis of this evaluation, people change their investments in the different funds depending on their investment strategy. The results of Kroll et al. (1988a) suggest that cyclical investment strategies should be distinguished from countercyclical investment strategies. Individuals with a cyclical (countercyclical) investment strategy increase (decrease) their investments in a fund when its return increases, and vice versa. A cyclical investment strategy can explain positive recency effects, and a countercyclical investment strategy can explain negative recency effects. The magnitude of the investment change—that is, the *step size*—depends on the relative magnitude of the changes in returns compared with the largest changes observed up to that point.

Mathematically, the LA model is defined as follows. In the first period, as with the RL model, resources are allocated equally across the available investment funds. In the second period, the percentage, $PR_2(i)$, of the resources invested in each fund is changed depending on the fund's returns relative to the average returns in the first period (because no previous returns are known). The step size in the second period is defined by $s_t(i) = r_{t-1}(i) \cdot s_1$, where $r_{t-1}(i)$ is the return on fund i in the first period, and s_1 is a free parameter called the *initial step size*. If the return on fund i was above the average return of all available funds, the percentage of the resources allocated to fund i is increased by $s_2(i)$ when following a cyclical investment strategy, whereas it decreases by $s_2(i)$ when following a countercyclical investment strategy, and vice versa when the fund performed below average.

In the third and all following periods, each fund is evaluated by comparing its current return with the return in the previous period. If the return increased, the percentage allocated to the fund is increased when following a cyclical investment strategy and decreased when following a countercyclical investment strategy, and vice versa when the fund's return decreased:

$$PR_t(i) = PR_{t-1}(i) + s_t(i) \cdot c, \tag{3}$$

where c is a free investment strategy parameter, set to the value 1 (expressing a cyclical investment strategy) or to the value -1 (expressing a countercyclical investment strategy). The step size of the new percentage is defined by

$$s_t(i) = \frac{r_{t-1}(\mathrm{i}) - r_{t-2}(i)}{\max[|r_{m+1}(j) - r_m(j)|;\ m = 1,\ldots,t-2;\ j = 1,\ldots,N]} \cdot s_1, \tag{4}$$

where the denominator determines the maximum change of returns observed across the series of changes for all funds so far. In the case in which the LA model predicts a percentage invested in one fund above or below the maximum of 100% or the minimum of 0%, the prediction is set to 100% or 0%. In addition, if the predicted percentages do not add up to 100%, they are normalized by dividing them by their total sum. In sum, LA has two free parameters: (a) the initial step size, s_1, which determines the baseline change of the investments, and (b) the investment strategy parameter, c, which determines how an individual investor reacts to the changes in a fund's returns. The LA model also includes the $1/n$ strategy as a special case. When the LA model's initial step size, s_1, has a value of 0, the model makes predictions identical to the $1/n$ strategy.

The Relationship of the Two Learning Models

The two models represent reasonable implementations of the two learning approaches considered here. Any empirical test of the two models, strictly speaking, only allows conclusions to be drawn about the empirical accuracy of these specific learning models. However, with this restriction in mind, both of these learning models are sufficiently flexible (in their free parameters) to predict a variety of learning processes. What are the different predictions that can be derived from the two learning models? In general, the RL model predicts that the percentage allocated to the different alternatives depends on the total store of previous reinforcements of these alternatives. This implies that, after sufficient learning opportunities, the alternatives that outperform the others on average should receive an increasing percentage of the decision maker's resources. In contrast, the LA model predicts that people only compare the outcome of the current period with the outcome of the previous period and ignore all other outcomes. This implies a strong path dependency, so that, depending on the starting point of the learning process and the magnitude of fluctuations of the funds' returns, learning does not necessarily lead to an increased percentage of resources being allocated to the best performing fund. However, the models' specific predictions depend on their parameters, so particular parameter values could lead to similar behavior for both models. For instance, if a high forgetting rate is used in the RL model, the funds' current returns strongly influence the succeeding allocation, implying a positive recency effect, which can also be predicted with the LA model. In contrast, the RL model cannot predict negative recency effects (see also Erev & Barron, 2005), whereas the LA can predict both positive *and* negative recency effects. Nevertheless, because of the models' flexibility, one can expect a relatively good fit for both when the parameters are fitted to the data. Therefore, the models were compared with a generalization test (Busemeyer & Wang, 2000), which necessitates a two-stage procedure. First, in Experiment 1, each model was fitted to the individual learning data. Second, the parameter distributions estimated in Experiment 1 were used to generate model predictions for the learning conditions presented in Experiment 2. The accuracies of the models' predictions for Experiment 2 provided the basis for a rigorous model comparison. In addition, because the two models include the $1/n$ strategy as a special case, both experiments also provided a test of whether it is worthwhile to consider learning processes for predicting people's allocations compared with the static $1/n$ strategy.

Experiment 1

The decision problem in Experiment 1 consisted of making retirement savings decisions by allocating resources for a fixed period of time among four investment funds. To examine composition effects, a varying relative number of stock funds compared with bond funds was offered. The goal was to test whether the composition of the funds influenced participants' initial allocations and whether a potential effect remained constant over time. In addition, participants' risk attitudes were assessed to test the risk attitude hypothesis. According to this hypothesis, the more risk-averse participants should invest less of their resources in stock funds. Finally, the two learning models were tested against each other and against the $1/n$ strategy to describe participants' allocations.

Method

Participants. Eighty persons (48 women and 32 men) with an average age of 22 years participated in the experiment. The computerized task lasted approximately 1 hr, followed by the risk attitude assessment that took approximately another 20 min. Participants were students in various departments at Indiana University Bloomington. For their participation, they received a show-up payment of $2. All additional payments depended on the participants' performance: They were paid a percentage of the final portfolio value, resulting in an average payment of $13.70. In addition, participants received a small payment for the risk assessment task.

Procedure. The experiment had two parts. Participants started with the retirement savings task, followed by the risk assessment task. Each participant was randomly assigned to one of four between-subjects conditions. The first between-subjects factor, *fund composition*, varied the number of stock funds offered. In the *three-stock fund condition*, three out of four funds offered were stock funds, whereas in the *two-stock fund condition*, only two out of four funds offered were stock funds. The remaining funds were always bond funds. The second between-subjects factor varied the *investment horizon*, with an investment period of 12 years in the *short-horizon condition* and with an investment period of 35 years in the *long-horizon condition*. Both between-subjects factors were completely crossed, providing four experimental conditions.

For the retirement savings task, the funds' returns were based on historical data. The returns of one stock fund were based on the inflation-adjusted annual total returns of large U.S. company stocks from 1960 to 1994 (or to 1971 for the short-horizon condition), and the returns of one bond fund were based on the inflation-adjusted annual total returns of intermediate-term government bonds for the same period (Ibbotson Associates, 2003, pp. 264–273). For the long-horizon condition, the average return was 8.67% (SD = 15.79%) for stocks and 3.02% (SD = 7.22%) for bonds. The returns of the other stock funds (or bond funds) were determined by randomly drawing an error from a normal distribution, with a mean of 2.0% (1.5%) and a standard deviation of 0.50% (SD = 0.25%), which was, with an equal probability, either subtracted from or added to the original returns of the stocks or bonds of the particular year. Thus, all stock funds (or bond funds) had similar returns.

The participants received the following instructions: They were asked to imagine they were making retirement savings decisions for a period of 35 years (or 12 years). They could invest $5,000 of their yearly income. Four different investment funds were available, among which they could spread their portfolio. The value of their portfolio each year would depend on the change in value of the investment funds they chose. Additionally, the value of their portfolio would increase with their yearly contribution of $5,000. Every year, they would be able to change their allocation. They were told that

> a fund usually pools your money with that of many other people who have similar investment goals. Professional money managers use the pool of money to buy individual securities, such as stocks and bonds. You make money if the fund shares grow in value and when the fund distributes profits to its investors.

They were informed that "in general, stocks represent ownership or equity in a company, whereas bonds represent a loan to a corporation or government agency." The stock funds were described as follows: "Stock Fund A: This stock fund invests in a broadly diversified range of U.S. stocks." The bond funds were described as follows: "Bond Fund A: This bond fund invests in government bonds." The other funds (*Stock Fund B* or *C*; *Bond Fund B*) were described in exactly the same way. Participants were also told that "normally, stocks outperform other types of investments over the long term. However, in the short term, stocks tend to have wider price fluctuations than other securities." The instructions quoted here are from the descriptions of a genuine retirement savings plan (Fidelity Investments Company, 2000).

The participants were informed that besides their show-up fee of $2, they would earn a percentage of the value of their portfolio at the end of the investment horizon as an additional payment (0.003% in the long-horizon condition and 0.012% in the short-horizon condition). In addition, the instructions explained how to use the experimental software by showing a screenshot of the program (see Figure 1). Participants indicated their decisions by moving a tab up and down a vertical slide rule. For the first decision, the tabs were positioned at the bottom of the rule, indicating an investment of 0% in the four funds, such that participants had to move the tab to make their allocation. In all subsequent years, the tabs were positioned at the previous allocations, so that if participants did not want to change their allocations, they could simply press the *Accept* button to move on. This method was intentionally used to avoid triggering any change in allocations by participants; thus, the default was to make no changes at all. On the same page as the slide rule, a pie chart showed participants' allocations in graphic form, and a history window gave the percentage of their portfolio invested in each of the four funds each year and the return rate of each fund in percentages. A table was also presented showing their gains each year and the total value of their portfolio.

In the risk assessment task, the second part of the experiment, participants made 12 choices between two gambles. The first of these 12 pairs of gambles was always the same. This one, *Gamble A*, led with a probability of 1.0 to a payoff of $10; thus, it represented a sure outcome. The second one, *Gamble B*, always led to one of two possible outcomes: $4 or $20. Across the 12 pairs, the only variation was in the probability of occurrence of these two possible payoffs. The first choice was between Gamble A and Gamble B, for which the probability of obtaining $20 was .66 and the probability of obtaining $4 was .34. The second choice was between Gamble A again and Gamble B, this time with the probability of obtaining the maximum payoff of $20 reduced by 4 percentage points to .62. For the remaining 10 choices, the probability of obtaining the maximum payoff for Gamble B was decreased further. Table 1 provides all 12 pairs of gambles. The expected value of the risky gamble decreased across the 12 gambles, starting with an expected value of $14.56 and ending with an expected value of $5.60. A persons' risk attitude could be measured by the choices between the gambles: An extremely risk-averse person will always choose the safe gamble, and an extremely risk-seeking person will always choose the risky gamble. A risk-neutral person will always choose the gamble with the larger expected payoff, choosing the risky gamble for the first 8 pairs of gambles and choosing the safe gamble for the remaining 4 gambles. The risk attitude measurement follows the logic of a Guttman scale (Guttman, 1950): Because of the increasing expected value of Gamble B, choosing Gamble B at a later stage (see Table 2) implies that a consistent decision maker had chosen Gamble B for all preceding pairs of gambles. Thus, to measure a person's risk attitude, it is theoretically sufficient to identify the pair of gambles at which point the decision maker switched from the choice of the risky Gamble B to the safe Gamble A. However, to take the probabilistic nature of people's preferences into account (see Rieskamp, Busemeyer, & Mellers, 2006), the risk attitude was measured by simply counting how often a participant had chosen the risky Gamble B, so lower values represent greater risk aversion. This gamble paradigm for measuring people's risk attitudes has been used in similar ways in past research (e.g., Donkers, Melenberg, & Van Soest, 2001; Holt & Laury, 2002; Lusk & Coble, 2005; Murnighan, Roth, & Schoumaker, 1988; Pennings & Smidts, 2000). For instance, Lusk and Coble were able to show that risk attitudes measured with the gamble paradigm are correlated with consumers' preferences for genetically modified food (i.e., the likelihood to reject genetically modified food increased by 15% for average risk-averse as opposed to risk-neutral individuals). Pennings and Smidts (2000) showed that risk attitudes measured with the gamble paradigm were positively correlated with managers' financial market behavior (e.g., manager's risk attitudes explained 4% of the variance of using futures to cover risk). Along these lines, Donkers et al. (2001) recommended using the gamble paradigm to measure risk attitudes in predicting asset allocation decisions. Because risk

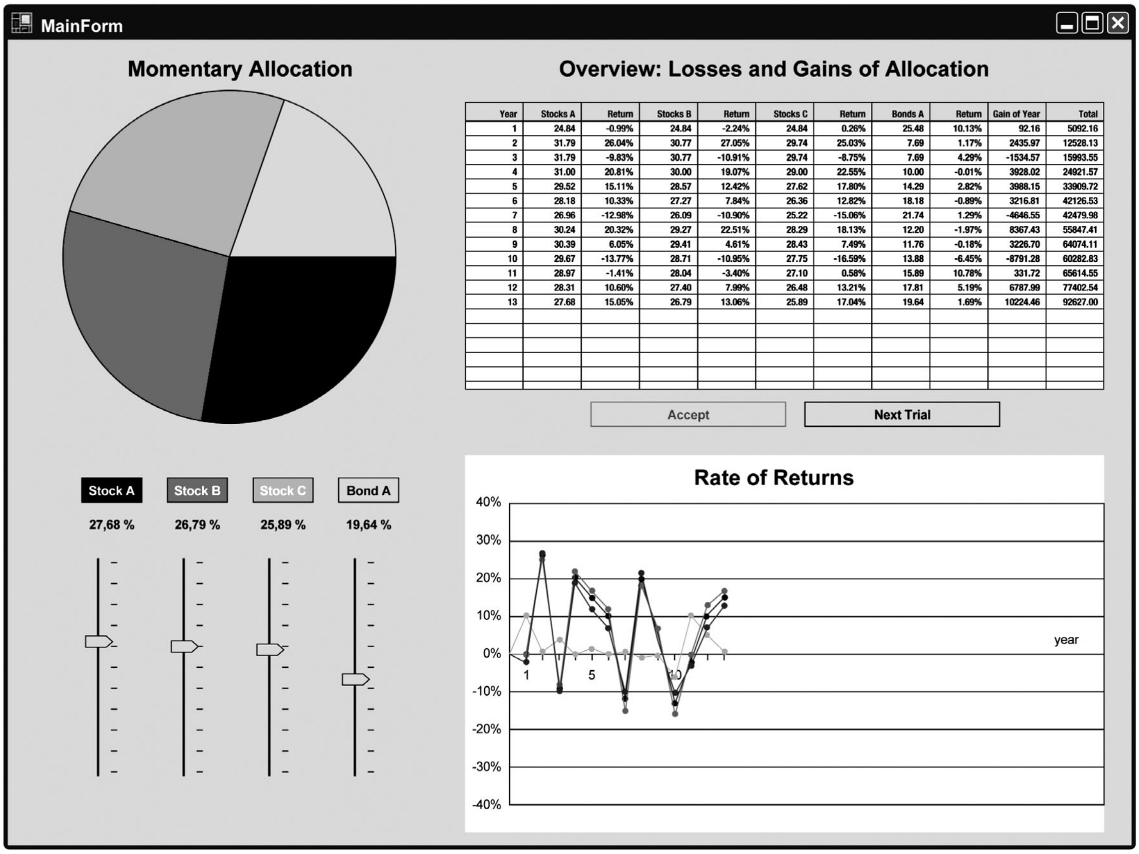

Figure 1. Screenshot of the experimental software used in Experiment 2. The software used in Experiment 1 was very similar, the only difference being that the graph showing the return on investment of the different funds in the lower right corner of the screen was missing. Participants made their decisions by moving a tab up and down a slide rule below each stock or bond fund. The pie chart reflected their allocations.

attitudes can differ across domains (Weber, Blais, & Betz, 2002), when focusing on the financial domain, use of the gamble paradigm appears to be a valid risk measurement technique.

The instructions for the risk assessment task explained to the participants that they would make 12 choices between pairs of gambles. Their choices would express their subjective preferences, and there would be no right or wrong choice. Participants were told that after the experiment, 1 of their decisions would be chosen at random and, depending on which gamble they had chosen, they would receive 10% of the particular outcome of the gamble as payment. Because of the primary focus of the present experiment on the learning processes in the retirement savings task, the risk assessment task was performed after the retirement savings task to exclude any potential effect of the risk assessment task on the behavior in the retirement saving task. However, to avoid carryover effects, participants were given no information about their performance in the retirement savings task before the risk assessment task. Only after the completion of both tasks were participants paid: First, they received payment according to their performance in the retirement savings task; second, 1 decision from the risk assessment task was selected randomly, the gamble chosen was played out, and the corresponding payment was made.

Results

The results were analyzed in two steps. First, participants' allocations, potential learning effects, and the composition of the various funds were analyzed. Further, the potential relationship between participants' risk attitudes and their allocations were examined. Second, the $1/n$ strategy and the two learning models were applied to describe participants' allocations.

Learning and composition effects. Figure 2 shows the average investment in stocks for all periods, differentiated with respect to the composition of the funds offered and the time horizon. It becomes clear that the investments in stocks were substantially higher in the three-stock fund condition compared with the two-stock fund condition. Participants' allocations express their subjective preferences, so there is no outside criterion by which to judge whether the allocations of an individual were good or bad. Nevertheless, to put the performance of the participants into perspective, their average return rates can be compared with the average return rates of 6.2% (5.9%) for the stock funds for the 12-year (35-year) investment period and the average return rate of 2.8% (2.8%) for the bond funds for the 12-year (35-year) investment period. Participants in the three-stock fund condition had an average return rate of 5.5% (SD = 4.6%) for the 12-year period and of 5.3% (SD = 0.5%) for the 35-year period. Participants in the two-stock fund condition had an average return rate of 3.8% (SD = 1.0%) for the 12-year period and of 4.5% (SD = 0.7%) for

Table 1
Risk Assessment Task

Pair	Gamble[a] A	Gamble[a] B	Participants choosing Gamble B (%) Experiment 1	Participants choosing Gamble B (%) Experiment 2
1	{$10, 1.0}	{$20, .66; $4, .34}	89	98
2	{$10, 1.0}	{$20, .62; $4, .38}	79	93
3	{$10, 1.0}	{$20, .58; $4, .42}	48	60
4	{$10, 1.0}	{$20, .55; $4, .45}	45	48
5	{$10, 1.0}	{$20, .52; $4, .48}	40	35
6	{$10, 1.0}	{$20, .49; $4, .51}	20	28
7	{$10, 1.0}	{$20, .46; $4, .54}	21	15
8	{$10, 1.0}	{$20, .42; $4, .58}	11	5
9	{$10, 1.0}	{$20, .37; $4, .63}	18	5
10	{$10, 1.0}	{$20, .30; $4, .70}	11	0
11	{$10, 1.0}	{$20, .20; $4, .80}	8	0
12	{$10, 1.0}	{$20, .10; $4, .90}	8	0

Note. In Experiment 1, U.S. dollars were used, whereas in Experiment 2, euros were used as the currency (at the time of the experiment, €1 = ~$1).
[a] The monetary values represent the payoffs that could occur; the value following each represents the probability with which that payoff occurs.

the 35-year period. Thus, on average, participants' performance was in between the performance of an investment purely in stocks and one purely in bonds, and participants with three stock funds performed better than participants with two stock funds, $t(78) = 2.41$, $p < .05$, $d = 0.54$, representing a medium effect according to Cohen (1988). For this and all following statistical tests, an alpha level of .05 was chosen to determine statistical significance.

A repeated measures analysis of variance (ANOVA) was conducted, with the percentage of resources allocated to stocks as the dependent variable (limited to the first 12 periods), the fund composition and the investment horizon as two between-subjects factors, and the periods as a within-subject factor. The results are summarized in Table 2. Participants' allocations were strongly influenced by the composition of the funds: On average, participants invested 75% ($SD = 4\%$) in stock funds in the three-stocks condition compared with 51% ($SD = 5\%$) in the two-stocks condition. In contrast, for the two investment horizon conditions, no significant differences in participants' investments were observed, such that on average, participants invested 62% in stocks with a 12-year horizon compared with 64% in stocks with a 35-year horizon. No significant interaction between the composition of the funds and the investment horizon was identified.

Participants changed their allocations across the 12 periods, with a maximum average investment in stocks of 71% ($SD = 20\%$) in the first period and a minimum average investment in stocks of 56% ($SD = 23\%$) in the fifth period. However, no clear trend in the investments could be observed. An interaction between the investment horizon and the time periods was observed, such that in the first four periods, participants with the 12-year horizon invested more in stocks compared with participants with the 35-year horizon. In contrast, in the remaining periods (with the exception of the last two periods), participants with the 12-year horizon invested a lower percentage in stocks compared with participants with the 35-year horizon. No other significant interactions were obtained.

Table 1 shows the results of the risk assessment task. Consistent with the logic of the task, the majority of participants had chosen the risky Gamble B for the first pair, whereas a much smaller number had chosen the risky Gamble B for the last pair. With the exception of Pairs 7 and 9, the choice percentages of the risky Gamble B decreased monotonically across the 12 pairs. To assess the reliability of the task for each participant, the number of inconsistent decisions was determined. For instance, when a participant had chosen Gamble B for Pair 2 but Gamble A for Pair 1 (for a list of the pairs and further details, see Table 1), the choice of Gamble A was considered an inconsistent decision. Thus, when choosing Gamble B for any pair, a consistent decision maker had to have chosen it as well for all preceding pairs. Likewise, when choosing Gamble A for any pair, a consistent decision maker had to choose Gamble A for all succeeding pairs. Following these two rules, the minimum number of inconsistent decisions was determined for each participant. On average, participants were very consistent, with an average number of 0.7 ($SD = 1.3$) inconsistent decisions, and for 69% of all participants, no inconsistent decisions were observed. Apparently, the participants understood the task very well, representing a reliable risk measurement.

To test the risk attitude hypothesis, the partial correlation between participants' risk attitudes and the average percentage invested in stocks across the first 12 periods was determined by statistically controlling for the two between-subjects factors, fund composition and investment horizon. The partial correlation was .31 ($p < .05$, $df = 76$, corresponding to a medium effect size, f, of .32 according to Cohen, 1988), supporting the risk attitude hypothesis.

The average investments in stocks shown in Figure 2 indicate that most participants could have applied the $1/n$ strategy for making their investments. Therefore, to test whether the $1/n$ strategy predicts the investments in stocks on an individual level, for each participant it was determined whether his or her investment corresponds with the 75% (50%) investment in stocks predicted by the $1/n$ strategy for the three-stock (two-stock) fund condition, with a tolerable deviation of plus or minus 5%. For the first

Table 2
Analysis of Variance: Experiment 1

Source	*df*	*MS*	*F*	*f*[a]
Between subjects				
Fund composition (F)	1	132,728	50.35*	0.26
Investment horizon (H)	1	621	.24	0.06
F × H	1	451	.17	
S within-group error	76	2,636		
Within subject				
Years (Y)	6.3[b]	2,393	6.28*	0.28
Y × F	6.3[b]	485	1.27	
Y × H	6.3[b]	872	2.29*	0.32
Y × F × F	6.3[b]	672	1.76	
Y × S within-group error	478[b]	381		

Note. S = Subjects.
[a] Cohen's effect size measure (see Cohen, 1988), with 0.10, 0.25, and 0.40 representing a small, a medium, and a large effect size, respectively.
[b] Adjusted degrees of freedom according to the Greenhouse–Geisser correction accounting for violation of sphericity.
* $p < .05$.

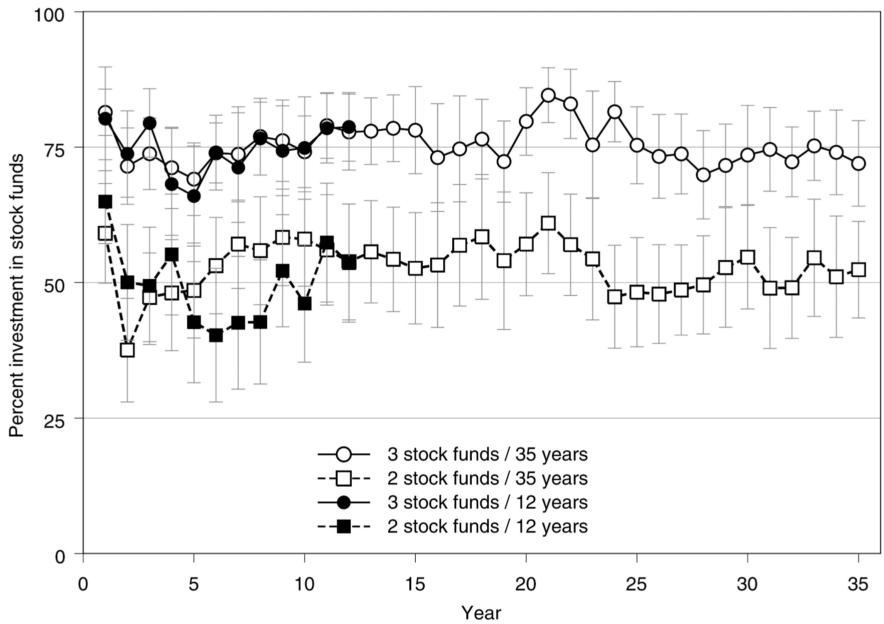

Figure 2. Average investment in stock funds in Experiment 1 across all periods, differentiated with respect to the composition of the funds offered and the time horizon. Error bars represent 95% confidence intervals.

investment period, 38% of the participants in the three-stock fund condition invested around 75% of their resources in stocks, whereas in the two-stock fund condition, 13% of the participants invested around 50% of their resources in stocks. Participants' investments corresponding to the investment predicted by the $1/n$ strategy were 24% (11%) across all periods for the three-stock (two-stock) fund condition. Thus, although the average investments in stocks corresponded with the $1/n$ strategy's prediction, the investments of the majority of participants deviated from the predicted investments.

Nevertheless, the composition of the funds offered had a strong influence on participants' investment decisions, and this effect remained constant across the whole investment period (see also Figure 2). However, participants did change their investments across the investment period, which can be illustrated by the frequency with which they changed their allocations. After their first allocation, on average, participants changed their allocations in 89% (SD = 14%) of all subsequent periods in the long-horizon condition and in 86% (SD = 10%) of all subsequent periods in the short-horizon condition. The variability of participants' investments in stocks across the 35-year and 12-year periods, respectively, was, on average, SD = 13% (12%). The average absolute change of the percentage invested in stocks for two subsequent periods was 9% for the long-horizon investment condition and 12% for the short-horizon condition. Thus, participants changed their investments substantially, although no clear trend across the whole investment period was observed.

To examine whether the changes of investments in stocks depended on the returns on the stocks in the previous periods, the correlation between the change of the investment in stocks in the current period—that is, $\Delta PR_t(\text{Stocks}) = PR_t(\text{Stocks}) - PR_{t-1}(\text{Stocks})$—and the change of the average rate of returns for the stock funds in the previous period—that is, $\Delta r_{t-1}(\text{Stocks}) = r_{t-1}(\text{Stocks}) - r_{t-2}(\text{Stocks})$—was determined. The average correlation was .11 (SD = .42; here and for all following average correlations, the correlations were first Fisher Z-transformed for every participant and afterward retransformed). The correlation varied between −.89 to .81, indicating large individual differences. For 45% (55%) of the participants, negative (positive) correlations were observed. A negative correlation indicates a negative recency effect: Participants responded to an increase in stock returns with a decreased investment. In contrast, a positive correlation indicates a positive recency effect: Participants responded to an increase in stock returns with an increased investment. The substantial positive and negative correlations indicate individual heterogeneity in responding to stock fund returns. In addition, the correlation between the change of the investment in stocks in the current period and the change of the rate of returns for the stock funds two periods before the current period—that is, $\Delta r_{t-2}(\text{Stocks}) = r_{t-2}(\text{Stocks}) - r_{t-3}(\text{Stocks})$—was determined. The average correlation was −.02 (SD = .34), ranging from −.73 and .73. In sum, although the average correlations indicate that the past performance of stocks did not have a strong influence on investments, the large individual differences show that the participants were apparently influenced differently, which is also assumed by the LA model. Can participants' investment changes be explained by the two learning models?

Comparison of the learning models. First, the two learning models were tested to determine whether they outperformed the $1/n$ strategy in describing participants' investments. Second, the two learning models were compared against each other. Because both learning models have free parameters, these were fitted for each individual. First, a set of parameter values was selected for a model for each individual. Using the model and parameters, a prediction was generated for each new trial, conditioned on the funds' past returns before that trial. The models predict one single allocation, which could be compared with the participants' allocations by computing the euclidean distance between the predicted and the observed allocations. To assess the overall fit of a model for a given individual, the deviation was squared for each trial, and the average was taken across all trials; that is, the average *squared error* (*SE*) was taken as a goodness-of-fit criterion (*SE* had a possible range of 0–20,000). It is important to stress that although the models were fitted to each participant, participants' investments were not used to determine any of the models' predictions. To optimize the parameters for each participant and model, reasonable parameter values were first selected by a grid-search technique, and thereafter the best-fitting grid values were used as a starting point for a subsequent optimization using the Nelder–Mead simplex method (Nelder & Mead, 1965). For the optimization process, the parameter values for the RL model were restricted to an initial strength parameter (w) between 0 and 10, a forgetting rate (ϕ) between 0 and 1, and a sensitivity parameter (θ) between 0 and 1. The parameter values for the LA model were restricted to an initial step size (s_1) between 1 and 141, and the investment strategy parameter (c) could take the value -1 or 1. The above procedure was applied to each of the 80 participants to obtain 80 sets of parameter estimates.

For the RL model, this produced the following means and standard deviations for the three parameters: average initial strength of $w = 0.32$ ($SD = 1.36$), average forgetting rate of $\phi = 0.05$ ($SD = 0.16$), and average sensitivity parameter of $\theta = 0.01$ ($SD = 0.02$). The RL model reached an average fit of $SE = 728$ ($SD = 773$). For the LA model, the estimation procedure led to an average initial step size of $s_1 = 7.0$ ($SD = 8.6$) and to a cyclical investment strategy for 58.75% of the participants, so that for these participants, a positive recency effect was predicted. For the remaining participants, there followed a countercyclical investment strategy, predicting a negative recency effect. The LA model reached an average fit of $SE = 603$ ($SD = 628$). Finally the 1/*n* strategy, without any free parameters, obtained an average fit of $SE = 848$ ($SD = 876$).

The models were compared with each other by comparing their fit for each individual. First, the RL model was compared with the 1/*n* strategy. The RL model was better at describing participants' investments for 66% of all participants ($Z = 4.94$, $p < .05$; according to a Wilcoxon signed ranks test). Likewise, the LA model was better at describing participants' investments for 89% of the participants compared with the 1/*n* strategy ($Z = 7.13$, $p < .05$). Thus, both learning models outperformed the 1/*n* strategy at predicting the investments by changing allocations because of the funds' performances. When comparing the two learning models, LA outperformed RL for 78% of all participants ($Z = 4.85$, $p < .05$). Thus, overall, the LA model was best at describing individuals' investments.

Figure 3 shows the average investments in the stock funds of the participants in the 35-year investment period. Additionally, it shows the average investment predicted by the two learning models when fitted to each participant. Both models described the investment process adequately. However, for the last third of the investment period, the RL model predicted a slight upward trend in the investments, which was not observed. In contrast, and consistent with the LA model, the investment in stock funds on average remained constant. However, although the average investments in stocks predicted by both models did not change substantially across the investment period, the predictions did differ when the prediction for each individual was examined.

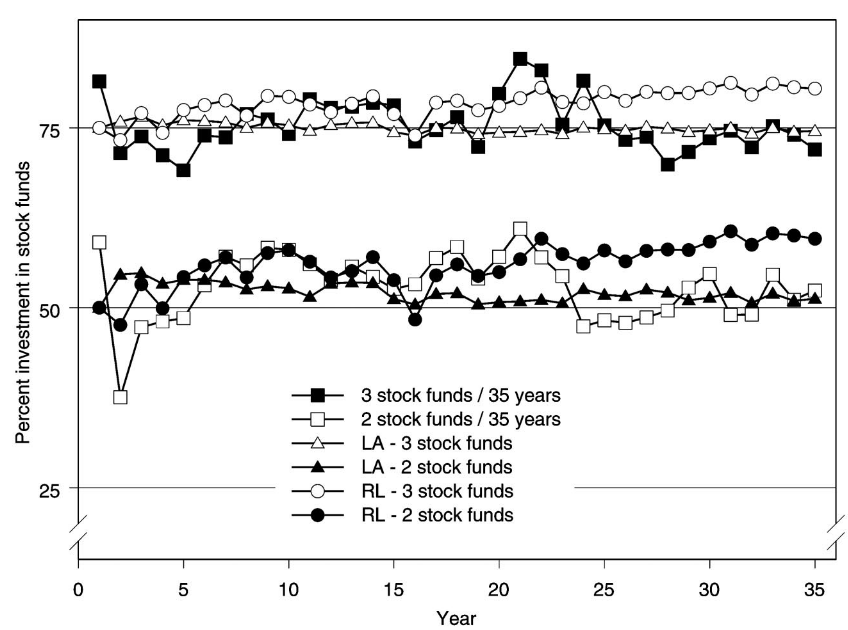

Figure 3. Participants' average investments in the stock funds as compared with the average predicted investment of the local adaptation learning (LA) and reinforcement learning (RL) models in Experiment 1 for the long-horizon (35-year) condition.

Summary of Experiment 1

Experiment 1 showed that people were strongly influenced by the composition of the investment funds offered. In the case in which the majority of funds offered were stock funds, participants invested a large percentage of their resources in stocks, whereas the investment in stocks was much lower when a smaller number of stock funds were offered. Surprisingly, the composition effect remained constant across the whole investment period. People's investments were also influenced by their risk attitudes, such that risk-seeking individuals made higher investments in stocks. No evidence was found that participants made different investments when they had a shorter investment horizon than when they had a longer investment horizon.

Although no evidence was found that the average percentage of participants' resources invested in stock funds increased or decreased over the whole investment period, participants still changed their investments significantly across time. Therefore, the two learning models proved better able to describe participants' investments than did the 1/*n* strategy. In general, one would expect that the RL model would predict an increasing percentage of resources invested in stocks across the investment horizon, given the superior performance of the stock funds. Moreover, this increase should be larger for the two-stock fund condition than for the three-stock fund condition. In fact, the RL model (when fitted to participants' investments) predicted a slight increase of about 5% of resources being invested in stock funds for the three-stock fund condition and an increase of about 10% being invested in stock funds for the two-stock fund condition. This moderate increase is attributable to the RL model's flexibility. Because the participants, on average, did not increase their investments in stock funds, parameter values for the RL model were obtained that could account for this result. The small value obtained for the sensitivity parameter means that the model was fairly insensitive to the funds' performance differences. Contrary to the RL model, the predictions of the LA model depended neither on the funds' absolute performance nor on their absolute cumulative performance. Rather, they depended on the recent changes in the funds' returns leading to positive or negative recency effects. If a fund's return increased compared with the previous period, the percentage of resources invested in it is increased when a cyclical investment strategy is used, leading to a positive recency effect, or it is decreased when a countercyclical investment strategy is used, leading to a negative recency effect. Thus, the model assumes that people's investments are strongly influenced by the recent changes in funds' returns. Contrary to the RL model, which can only predict positive recency effects, the LA can also predict negative recency effects, which were observed for a substantial percentage of the participants. Thereby, LA outperformed RL and the 1/*n* strategy in predicting participants' investments.

However, can this advantage be generalized to an independent situation for which the models' parameters are not fitted? This question was addressed in Experiment 2. The model comparisons for Experiment 1 neglected the models' complexities. For instance, the 1/*n* strategy has no free parameters, so it has no flexibility in describing changes in investments. Likewise, the two learning models have different numbers of free parameters, indicating different flexibilities. However, when judging the models' complexity on the number of free parameters, the simple LA model with only two free parameters achieved a better fit than did the more complex RL model with three free parameters, providing additional support for LA. Nevertheless, because the model's complexity was not explicitly taken into account, Experiment 2 provides the crucial generalization test of the models when making predictions for a new situation without estimating parameters.

Experiment 2

The surprising result of Experiment 1 was the strong effect of the composition of the funds offered, which was not only observed in the first investment period but also remained constant across the entire investment period, over which participants were provided with extensive learning opportunities. However, in Experiment 1, participants could see their funds' performance in a detailed table on the computer screen. This table also presented information on participants' previous investments in each of the four funds, their overall gain in the current period, and the total value of their portfolio. Thus, because of the large amount of information, participants might have found it difficult to detect the performance advantage of the stock funds. To rule out this explanation for the small overall changes in participants' investments, additional feedback about the funds' performance was provided graphically in Experiment 2, as shown in Figure 1. Here, the annual rate of return for each fund was presented on a graph with a colored line that had markers representing the returns for each fund. The graph was designed to make it easier to detect the performance advantages of the stock funds and, potentially, change the allocations over time.

The second and equally important goal of Experiment 2 was to provide a generalization test of the two learning models. For this purpose, the two learning models' predictions on the basis of the parameter estimates of Experiment 1 and of the 1/*n* strategy were compared with participants' average investments in Experiment 2. Because the models' parameter values are not optimized on the basis of the participants' behavior, Experiment 2 provided a strong empirical generalization test of the models, which has often been called for but seldom accomplished (Busemeyer & Wang, 2000). Although the experimental task was very similar to the task of Experiment 1, the experiment was conducted in a different country—Germany—whereas Experiment 1 was conducted in the United States. People's savings behavior differs substantially between these two countries, with Americans investing much higher percentages of their savings in the stock market than do Germans (e.g., Guiso, Haliassos, Jappelli, & Claessens, 2003), indicating a higher risk aversion for Germans. Because of these differences in savings behavior, Experiment 2 represented a crucial generalization test of the two learning theories.

Method

Participants. Forty persons (20 women and 20 men) with an average age of 25 years participated in the experiment. The duration of the computerized task was approximately 1 hr. Most participants (80%) were students in various departments at the Free University of Berlin, Germany. For their participation, they received a show-up payment of €2 (~U.S.$2). All additional payment depended on the participants' performance; they received 0.002% of their final portfolio value. This resulted in an average payment of €11. In addition, participants received payment for the risk assessment task.

Procedure. In the first part of the experiment, participants performed the retirement savings decision task, and in the second, their risk attitudes were assessed. Participants were randomly assigned to one of two between-subjects conditions. The conduct of the experiment was virtually identical to that of Experiment 1, with only two changes: First, all participants had an investment horizon of 35 years. Second, during the experiment, participants received additional feedback in graphs about the annual return rates of the funds (see Figure 1). The funds' returns were identical to the returns used in Experiment 1. The instructions were virtually identical to those provided in Experiment 1, the only difference being that the graphs on the computer screen were explained and euros were used as the currency instead of U.S. dollars. Again, participants were informed that besides their show-up fee, they would earn 0.002% of the value of their portfolio at the end of the investment period as an additional payment.

In the second part of the experiment, participants performed the risk assessment task. The task was virtually identical with that used in Experiment 1, the only difference being that euros were used as the currency. Again, a participant's risk attitude was measured by the frequency with which the risky gamble was preferred over the sure payoff. At the end of the experiment, after the risk assessment task, participants were paid according to their performance in the retirement savings task. Afterward, 1 pair of gambles from the risk assessment task was drawn randomly, the chosen gamble was played out, and 10% of the payoff obtained was paid to the participant.

Results

The results were analyzed in two steps. First, participants' allocations and potential learning and composition effects were analyzed, considering the relationship between risk attitudes and allocations. Second, the $1/n$ strategy and the two learning models were applied to describe participants' allocations.

Learning and composition effects. Figure 4 shows the average investment in stocks for all periods, differentiated with respect to the composition of the funds offered. It became clear that the investments in stocks were substantially higher in the three-stock fund condition than in the two-stock fund condition, replicating the findings of Experiment 1. Similar to Experiment 1, participants' average return rates were determined and compared with the average return rates on stocks (bonds) of 6.2% (2.8%) across the 35-year investment period. Participants with a three-stock fund had an average return rate of 4.9% (SD = 0.7%) compared with participants with a two-stock fund, who had an average return rate of 4.2% (SD = 0.5%). Thus, participants' performance was, on average, in between that of an investment purely in stocks or purely in bonds, and participants with three-stock funds performed better than those with two, $t(38) = 3.45$, $p < .05$, $d = 1.09$, representing a strong effect, as described by Cohen (1988).

A repeated measures ANOVA was conducted, with the percentage of the resources invested in the stock funds as the dependent variable, the fund's composition as a between-subjects factor, and the periods as a within-subject factor. Table 3 provides an overview of the results. Participants' allocations were strongly influenced by the funds' composition: Participants, on average, invested 65% (SD = 15%) in stocks in the three-stock fund condition compared with 45% (SD = 10%) in the two-stock fund condition. The average investments across the investment period differed across the 35 investment periods. No significant trend in the investments could be observed. An interaction between the funds' composition and the investment periods was not observed. Comparing Experiment 2 with Experiment 1, the average investment in stocks of 55% in Experiment 2 was lower than that of 64% in the American sample in Experiment 1, indicating higher risk aversion in the German sample.

Table 1 shows the result of the risk assessment task. Consistent with the logic of the risk assessment measure and similar to Experiment 1, the majority of participants had chosen the risky

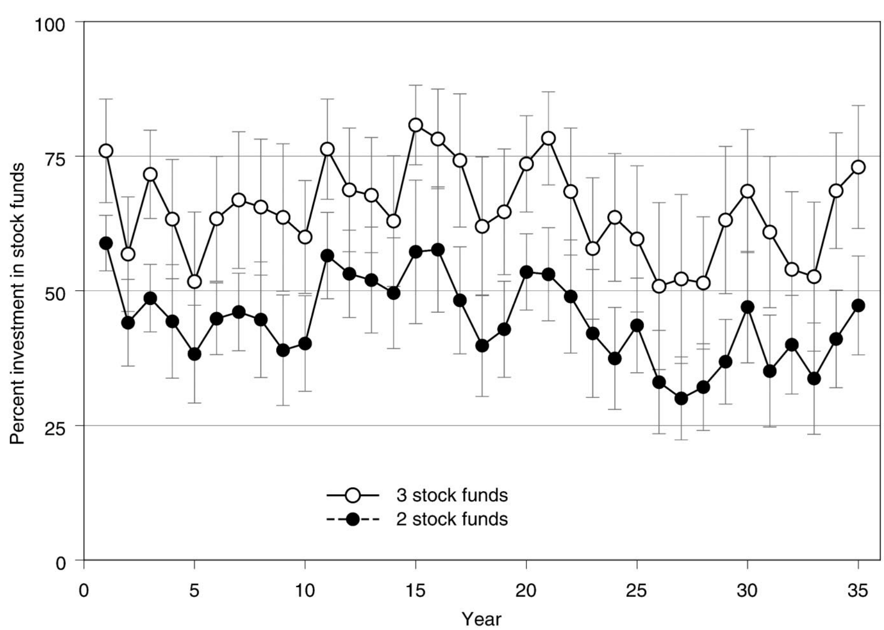

Figure 4. Average investments in stock funds in Experiment 2 across all periods, differentiated with respect to the composition of the funds offered. Error bars represent 95% confidence intervals.

Table 3
Analysis of Variance: Experiment 2

Source	*df*	*MS*	*F*	f^a
		Between subjects		
Fund composition (F)	1	4,196,831	711*	0.78
S within-group error	38	5,905		
		Within subject		
Years (Y)	8.6[b]	9,561	5.52*	0.80
Y × F	8.6[b]	679	0.39	
Y × S within-group error	328[b]	1,732		

Note. S = Subjects.
[a] Cohen's effect size measure (see Cohen, 1988), with 0.10, 0.25, and 0.40 representing a small, a medium, and a large effect size, respectively.
[b] Adjusted degrees of freedom according to the Greenhouse–Geisser correction accounting for violation of sphericity.
* $p < .05$.

Gamble B for the 1st pair of gambles. No participant had chosen the risky Gamble B for the last 3 pairs of gambles, indicating again that the German sample was more averse to risk. The choice percentages of the risky gamble monotonically decreased across the 12 pairs, with the exception of Pair 9, for which the percentage remained stable. As in Experiment 1, the reliability of the risk assessment measurement was assessed by determining the number of inconsistent decisions for each participant. On average, Experiment 2 participants were very consistent, even more than those in Experiment 1, with an average number of only 0.1 (*SD* = 0.4) inconsistent decisions and no inconsistent decisions observed for 93% of all participants. Thus, the task was apparently well understood by the participants, suggesting a reliable risk measurement.

The correlation between participants' risk attitudes and the average percentage invested in stocks across the 35 periods was determined by controlling for the stock composition factor; the partial correlation was −.27 (*ns*, *df* = 37). Thus, these results do not replicate the positive correlation between risk attitude and the investment in stocks observed in Experiment 1.

Participants' investments were strongly influenced by the composition of the funds offered. Could the average investments in stocks be a result of the application of the 1/*n* strategy? The average results shown in Figure 4 already indicate that many participants apparently did not select the 1/*n* strategy because the average investments were lower than the 75% (50%) predictions of the 1/*n* strategy in the three-stock (two-stock) fund condition. However, to test whether the 1/*n* strategy predicts the investments in stocks on an individual level, it was determined whether each participant's investment corresponded with the 75% (50%) investment in stocks predicted by the 1/*n* strategy for the three-stock (two-stock) fund condition, with a tolerable deviation of plus or minus 5%. For the first investment period, 28% of the participants in the three-stock fund condition invested around 75% of their resources in stocks, whereas in the two-stock fund condition, 18% of the participants invested around 50% of their resources in stocks. The average percentage of participants whose investments corresponded to the investment predicted by the 1/*n* strategy was 22% (11%) across all periods for the three-stock (two-stock) fund condition. Thus, the investments of the majority of participants, as in Experiment 1, deviated from the investments predicted by the 1/*n* strategy.

Nevertheless, the composition of the funds did affect participants' investments. This effect was observed across the whole investment period (see also Figure 4). Thus, there is no evidence that the graphic feedback about the funds' performance changed participants' investments. Although, on average, the investments in the stock funds did not differ across the investment periods, looking at the individual investments, participants changed their investments substantially on a period-to-period basis. After their first allocation, participants changed their allocations on average in 84% (*SD* = 16%) of all of the 34 subsequent periods in the three-stock fund condition and in 80% (*SD* = 23%) of all of the 34 subsequent periods in the two-stock fund condition. The variability of participants' investments in stocks across the 35 periods was on average *SD* = 19%. Likewise, the absolute change in the percentage participants invested in stocks for the 2 subsequent periods was, on average, 15%.

To examine whether the changes in investments in stocks depended on the stocks' performance in the previous periods, the correlation between the change of the stocks' average rates of return in the previous period and the change in the investments in stocks in the current period was determined. The average correlation was .02 (*SD* = .44), ranging from −.78 to .63, indicating large individual differences. For 55% (44%) of the participants, negative (positive) correlations were observed. Thus, for roughly half of the participants, the positive correlation indicated a positive recency effect, and for the other half, the negative correlation indicated a negative recency effect. In addition, the correlation between the change of the investment in stocks in the current period and the change in the rate of return for the stock funds 2 periods before the current period was determined. The average correlation was −.04 (*SD* = .29), ranging from −.70 and .52. In sum, similar to Experiment 1, the average correlations do not indicate that the stocks' past performance affected participants' investments, but the large individual differences show that the participants were apparently influenced differently because of a negative and a positive recency effect. How can the changes in investments be explained?

Model comparison. To predict participants' investments, the predictions of the 1/*n* strategy and the two learning models were compared with participants' investments. The learning models' predictions were generated by simulating 10,000 agents on the basis of the parameter estimates of Experiment 1. For each agent, a set of parameter values was randomly drawn from a normal distribution with a mean and a standard deviation according to the distribution of the estimated parameters from Experiment 1. For the investment strategy parameter of the LA model, a cyclical investment strategy was assumed for 58.75% of the agents, and a countercyclical investment strategy was assumed for the remaining 41.25% of the agents, according to the percentages estimated from Experiment 1. The 10,000 investments were averaged for each condition and compared with the participants' average investments by determining the average squared euclidian distance between them. Thus, the average *SE* between the prediction and the observation was determined.

The 1/*n* strategy predicts a constant investment of 50% and 75% in stock funds in the two-stock fund condition and the three-stock fund condition, respectively. In fact, the participants on average invested only 45% or 65%, respectively, of their resources in stocks. Comparing the 1/*n* strategy's prediction across all periods with participants' investments, the 1/*n* strategy reached an average fit of $SE = 98$ for the two-stock fund condition and of $SE = 258$ for the three-stock fund condition. The RL model predicts an average investment of 61% and 82% in stock funds in the two-stock fund and three-stock fund conditions, respectively. Thus, on average, RL overestimates the percentage invested in stocks. Comparing the RL model's predictions across all periods with participants' investments, it reached a lower fit than did the 1/*n* strategy, with average fits of $SE = 440$ in the two-stock fund condition and $SE = 569$ in the three-stock fund condition. In comparison, the LA model predicts an average investment of 48% and 65% in the two-stock fund and three-stock fund conditions, respectively, thereby coming closest to the behavior observed. Comparing the LA model's predictions across all periods with participants' investments, it reached average fits of $SE = 94$ for the two-stock fund condition and $SE = 152$ for the three-stock fund condition. Thus, LA outperformed the 1/*n* strategy and RL in describing participants' investments. In sum, the LA model not only performs better at describing participants' investments when the model is fitted to the data as shown in Experiment 1, it is also better at predicting independent behavior.

Figure 5 shows the average investments observed in stock funds across the 35 periods and the average investment predicted by the two learning models, differentiated for the two experimental conditions. For better comparison, a moving average with a window size of 3 periods was used. Both learning models predict participants' investments only to a certain degree. However, when both models are compared, RL makes an unsatisfactory prediction, because contrary to participants' actual investments, it predicts an increasing percentage invested in stock funds for the first 10 periods, followed by a decreasing percentage, and finally (for the second half of the investment period) an increasing percentage invested in stocks; all of these predictions were inconsistent with the participants' actual investments. In contrast, the LA model predicts a decrease in the percentage invested in stocks across the investment period, consistent with the investments observed. However, the LA model cannot account for the investment fluctuations that occur during the investment period.

Summary of Experiment 2

Experiment 2 replicated the main findings of Experiment 1: People were strongly influenced by the composition of the investment funds offered. If the majority of the funds offered were stock funds, participants invested a large percentage of their resources in stocks, whereas the investment in stocks was much lower when a smaller number of stock funds were offered. This effect remained constant across the whole investment period. This is an important generalized finding, because the experiment was conducted under different conditions: First, participants received additional graphic feedback about the annual performance of the different funds. With this feedback, it should have been easier to detect performance advantages among them. Accordingly, participants, in particular in the two-stock fund condition, should have increased the percentage of their resources invested in stocks. However, this was not observed. Second, with Germans as the sample of participants in Experiment 2, the findings of Experiment 1 were further generalized. Germans are generally more risk averse in their savings behavior than are Americans (e.g., Guiso et al., 2003). In sum, the composition of the funds offered in the retirement savings situation had a strong effect on participants' allocations.

The second goal of Experiment 2 was to provide a generalization test for the two learning models and the 1/*n* strategy. The predictions of the two learning models were generated on the basis of the parameter estimates of Experiment 1; thus, participants' behavior in Experiment 2 was not used to fit any parameters. Because Experiment 2 was conducted under different conditions, this provided a rigorous model-comparison test. Here, in contrast to Experiment 1, the 1/*n* strategy outperformed the reinforcement learning model in predicting participants' investments. This was because of the RL model's prediction of an increased investment in stock funds, which did not gain empirical support. The local adaptation learning model was best at predicting participants' investments for both experimental conditions. This illustrates that participants changed their allocations frequently but only on the basis of the funds' recent performance changes. The correlation between the stock funds' performance change in the previous period and participants' changed investments in the current period indicated a positive and a negative recency effect, as observed in Experiment 1.

General Discussion

Retirement savings decisions are among the most important financial decisions people make. Furthermore, these decisions can have dire consequences given the long period over which retirement savings are made and taking into account compound interest. However, despite their importance, retirement savings decisions have received relatively little attention in the psychological research. Most of the research on such decisions has been conducted by economists, who focus less on the cognitive processes behind decisions and more on the behavioral deviations from standard normative models. However, an understanding of the cognitive processes that lead to a decision can enable better predictions of behavior. The first goal of this study was to examine the extent to which people are influenced by the composition of the funds offered in retirement savings plans when they have the chance to change their allocations repeatedly. The second goal was to test two learning theories predicting repeated retirement savings decisions.

In both experiments, participants repeatedly made retirement savings decisions, allocating a given amount of resources to different funds. The task was difficult because the rates of return were unknown and the returns varied substantially. In both experiments, the composition of the funds offered had a strong effect on participants' allocations, such that a larger percentage of the financial resources were invested in stock funds when more stock funds were offered. Surprisingly, this effect remained stable across the whole investment period. Although participants' changed their allocations substantially from period to period across the entire investment period, no substantial trends in the investments were observed. Regardless of whether participants made their choices in the condition with a large or a small percentage of stock funds, they increased or decreased their investments in stocks to a similar extent, such that the difference between the two conditions did not change. In Experiment 1, no evidence was found that participants'

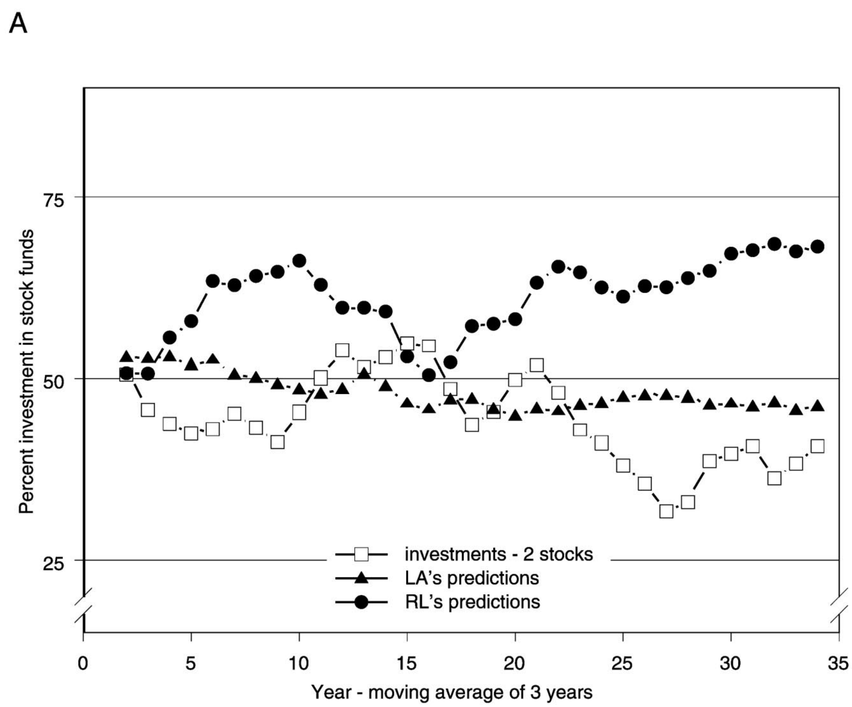

Figure 5. Average investments in the stock funds as compared with the average predicted investment of the local adaptation learning (LA) and reinforcement learning (RL) models for the two-stock fund (A) and three-stock fund (B) conditions in Experiment 2.

investments differed for the long-horizon (35-year) or short-horizon (12-year) investment horizon conditions. This is surprising given that people have a tendency to avoid losses (Kahneman & Tversky, 1979). Because a shorter investment horizon leads to a higher probability of obtaining a total loss when investing in stock funds, a smaller investment in stock funds could have been expected with a shorter investment horizon (for a discussion of this topic, see also Lopes, 1996; Samuelson, 1963). The results concerning the relationship between people's risk attitudes and their investment in stocks were mixed: Whereas in Experiment 1, a positive correlation between participants' risk aversion and their investments in bond funds was observed, no significant correlation was observed in Experiment 2. However, these results also have to be taken cautiously, because the risk assessment task was performed after the retirement investment task, so carryover effects cannot be ruled out.

The strong composition effect of the fund offered was consistent with people's tendency to allocate a resource equally among the available options, which has been found not only for retirement savings decisions (Benartzi & Thaler, 2001) but also for other resource allocation tasks (e.g., Fox et al., 2005; Langholtz et al., 1993, 1994, 1995). Hertwig, Davis, and Sulloway (2002) have even argued that people use this kind of "equality heuristic" when allocating resources to their children's development, which results in middle children receiving different amounts of resources than do firstborn and lastborn children. In principle, when a $1/n$ strategy is applied to allocation of a resource among different financial options, the strategy appears reasonable, assuming that a person's beliefs about the expected returns do not differ among the different options. Because most people do not have specialized knowledge or training in financial markets, this assumption is not implausible. However, the $1/n$ strategy is not sufficient to explain participants' changed allocations across the investment period. Therefore, two learning models were examined for predicting participants' investments.

The Learning Models

Recently, several learning theories have been proposed for decision-making problems (e.g., Börgers & Sarin, 1997; Busemeyer & Myung, 1992; Camerer & Ho, 1999a, 1999b; Erev & Roth, 1998; Rieskamp, 2006; Rieskamp et al., 2003; Rieskamp & Otto, 2006; Selten & Stöcker, 1986; Stahl, 1996). Most of them build on the basic idea that people do not solve a problem from scratch but adapt their behavior on the basis of experience (for a discussion on decisions based on experience, see also Hertwig, Barron, Weber, & Erev, 2004). The theories differ with respect to the learning mechanism assumed. The RL approach assumes that the decision alternatives can be assigned an overall evaluation. To make this evaluation, people integrate their experience, and alternatives that are evaluated positively are more likely to be selected. The LA approach assumes not that people make decisions on the basis of a global evaluation of the alternatives but, rather, that they adapt their decisions locally so that the success or failure of a preceding decision may lead to a slight modification of the subsequent decision (see also Busemeyer & Myung, 1987; Rieskamp et al., 2003; Selten & Stöcker, 1986). Because of this local adaptation process, when a cyclical or countercyclical investment strategy is used, the LA model is able to predict both positive and negative recency effects, whereas the RL model is only able to predict positive recency effects. To test the two models against each other in Experiment 1, the models' parameters were estimated separately for each participant, and the model fits were compared with the individual data. In Experiment 2, the estimated parameters from Experiment 1 were used to generate a priori predictions of the models tested.

Which learning theory was better at predicting participants' behavior? Both were better at describing participants' investments than was the $1/n$ strategy in Experiment 1. However, because the two learning models have free parameters and include the $1/n$ strategy as a special case, this finding is not surprising. Comparing the two learning theories, LA was better at describing the majority of participants' investments. For a substantial percentage of them, negative recency effects were observed, which can only be predicted by the LA model. However, because the models differ in complexity, Experiment 2 represents the critical model-comparison test: Here, LA was superior to RL at predicting participants' decisions.

How do the models explain participants' investments, and in particular, can they explain why the composition effect remained constant over time? In general, considering the mechanisms of the RL model, one would expect the model to predict an increasing percentage of resources being invested in stocks because of the better performance of the stock funds than of the bond funds. However, when the model was fitted to participants' investments in Experiment 1, it only described a moderate increase in the investment in stocks. This result is attributable to the fitting process: Because participants did not increase the percentage of resources invested in stock funds, the models' parameters were optimized in such a way as to account for this finding. By using, on average, small values for the sensitivity parameter (i.e., with an average of 0.01), the model could predict only a moderate increase in the investments in stocks. Nevertheless, even for this moderate increase, no empirical support was found. Thus, the RL model does not provide a good explanation as to why the composition effect remained constant over time.

The LA model provides a better account of participants' investments. It does so via two mechanisms: First, it assumes that people do not keep track of the whole history of the outcomes of past decisions, so it does not memorize the funds' past performances. Instead, it only compares the funds' recent performance with their past performance and with the maximum performance increase observed up to that point. Thus, in a situation in which the performance of investment options varies widely, this model is able to predict substantial changes in investments that do not have a clear trend. In particular, because the model makes predictions on the basis of performance *changes* instead of absolute performance, it is able to predict a decreasing percentage invested in an option even when that option, on average, performs much better than another. Considering the performance changes of the 35 investment periods used in Experiments 1 and 2 (starting from the 2nd period), in 17 periods (16 periods) the performance of the stock funds (bond funds) increased, and in the remaining 17 periods (18 periods) the performance decreased. Consistently, an individual who changed his or her investments on the basis of the performance changes would in approximately half of the periods increase and in the other half decrease the investment in the corresponding funds, so over a longer period, the average allocation would not necessarily change substantially. Therefore, this allocation mechanism can explain why the composition effect remained constant across the investment periods. The different investment strategies—cyclical or countercyclical—assumed by the LA model do not interfere with this explanation. An individual following a cyclical investment strategy will show a different reaction in the next period to the improved performance of an allocation option than will an individual following a countercyclical investment strategy. When the performance changes are counterbalanced, both strategies lead, on average, to similar percentages being allocated to the different investment options. However, the cyclical investment strategy leads to a positive recency effect, and the countercyclical investment strategy leads to a negative recency effect.

When the LA model was fit to the 80 participants in Experiment 1, for approximately 60% of the participants, a cyclical investment strategy was found; this percentage is larger than the 33% of participants engaging in a cyclical strategy found by Kroll et al. (1988a). However, their task was different, which might explain the different percentages. In sum, the principle of the LA model seems to correspond with the learning behavior observed among the individuals in both experiments. Investments were made on the basis of a fund's recent "success" or "failure." In an environment with large stochastic outcomes, such as a financial market, the principle predicts an increase and decrease in investments and no convergence to a particular allocation. This is exactly what was found in both experiments.

To what extent can the results of the current experiments be generalized to different learning models? The two models are instantiations of the two learning approaches. Both types have been applied to resource allocation tasks in previous research (see, e.g., Erev & Gopher, 1999; Rieskamp et al., 2003). The designs of the two models presented here were the best ones found for the retirement savings decision task. Moreover, given the flexibility of the models implemented—that is, their free parameters—slight modifications would not lead to substantially different results. Thus, it appears unlikely that modified versions of these two learning models would challenge the claim that LA is better than RL at predicting the specific learning processes involved in the situation of retirement savings. To what extent can the results of the current experiments be generalized to other decision problems? It should be emphasized that these conclusions are restricted to the decision problem considered here. It can be expected that in similar decision situations with large numbers of possible allocation options, the LA model would better describe the learning process. In such situations, it is hard to form overall evaluations for all possible allocations; instead, a directed learning process is much easier to apply. However, there are many situations for which RL models better describe learning processes. For instance, there is extensive empirical evidence that RL is appropriate to describe learning processes in situations with a small number of categorically different alternatives (Erev & Barron, 2005; Erev & Roth, 1998; Yechiam, & Busemeyer, 2005). In these situations, it is easy to memorize the accumulated experience for the small number of alternatives. Furthermore, it is difficult to apply an LA model when the set of alternatives does not provide a natural order to define directions for changes. In general, RL has successfully described learning in various domains, ranging from the learning of touch-typing (Yechiam, Erev, Yehene, & Gopher, 2003) to eye movement during reading (Reichle & Laurent, 2006).

Conclusions for Retirement Savings Decisions

One main finding of the current experiments is the strong composition effect of the funds offered on participants' retirement savings decisions. The effect remained constant even when there was a substantial learning opportunity. The general conclusions that can be drawn from this result are, however, subject to several limitations. First, participants in these experiments did not make real retirement savings decisions but had to imagine the retirement savings situation. Contrary to real retirement savings decisions, the monetary incentives were relatively small, and the experimental task lasted for only about 1 hr. However, because of these differences, the observed composition effects are even more surprising. The relatively small monetary incentives should have made participants more willing to change their initial allocations, thereby potentially overcoming the composition effect. In fact, participants did change their allocations significantly, but this did not weaken the composition effect. In addition, if participants' allocations do not change substantially when there is frequent feedback during a short time period, it is rather unlikely that substantial changes will occur in real retirement savings decisions in situations in which feedback on investments is likely to be infrequent. Thus, it appears important to take the role of composition effects for retirement savings decisions into account.

Nevertheless, the composition effect observed might be weaker for real retirement savings decisions when more allocation options are provided. As discussed above, empirical evidence for composition effects in real retirement saving plans is mixed (see Benartzi & Thaler, 2001; Huberman & Jiang; 2006). In Experiments 1 and 2, participants were offered a choice of only four funds. Although this is not an unrealistic number, many retirement savings plans offer their participants many more. In cases, of—for instance—20 funds, it appears less likely that people will allocate their resources in equal percentages among all the funds offered. Rather, people are likely to make a preselection of the funds (Huberman & Jiang, 2006) on the basis of their types and past performances. Such a preselection will, in turn, moderate composition effects. Finally, it is uncertain how improved learning opportunities could affect the composition effect. Participants in this study were not provided the learning opportunity to make repeated retirement savings decisions over a whole investment period. It would be interesting to study whether the opportunity to make retirement savings decisions repeatedly—for instance, 10 times for the whole 35-year investment period—would lead to stronger learning effects and, potentially, overcome composition effects. If this were the case, it could pave the way for training tools that could be used to advise individuals on making their retirement savings decisions.

This study demonstrates that retirement savings decisions represent a complicated decision problem. In such situations, context factors may have an influence that the individual would prefer to avoid. Through a rigorous test of two learning models representing two of the central approaches in the recent learning literature, an insight into individual decision processes has been provided. The LA approach helps to explain why learning did not lead to a substantial change in participants' allocations and why positive and negative recency effects occur. In other domains, another learning mechanism or learning mechanisms might govern behavior, and each learning model might have its own domain in which it works well. Identifying these domains will be a promising enterprise for future research.

References

Agnew, J., Balduzzi, P., & Sunder, A. (2003). Portfolio choice and trading in a large 401(k) plan. *American Economic Review, 93,* 193–215.

Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition, 32,* 1369–1378.

Ball, C. T., Langholtz, H. J., Auble, J., & Sopchak, B. (1998). Resource-allocation strategies: A verbal protocol analysis. *Organizational Behavior and Human Decision Processes, 76,* 70–88.

Benartzi, S., & Thaler, R. H. (2001). Naive diversification strategies in defined contribution saving plans. *American Economic Review, 91,* 79–98.

Börgers, T., & Sarin, R. (1997). Learning through reinforcement and replicator dynamics. *Journal of Economic Theory, 77,* 1–14.

Busemeyer, J. R., & Myung, I. J. (1987). Resource allocation decision-making in an uncertain environment. *Acta Psychologica, 66,* 1–19.

Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General, 121,* 177–194.

Busemeyer, J. R., Swenson, K., & Lazarte, A. (1986). An adaptive approach to resource allocation. *Organizational Behavior and Human Decision Processes, 38,* 318–341.

Busemeyer, J. R., & Wang, Y.-M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology, 44,* 171–189.

Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning.* New York: Wiley.

Camerer, C., & Ho, T.-H. (1999a). Experience-weighted attraction learning in games: Estimates from weak-link games. In D. V. Budescu & I. Erev (Eds.), *Games and human behavior: Essays in honor of Amnon Rapoport* (pp. 31–51). Mahwah, NJ: Erlbaum.

Camerer, C., & Ho, T.-H. (1999b). Experience-weighted attraction learning in normal form games. *Econometrica, 67,* 827–874.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences.* Hillsdale, NJ: Erlbaum.

Cutler, D. M., Poterba, J. M., & Summers, L. H. (1990). Speculative dynamics and the role of feedback traders. *American Economic Review, 80,* 63–68.

Donkers, B., Melenberg, B., & Van Soest, A. (2001). Estimating risk attitudes using lotteries: A large sample approach. *Journal of Risk and Uncertainty, 22,* 165–195.

Dorfman, D. D., Saslow, C. F., & Simpson, J. C. (1975). Learning models for a continuum of sensory states reexamined. *Journal of Mathematical Psychology, 12,* 178–211.

Edwards, W. (1961). Probability learning in 1000 trials. *Journal of Experimental Psychology, 62,* 385–394.

Erev, I. (1998). Signal detection by human observers: A cutoff reinforcement learning model of categorization decisions under uncertainty. *Psychological Review, 105,* 280–298.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112,* 912–931.

Erev, I., & Gopher, D. (1999). A cognitive game-theoretic analysis of attention strategies, ability, and incentives. In D. Gopher & A. Koriat (Eds.), *Attention and performance XVII: Cognitive regulation of performance: Interaction of theory and application* (pp. 343–371). Cambridge: MIT Press.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review, 88,* 848–881.

Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review, 57,* 94–107.

Estes, W. K. (1964). Probability learning. In A. W. Melton (Ed.), *Categories of human learning* (pp. 89–128). New York: Academic Press.

Fidelity Investments Company. (2000). *Choosing your retirement investments* [Brochure]. Boston, MA: Author.

Fox, C. R., Ratner, R. K., & Lieb, D. S. (2005). How subjective grouping of options influences choice and allocation: Diversification bias and the phenomenon of partition dependence. *Journal of Experimental Psychology: General, 134,* 538–551.

Gingrich, G., & Soli, S. D. (1984). Subjective evaluation and allocation of resources in routine decision making. *Organizational Behavior and Human Decision Processes, 33,* 187–203.

Guiso, L., Haliassos, M., Jappelli, T., & Claessens, S. (2003). Household stockholding in Europe: Where do we stand and where do we go? *Economic Policy, 18,* 125–170.

Guttman, L. (1950). The basis of scalogram analysis. In S. A. Stouffer et al. (Eds.), *Measurement and prediction: Studies in social psychology in World War II* (Vol. IV, pp. 60–90). Princeton, NJ: Princeton University Press.

Hake, H. W., & Hyman, R. (1953). Perception of the statistical structure of a random series of binary symbols. *Journal of Experimental Psychology, 45,* 64–74.

Harley, C. B. (1981). Learning the evolutionary stable strategy. *Journal of Theoretical Biology, 89,* 611–633.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15,* 534–539.

Hertwig, R., Davis, J. N., & Sulloway, F. J. (2002). Parental investment: How an equity motive can produce inequality. *Psychological Bulletin, 128,* 728–745.

Hirshleifer, D. (2001). Investor psychology and asset pricing. *Journal of Finance, 56,* 1533–1597.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review, 92,* 1644–1655.

Huberman, G., & Jiang, W. (2006). Offering versus choice in 401(k) plans: Equity exposure and number of funds. *Journal of Finance, 61,* 763–801.

Ibbotson Associates. (2003). *Stocks, bonds, bills, and inflation 2003 yearbook.* Chicago: Author.

Jarvik, M. E. (1951). Probability learning and a negative recency effect in the serial anticipation of alternative symbols. *Journal of Experimental Psychology, 41,* 291–297.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291.

Kroll, Y., Levy, H., Markowitz, H. M. (1984). Mean-variance versus direct utility maximization. *Journal of Finance, 39,* 47–61.

Kroll, Y., Levy, H., & Rapoport, A. (1988a). Experimental tests of the mean-variance model for portfolio selection. *Organizational Behavior and Human Decision Processes, 42,* 388–410.

Kroll, Y., Levy, H., & Rapoport, A. (1988b). Experiment tests of the separation theorem and the capital asset pricing model. *American Economic Review, 78,* 500–519.

Langholtz, H., Ball, C., Sopchak, B., & Auble, J. (1997). Resource-allocation behavior in complex but commonplace tasks. *Organizational Behavior and Human Decision Processes, 70,* 249–266.

Langholtz, H., Gettys, C., & Foote, B. (1993). Resource-allocation behavior under certainty, risk, and uncertainty. *Organizational Behavior and Human Decision Processes, 54,* 203–224.

Langholtz, H., Gettys, C., & Foote, B. (1994). Allocating resources over time in benign and harsh environments. *Organizational Behavior and Human Decision Processes, 58,* 28–50.

Langholtz, H., Gettys, C., & Foote, B. (1995). Are resource fluctuations anticipated in resource allocation tasks? *Organizational Behavior and Human Decision Processes, 64,* 274–282.

Lopes, L. L. (1996). When time is of the essence: Averaging, aspiration, and the short run. *Organizational Behavior and Human Decision Processes, 65,* 179–189.

Luce, R. D. (1959). *Individual choice behavior.* New York: Wiley.

Lusk, J. L., & Coble, K. H. (2005). Risk perceptions, risk preference, and acceptance of risky food. *American Journal of Agricultural Economics, 87,* 393–405.

Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance, 7,* 77–91.

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

Murnighan, J. K., Roth, A. E., & Schoumaker, F. (1988). Risk aversion in bargaining: An experimental study. *Journal of Risk and Uncertainty, 1,* 101–124.

Myers, J. L. (1976). Probability learning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes, Vol. 3: Approaches to human learning and motivation* (pp. 171–205). Hillsdale, NJ: Erlbaum.

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal, 7,* 308–313.

Northcraft, G. B., & Neale, M. A. (1986). Opportunity costs and the framing of resource allocation decisions. *Organizational Behavior and Human Decision Processes, 37,* 348–356.

Pennings, J. M. E., & Smidts, A. (2000). Assessing the construct validity of risk attitude. *Management Science, 46,* 1337–1348.

Reichle, E. D., & Laurent, P. A. (2006). Using reinforcement learning to understand the emergence of "intelligent" eye-movement behavior during reading. *Psychological Review, 113,* 390–408.

Rieskamp, J. (2006). Perspectives of probabilistic inferences: Reinforcement learning and an adaptive network compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 1371–1384.

Rieskamp, J., Busemeyer, J. R., & Laine, T. (2003). How do people learn to allocate resources? Comparing two learning theories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 1066–1081.

Rieskamp, J., Busemeyer, J. R., & Mellers, B. (2006). Extending the bounds of rationality: A review of research on preferential choice. *Journal of Economic Literature, 44,* 631–661.

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General, 135,* 207–236.

Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior, 8,* 164–212.

Samuelson, P. A. (1963). Risk and uncertainty—A fallacy of large numbers. *Scientia, 98,* 108–113.

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty, 1,* 7–59.

Selten, R. (1991). Evolution, learning, and economic behavior. *Games and Economic Behavior, 3,* 3–24.

Selten, R., & Stöcker, R. (1986). End behavior in sequences of finite prisoner's dilemma supergames: A learning theory approach. *Journal of Economic Behavior and Organization, 7,* 47–70.

Sharpe, W. F. (1981). *Investments.* Englewood Cliffs, NJ: Prentice Hall.

Stahl, D. O. (1996). Boundedly rational rule learning in a guessing game. *Games and Economic Behavior, 16,* 303–330.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction.* Cambridge, MA: MIT Press.

Thomas, E. A. C. (1973). On a class of additive learning models: Error correcting and probability matching. *Journal of Mathematical Psychology, 10,* 241–264.

Weber, E. U., Blais, A.-R., & Betz, N. E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making, 15,* 263–290.

Winefiel. A. (1966). Negative recency and event-dependence. *Quarterly Journal of Experimental Psychology, 18,* 47–54.

Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review, 12,* 387–402.

Yechiam, E., Erev, I., Yehene, V., & Gopher, D. (2003). Melioration and the transition from touch-typing training to everyday use. *Human Factors, 45,* 671–684.

Received March 4, 2005
Revision received September 11, 2006
Accepted September 15, 2006 ■