Myopic Loss Aversion in Investment Behavior

In 1997, Richard Thaler, Daniel Kahneman, Amos Tversky, and Alan Schwartz investigated the effect of myopic loss aversion on investment behavior. Nearly two decades later, this undergraduate thesis research replicates an experiment Thaler et. al. conducted for their investigation. This replication's experimental design is closely aligned with the original experimental design with a few modifications that addresses a few minor flaws. The same statistical analysis used by Thaler et. al. is applied here to compare both similar and conflicting results between the original and the replicated experiment. Additional statistical analysis not conducted by Thaler et. al. is also performed to improve upon the original research. Finally, a mathematical application of several probability theorems including the weak law of large numbers, central limit theorem, and cumulative density are applied in order to quantify rational investment behavior given return distributions that are stochastically determined.

Background

It is well known in the investment world that equities (stocks) carry a large risk premium over low-risk investments such as bonds or treasury bills. That is, the risk premium entails that equities must have a higher average return than low-risk investments to compensate for the high volatility (risk) of the investment. In any given moment, equities are expected to give a higher return than low-risk investments but have a much higher variance of returns. It's entirely possible that equities lose all of their value in a given moment whereas the possibility of a U.S. Treasury Bill doing the same is next to zero.

This relationship between expected return and variance of return, leads to an interesting application of the Weak Law of Large Numbers (WLLN). Equities are at high risk of underperforming low-risk investments in a short time period, but are nearly guaranteed to outperform them in the long-run (how long is not well defined) due to equities' higher expected returns. The long time period erodes the effect of high volatility in returns. This is why professional investment advisers generally advise younger investors to allocate more of their portfolio in equities and older investors to start shifting to lower-risk investments such as bonds or treasury bills. Younger investors have a longer time horizon and therefore have the advantage of the WLLN. Older investors do not have such a luxury.

A simple example of this would be this: Would you accept this bet? I flip a coin. Heads, you win $200. Tails, you lose $100. The expected value of this bet is positive $50. But due to the high risk of the bet, many people would not take this bet since they are averse to the possible major loss. What if I offered you the chance to play this game a hundred or a thousand times? Most people probably would (and should) take the bet since the WLLN is now in their favor.

This recognition of the advantage that a large time period or large sample size gives to a risk-averse decision maker seems to be lost in the amateur investment world. Most amateur investors are unable to think with a long time horizon when it comes to managing their investments. It has been theorized that major losses induce myopic decision making which causes investors to behave this way. This phenomenon is known as myopic loss aversion. Myopic loss aversion stems from the ideas of cumulative prospect theory (CPT), a theory that states that losses carry much greater disutility than the utility from a gain of equal magnitude (Tversky & Kahneman, 1992). As an example, a large percentage of Millennials are very afraid of the stock market and prefer to hold cash in low yielding but very safe investments such as bank accounts (Egan, 2015). This is likely due to seeing both the 2000 Dotcom bubble crash and the 2008 Financial Crisis in which stocks tanked heavily, incurring huge losses to portfolios. These investors are only thinking in the short term however. Had Millennials decided to take the long term view rather than the myopic view, they could have ridden the six year bull market in equities after 2009.

Thaler, Tversky, Kahneman, and Schwartz investigated the effects of myopic loss aversion in investment behavior in an experiment based paper: “The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test” (1997). In this experiment, they gathered 80 undergraduate students and had them sit through an investment simulator in which they chose between two funds: a high variance high reward stock fund that had a significant probability of returning a loss and a low variance low reward bond fund that always had positive returns. Thaler et. al. was interested in whether myopic loss aversion influenced investors to pursue the short term goal of minimizing losses by favoring the bond fund rather than pursuing the more beneficial long term goal of maximizing returns by favoring the stock fund. Furthering this question, they were interested in whether restricting investors' choice flexibility by making choices lock in for extended periods and also restricting investors' information through statistical aggregation of returns would actually improve investing performance. Finally, they were interested in whether investors favored the bond fund due to its lower overall variance rather than trying to avoid losses from the stock fund. This is a distinction between general *risk* aversion versus the more specific *loss* aversion.

To answer the last question, Thaler et. al. shifted the distribution of returns for both of the investments a fixed amount to the right. This was done by adding a fixed amount to any return the distribution generates. The variance of either fund's returns did not change since every point in the distribution moved the same distance to the right. However, if the distribution is shifted substantially to the right, almost all possibilities of it returning a loss are eliminated. If investors stop exhibiting myopic behavior in this case, then that behavior should be attributed to loss aversion rather than general risk aversion.

This mathematics-economics undergraduate thesis research is a close replication of the experiment performed by Thaler et. al. and attempts to answer the same questions that they posed. A successful replication of Thaler et. al.'s results would help increase confidence in the existence of myopic loss aversion and the remedies that address that tendency. Furthermore, the remedies of restricting information and flexibility run against traditional neoclassical economic models that place perfect information and maximum flexibility on a pedestal. Showing that these remedies are in fact helpful to investors would further the credibility of behavioral economics in modeling investment behavior.

Experimental Design

Subjects were informed that they would be participating in an investment simulator in which they had to decide how to allocate their virtual portfolio between two funds, Fund A (the stock fund) and Fund B (the bond fund). Subjects did not know the characteristics of either fund beforehand; they had to learn them through experience. Each period (or ten periods depending on assigned condition), subjects selected from a menu what percentage allocation they wanted for the next period (ten periods). The choices ranged from 100-0% to 0-100% in increments of ten percent. Each subject began with $100 in their virtual portfolio. The simulator lasted for two-hundred periods. After the two-hundredth period, subjects were asked to lock-in an allocation that would last for the next four-hundred periods. Subjects were told that their rewards will be based on their final portfolio value after the six-hundred total periods. The top 20% of portfolios would receive $25 as compensation, next 20%: $20, next 20%: $15, and so on until the minimum reward of $5.

The distribution used to generate returns for each period for Fund A and Fund B were the same distributions used by Thaler et. al. for their stock and bond fund respectively. Henceforth, Fund A will be known as the "Stock Fund" and Fund B will be known as the "Bond Fund" to avoid confusion. The Stock Fund's returns were generated from a normal distribution with a mean return of 1% and a standard deviation of 3.54%. The Bond Fund's returns were generated from a truncated normal distribution with a mean return of 0.25% and a standard deviation of 0.177% with truncation at 0% to prevent losses. It is not clear whether Thaler et. al. intended these returns to be continuously compounded returns or simple returns. Based on a best-guess interpretation of their description of their experimental design, this replication assumes those returns are simple returns. The distinction between the two types of returns is elaborated in the theoretical results section. Whether the returns are considered continuously compounded or simple is not likely to strongly affect experimental results. The relative mean and variance between the two funds are still preserved.

The subjects were split into three experimental groups, the monthly condition, the yearly condition, and the inflated-monthly condition. The monthly condition was the control group in which subjects saw every period's returns and were able to change their investment allocation every period. The yearly condition only allowed subjects to make decisions every ten periods, meaning each allocation was locked in for ten periods. They also only saw the average returns for both funds over ten periods rather than individual period's returns. Subjects in the yearly condition would be seen as at a disadvantage according to neoclassical economic theory. Subjects in the monthly condition could make every choice subjects from the yearly condition chose (by simply choosing the same allocation for ten straight periods) but they could also adjust their allocation every period. This was a flexibility that subjects in the yearly condition did not have. Subjects in the monthly condition also saw fund returns for every period which means they knew the average return over ten periods as well. Subjects in the yearly condition only saw the average return over ten periods and not the individual period returns. The inflated-monthly condition was just like the monthly condition except that both funds had every period's return increased by 5% (the inflation factor) over the course of the experiment (this was readjusted back down to regular levels when determining reward payouts as to not give the inflated-monthly group an unfair advantage) .

At the end of the experiment, all subjects filled out a short exit survey asking them about their familiarity with investing and markets. They were asked if they are or are potentially an economics major. They were asked if they are a part of the [relevant college's] Investment Club or any other investment related entity. They were asked if they own and actively make investments privately. Finally, they were asked about the Dotcom bubble and the Financial Crisis and what they thought was the worst percentage loss in value for stocks in those periods. The information from the exit survey was used to assist in the final data analysis for this experiment by controlling for investment knowledge factors and behaviors. Participants that were already familiar with investing and/or markets may not exhibit myopic loss aversion as strongly. Participants with particularly strong or exaggerated perceptions of recent major downturns may exhibit myopic loss aversion more strongly.

Thaler et. al. did not provide much description or resources on the creation of their investment simulator interface. The investment simulator used for this replication was engineered using the Shiny package within R which is a package intended for interactive data *visualization* but adapted here for interactive data *collection*. This simulator was designed to match Thaler et. al.'s brief description of their simulator with the additional portfolio value information that Thaler et. al. did not include in their version. The simulator accepts a csv file of returns for the funds which allows the experimenter to change the investing experience for the users. As the user progresses through the simulation, R tracks the choices and results of the user's investments locally. At the end of the simulation, to keep data collection automated and streamlined, the GoogleSheets package is used to directly append the user's information, choices, and results to the experimenter's online GoogleSheet. This allows for multiple users to use the simulator on their own server but then aggregate the data into one convenient spreadsheet.

Theoretical Results

An investor with a basic understanding of the relative expected value and variance of the returns of the two funds and the weak law of large numbers would probably figure out that the stock fund would outperform the bond fund over those four-hundred periods with near certainty. Recognizing that, the investor should be comfortable placing most if not all of their portfolio into the Stock Fund for their final choice. However, as shown later in the empirical results section, most people elect not to do so. In fact, the propensity to invest in the Stock Fund depends heavily on the assigned experimental condition. Given that the subjects in the monthly condition received more information over the initial two-hundred periods about the two funds than the subjects in the yearly condition did, one would expect that the subjects in the monthly condition would favor the Stock Fund more than subjects in the yearly condition. This again, did not turn out to be the case.

To prove mathematically that investing your entire portfolio in the Stock Fund for the final choice is the best decision, the following quantifies the probability of the Stock Fund underperforming versus the Bond Fund over the long run. Intuitively, we would expect the probability to decrease as the amount of periods (n) involved increases. We would also expect that as the amount of periods gets very large (more specifically, at around 400), this probability is essentially zero.

To begin, we define the returns to the Stock Fund and the Bond Fund for as r_A and r_B respectively. r_A is distributed Normal(mean=μ_A, s.d.=σ_A) and r_B is distributed Truncated Normal(mean=μ_B, s.d.=σ_B, min.=0). The possibility of analytically solving for the geometric mean of simple returns from these distributions depends on whether the returns are continuously compounded returns or simple returns. The distribution is analytically solvable if these returns are seen as continuously compounded returns (r) rather than simple returns (R). Appendix B will address the second case. Note that converting a continuously compounded return to a simple return *simply* requires exponentiation with base e.

A simple return (and adding 1) for the Stock Fund in period i will be denoted $e^{r_A^i}$ and similarly for the Bond Fund. Equation (1) states the definition of the geometric mean as applied to the simple return for the Stock Fund and then using the laws of exponentiation to arrive at the right-hand side. The number of periods or time horizon of the investment life is indicated by n. The Bond Fund's geometric mean return is derived similarly using its own simple return and is shown in equation (2).

(1) $\left(\prod_{i=1}^{n} e^{r_A^i}\right)^{\frac{1}{n}} = e^{\frac{\sum_{i=1}^{n} r_A^i}{n}}$

(2) $\left(\prod_{i=1}^{n} e^{r_B^i}\right)^{\frac{1}{n}} = e^{\frac{\sum_{i=1}^{n} r_B^i}{n}}$

After establishing equations (1) and (2), we take the ratio of the two equations to find the equation for the ratio of the geometric mean of returns for the Stock Fund and the Bond Fund in equation (3) and apply more laws of exponentiation.

(3) $\frac{e^{\frac{\sum_{i=1}^{n} r_A^i}{n}}}{e^{\frac{\sum_{i=1}^{n} r_B^i}{n}}} = e^{\left(\frac{\sum_{i=1}^{n} r_A^i}{n} - \frac{\sum_{i=1}^{n} r_B^i}{n}\right)}$

Recognize that the arithmetic mean of the continuously compounded returns for the Stock Fund is distributed Normal(mean=μ_A,s.d.=$\frac{\sigma_A}{\sqrt{n}}$). Also recognize that the arithmetic mean of the continuously compounded returns for the Bond Fund converges in distribution to a Normal(mean=μ_B, s.d.=$\frac{\sigma_B}{\sqrt{n}}$) distribution by applying the Central Limit Theorem. Given this, we know that the exponent of the right-hand side of equation (3) converges to the difference between two normal distributions which is also distributed normally. The asymptotic distribution of the exponent is then Normal(mean=$\mu_A - \mu_B$, s.d.=$\sqrt{\frac{\sigma_A^2+\sigma_B^2}{n}}$). This leads us to equation (4) which recognizes that the right-hand side of equation (3) is the exponentiation of a normal distribution with base e. This means that the ratio of the geometric mean of returns of the Stock Fund to the Bond Fund converges in distribution to a lognormal distribution.

(4) $e^{\left(\frac{\sum_{i=1}^{n} r_A^i}{n} - \frac{\sum_{i=1}^{n} r_B^i}{n}\right)} = Ratio_{ab}$ ~ LnNormal($\mu_R = \mu_A - \mu_B$, $\sigma_R(n) = \sqrt{\frac{\sigma_A^2+\sigma_B^2}{n}}$)

We are interested in the probability that the ratio is less than 1 meaning that the Stock Fund has underperformed the Bond Fund. Interestingly and perhaps also intuitively, the cumulative distribution function as a function of x for a lognormal distribution is the cumulative distribution function of the underlying normal distribution as a function of Ln(x).

(5) $C.D.F.Ratio_{ab}(x, \mu_R, \sigma_R(n))= \Phi(\frac{\ln(x)-\mu_R}{\sigma_R(n)})$ Where Φ(.) is the cumulative distribution function for a standard normal distribution

That is to say, the cumulative distribution function of a lognormal distribution is fully characterized by the underlying normal distribution. In application to this specific instance, this means that the cumulative distribution of the *ratio* of the *geometric mean of simple returns* from two funds is fully characterized by the cumulative distribution of the *difference* of the *arithmetic mean of continuously compounded returns* from the two funds.

Take the perspective of an experimental subject who has completed the initial two-hundred periods of investment simulation and now must make a choice that locks in for four-hundred periods. Clearly, one would want to invest their money in the fund with the higher expected return but also understand the risks of having that fund underperform. Given the computation of the parameters for $C.D.F.Ratio_{ab}$, we can establish the probability of the Stock Fund underperforming the Bond Fund over four-hundred periods as $C.D.F.Ratio_{ab}(1, 0.0075, 0.001772211)$. This will evaluate to a 0.001% probability that the Stock Fund underperforms the Bond Fund over four-hundred periods. It would take an enormous amount of risk-aversion, or an enormous amount of irrationality, for someone to put any amount of their portfolio into Fund B.

Given x = 1, the argument inside of Φ(.) is $\frac{-\mu_R}{\sigma_R(n)}$. As $n \to \infty$, $\sigma_R \to 0$, which means $\frac{-\mu_R}{\sigma_R(n)} \to -\infty$. In this case, $\Phi\left(\frac{-\mu_R}{\sigma_R(n)}\right) \to 0$. This means that as the number periods over which we are evaluating fund returns approaches infinity, the probability of the Stock Fund underperforming the Bond Fund approaches zero, otherwise speaking, impossible. Indeed, given that Φ(.) is monotonically increasing and that $\frac{-\mu_R}{\sigma_R}$ is monotonically decreasing as a function of n, the probability of the Stock Fund underperforming the Bond Fund decreases monotonically as n increases.

Thaler et. al. does not make fully clear whether the returns generated from the normal and truncated normal distributions are continuously compounded returns or simple returns. In finance theory, returns that come from normal distributions are generally continuously compounded returns. This allows for the simple returns, and therefore asset price, to be distributed lognormally. It is unreasonable to allow for the possibility of greater than 100% loss from an investment, since assets cannot be valued below zero. The lognormal distribution sets that limitation whereas the normal distribution does not (Zucchi, 2014). However, based on a best guess of Thaler et. al.'s work, it appears that they chose (incorrectly) to use a normal distribution (and truncated normal distribution) to model simple returns and asset prices. This experimental replication has assumed the same. Estimating the probability distribution of the ratio of the two funds when simple returns are distributed normally (or truncated normally) can only be done via simulation.

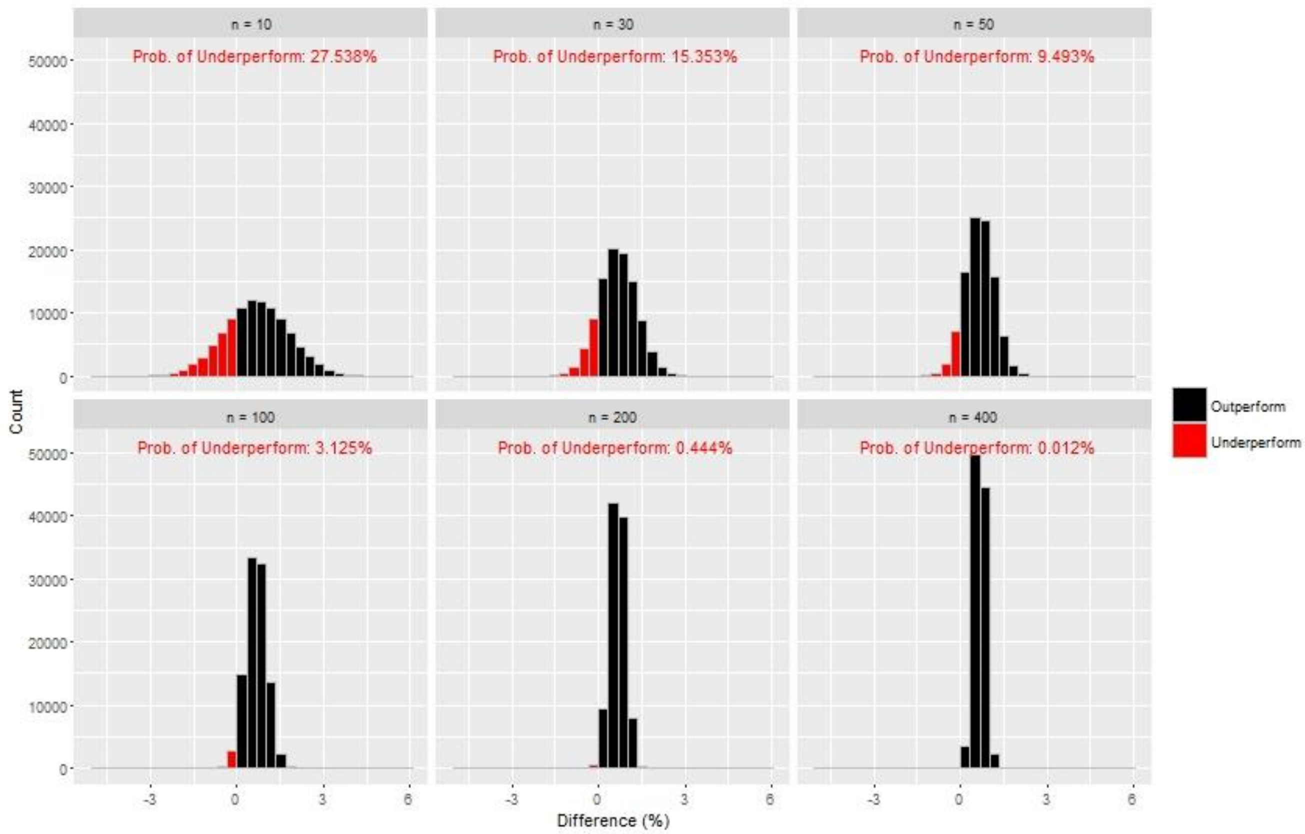

Figure 1. Simulated distribution of difference in geometric mean return between the Stock Fund and the Bond Fund over various sample sizes
Each distribution has 100,000 simulated observations. Red bars indicate underperformance (difference is negative) from the Stock Fund relative to the Bond Fund. On average, the Stock Fund outperformed the Bond Fund by 0.66% per period.

To simulate the distribution of relative performance between the Stock Fund and Bond Fund when assuming Thaler et. al.'s returns are simple returns, n observations (time horizon) are drawn from the distributions for both funds. The geometric mean of those n observations are computed for both funds and the difference between the two (Stock Fund – Bond Fund) is saved as one observation in the distribution. A ratio was used in Appendix A to ensure that the distribution could be solved analytically. Given that constraint no longer applies, the difference is used here which is a more common metric of comparing returns[1]. This process is repeated 100,000 times for n = 10, 30, 50, 100, and 400 for a total of 600,000 times. The resulting distributions are presented in Figure 1.

The resulting distributions appear to be normally distributed. Kolmogorov-Smirnov tests could not reject the hypothesis that the distribution is normal for all six distributions. Notably, the WLLN seemed to be in effect here. As n increases, the variance of the distribution decreases, clustering the distribution around the mean (roughly 0.66%). Due to this effect, the probability of the Stock Fund underperforming the Bond Fund ($x < 0$) decreases as n increases. This is very similar to the findings in Appendix A where the returns are assumed to be continuously compounded. With a four-hundred period time horizon, the probability of the

[1] Note that if the difference between two returns is negative, it implies that the ratio is also less than one. Thus the probability of underperformance remains the same whether difference or ratio is used.

Stock Fund underperforming the Bond Fund is only 0.012%, again showing that it would take an enormous amount of risk-aversion or irrationality for a subject to allocate to the Bond Fund in the final choice. Table 1 demonstrates that whether one assumes continuously compounded returns or simple returns does not affect the effect of long time horizons on reducing the probability of the Stock Fund underperforming the Bond Fund.

N (time horizon)	**Cont. Compounded**	**Simple**
10	25.170%*	27.538%
30	12.323%	15.353%
50	6.730%	9.493%
100	1.717%	3.125%
200	0.138%	0.444%
400	0.001%	0.012%

Table 1. Comparison of probability of underperformance of Stock Fund relative to Bond Fund for various time horizons
*Analytically solving for the distribution in appendix A required the application of the Central Limit Theorem which is only valid with large sample sizes. A sample size of 10 may not be sufficient.

Empirical Results

The following analysis of the results from the replicated experiment will present tables and figures with an alpha level of both 0.01 and 0.05. The analysis will be presented with an alpha level of 0.01 to mimic Thaler et. al.'s standard. Later on, statistical conclusions that change when a more conventional alpha level of 0.05 is used will be discussed.

Figure 2 displays the dispersion of final allocations across subjects. A few subjects were able to discern that it was in their best interest to put all of their money in the Stock Fund (none in the Bond Fund) in all three conditions. Figure 2 also shows that statistical aggregation of returns (yearly condition) and distribution shifting to mostly eliminate losses (inflated monthly condition) were very effective in causing subjects to allocate towards the Stock Fund instead of the Bond Fund. In the monthly condition, twelve of eighteen subjects chose to allocate at least 70 percent of their portfolio to the Bond Fund as their final allocation. Only two of sixteen and six of nineteen did so for the yearly and inflated monthly conditions respectively. Finally, although the yearly and inflated monthly conditions had relatively normally distributed data, the monthly condition had what appears to be a bimodal distribution.

A likely explanation for this case is that the monthly condition had the strongest effect in inducing myopic loss aversion which clumped most subjects' allocation choice to the right of the distribution. However, there were enough subjects shrewd enough to realize that the Stock Fund was most likely going to outperform the Bond Fund in the long run and chose to allocate their entire portfolio to the Stock Fund. These few subjects clumped to the left of the distribution, completing the bimodal distribution. All of these results are extremely similar to Thaler et. al.'s findings in their experiment.

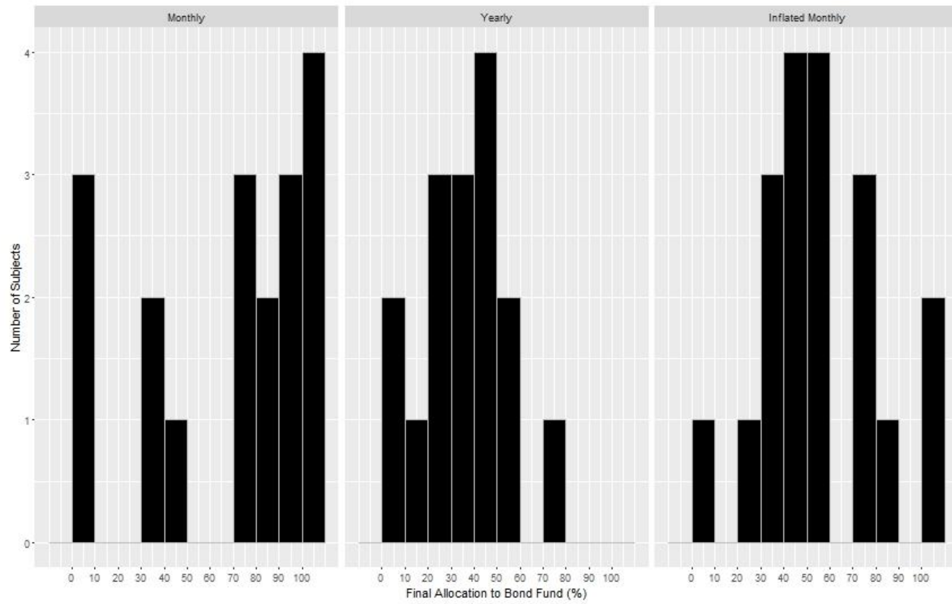

Figure 2. Distribution of final allocation choices for each condition

Statistics for the average allocation to the Bond Fund for the final choice and the last forty periods (this is the last forty choices for the monthly and inflated monthly conditions and the last four choices for the yearly condition; not counting the final choice that locked for four-hundred periods) are presented in Table 2 through Table 3B respectively. The means for each condition in both tables are nearly identical to the means in Thaler et. al.'s experiment. The one exception is that the average final choice to the Bond Fund for the inflated monthly condition here is much higher than what Thaler et. al. found (50.53% to 27.6%).

Welch's t-tests are applied to find any differences in average allocation to the Bond Fund between conditions. The results for the final choice are very similar to the findings Thaler et. al. found. For this replication, there is a statistical difference between the monthly and yearly conditions (p=0.003) but not between the yearly and inflated-monthly conditions (p=0.013). Thaler et. al. find a statistical difference between the monthly and inflated monthly condition whereas there is no statistical difference here (p=0.233). Although the experimental design randomly assigned students into conditions, there remained the possibility that one condition contained a significantly higher proportion of subjects with more investment experience. Adding these controls did not change any statistical significance conclusions but did reduce the p-value when comparing the monthly and inflated monthly conditions from 0.233 to 0.071. This may indicate that the effect of the inflated condition on subjects was masked by unlucky grouping of subjects with investment knowledge into one condition.

Condition	N	Mean	SD	SE
Monthly^{1}	18	63.33	36.78	8.67
Yearly^{a}	16	30.63	18.79	4.70
Inflated Monthly^{a1}	19	50.53	25.92	5.95

Table 2. Percent of portfolio allocated to the Bond Fund for final choice
Conditions with common letter (number) subscripts denote no statistical difference according to a Welch's t-test with α=0.01 (α=0.05).

Condition	N	Mean	SD	SE
Monthly[a]	18	57.57	28.10	6.62
Yearly[b]	16	30.00	12.91	3.23
Inflated Monthly[ab]	19	41.51	13.06	3.00

Table 3A. Percent of portfolio allocated to the Bond Fund for final 40 periods (averaged by subject)
Conditions with common letter (number) subscripts denote no statistical difference according to a Welch's t-test with $\alpha=0.01$ ($\alpha=0.05$).

Condition	N*	Mean	SD	SE**
Monthly	720	57.57	36.51	1.36
Yearly[a]	64	30.00	19.44	2.43
Inflated Monthly[a]	760	41.51	26.82	0.97

Table 3B. Percent of portfolio allocated to the Bond Fund for final 40 periods (cluster-robust standard errors)
Conditions with common letter (number) subscripts denote no statistical difference according to a Welch's t-test with $\alpha=0.01$ ($\alpha=0.05$). *Not fully independent **Unclustered standard errors

The fact that the monthly condition final choice distribution does not appear to be normal puts the validity of the t-test results in jeopardy. T-tests assume data is drawn from a normal distribution. To address this, non-parametric bootstrap resampling was performed to compare the differences of means involving the monthly condition data. The statistical significance conclusions from the resampling are not different from the statistical significance conclusions of the t-tests. The distributions of the resampled differences and 99 percent confidence intervals for the test-statistics are presented in Figure 3A-B.

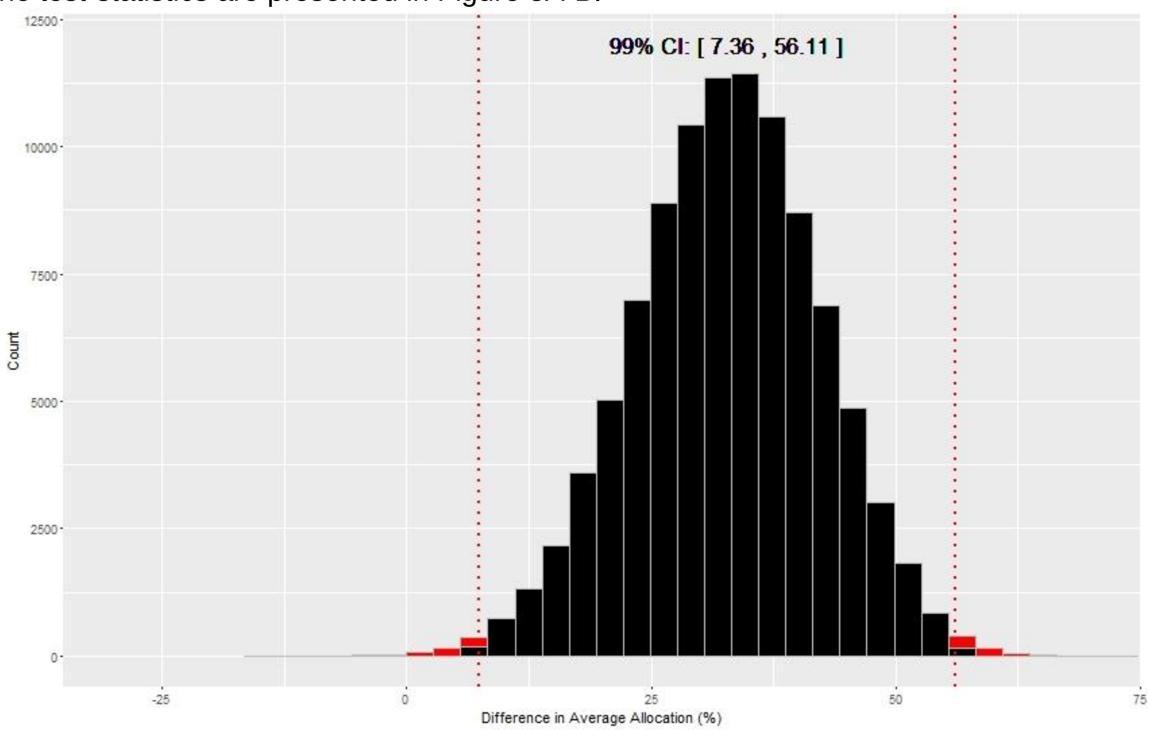

Figure 3A. Bootstrapped differences in average final choice allocation to the Bond Fund for the monthly and yearly conditions
Red bars represent values outside of the 99% confidence interval. 95% confidence interval: [13.47, 50.90]

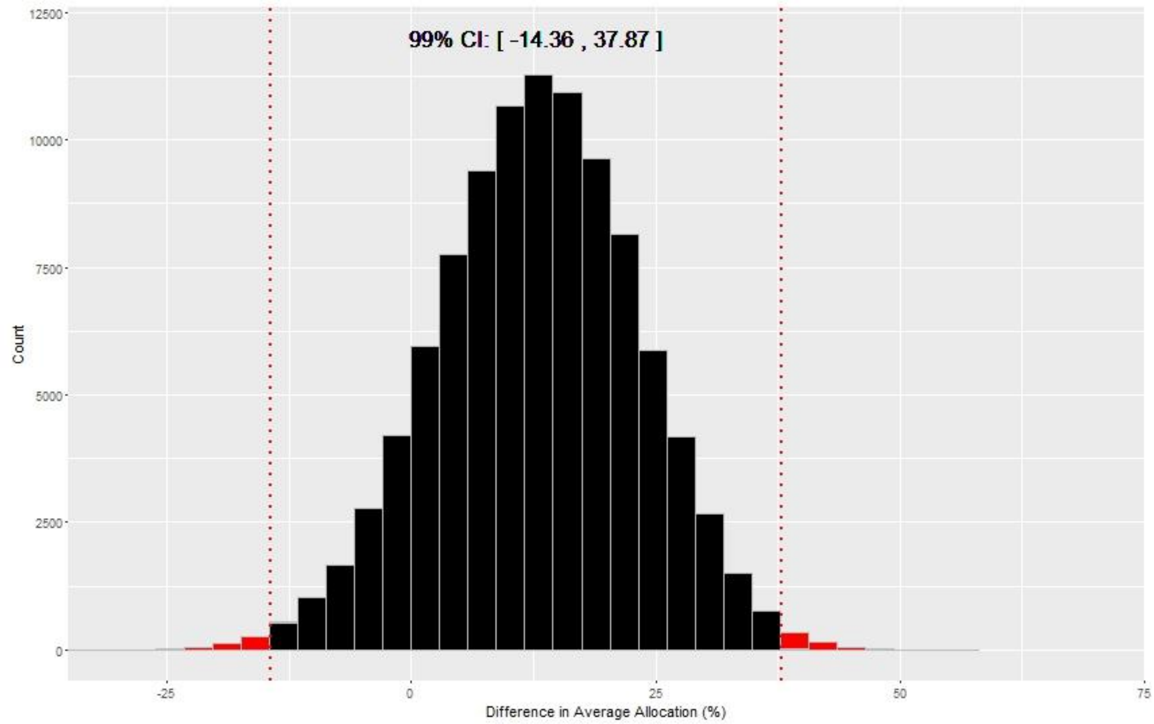

Figure 3B. Bootstrapped differences in average final choice allocation to the Bond Fund for the monthly and inflated-monthly conditions
Red bars represent values outside of the 99% confidence interval. 95% confidence interval: [-7.66, 32.31]

It is reasonable to suspect correlation of choices within subjects over time. Indeed, empirical autocorrelation functions (Figure 4 A-C) confirmed this suspicion for many subjects, with strong correlation between choices with short (1-2) lags. There is less evidence of autocorrelation for subjects in the yearly condition but this is likely a result of lack of statistical power due to lower sample size than the other two conditions. Correlation of choices within subjects violates the basic independence assumptions in the t-tests. Thaler et. al. addressed this issue by using the *average* allocation choice over the last forty periods for each subject (aggregation method). This method is applied to the data and presented in Table 3A. There is a statistical difference between the monthly and yearly conditions ($p<0.001$) but no difference between the monthly and inflated-monthly conditions ($p=0.037$) nor between the yearly and inflated-monthly conditions ($p=0.014$).

An issue with the aggregation method is that it throws away information on the dispersion of choices for each subject and treats the data as if each subject only made one choice. Table 2B displays the results of t-tests when using cluster-robust standard errors to account for autocorrelation within each subject. This changes the p-values of the above tests to $p<0.001$, $p=0.003$, and $p=0.012$ respectively. Whereas the aggregation method couldn't reject the null hypothesis that subjects in the monthly condition allocated differently from subjects in the inflated-monthly condition, the cluster-robust method could. These cluster-robust tests also provided intra-cluster correlation factors that were quite significant (monthly/yearly: 0.559, monthly/inflated-monthly: 0.445, yearly/inflated-monthly: 0.215) which further confirmed the need to adjust regular standard errors.

Paired sample t-tests were performed to compare the final choice with the average of the last forty periods for each condition. None of the three conditions had a statistical difference between the two choices (for either alpha level), an identical result to what Thaler et. al. found. This observed behavior is irrational. In the last forty periods, subjects were allocating assets for only one or ten periods at a time, with a very limited time horizon (less than forty periods). For the final choice, subjects were explicitly told that the time horizon for their final choice would be four-hundred periods, which when applying the WLLN, subjects should have taken the opportunity to accept more risk than they did in the last forty periods. This irrational behavior further strengthens the argument that these subjects were suffering from myopic loss aversion. The subjects were so myopic that they did not distinguish between an allocation for one or ten periods and four-hundred periods.

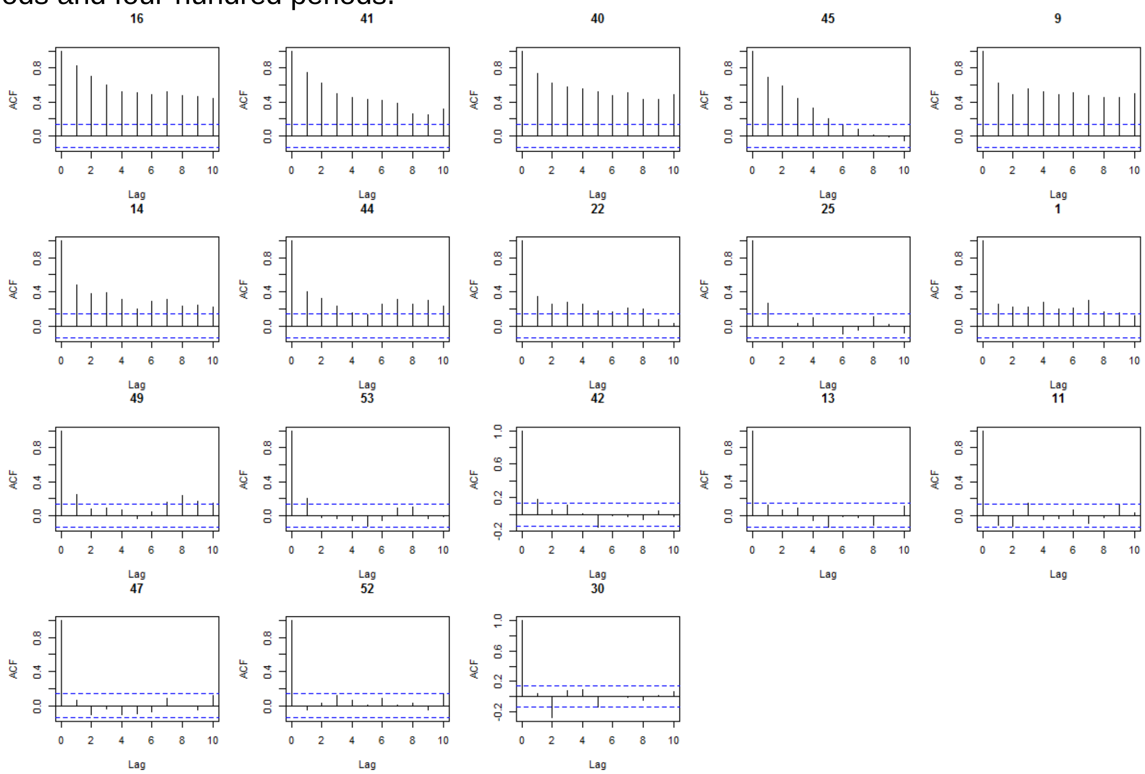

Figure 4A. Monthly condition subjects' autocorrelation functions (ACF) of allocation choices
ACFs that extended above or below the dotted lines indicate that they are statistically significantly different from zero. Numbers above each graph are the unique ID numbers for each subject. Graphs are organized by highest absolute ACF for lag 1.

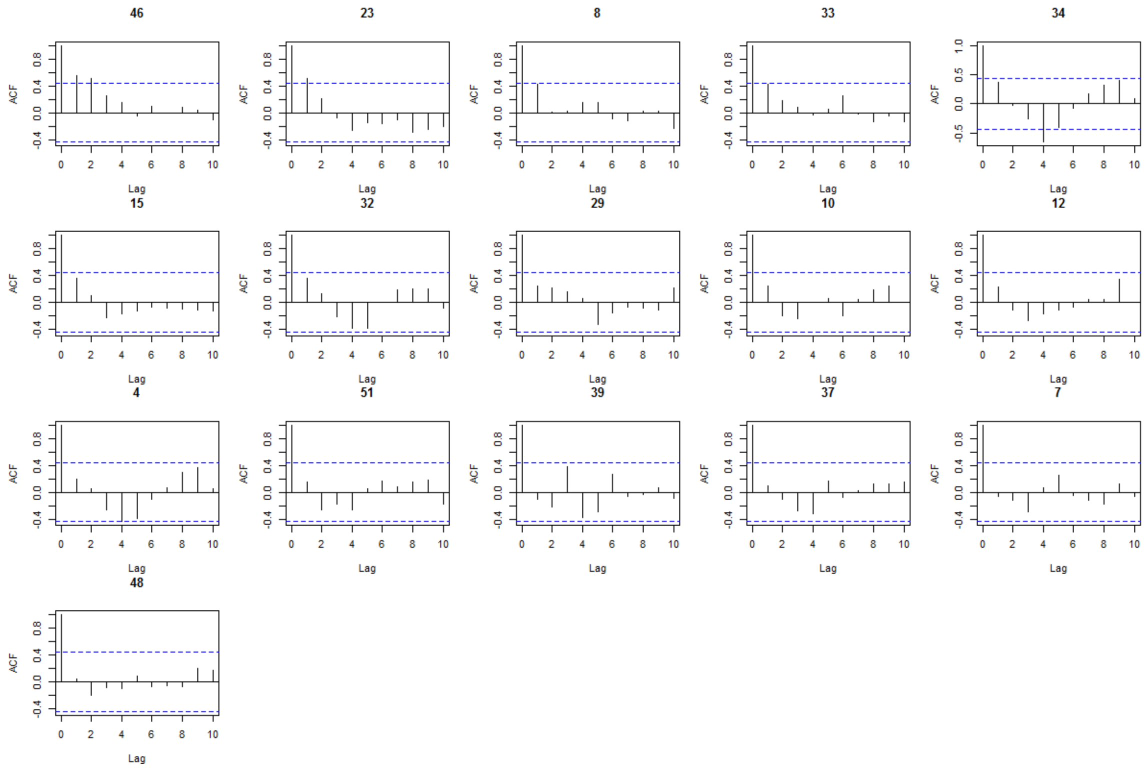

Figure 4B. Yearly condition subjects' ACFs of allocation choices

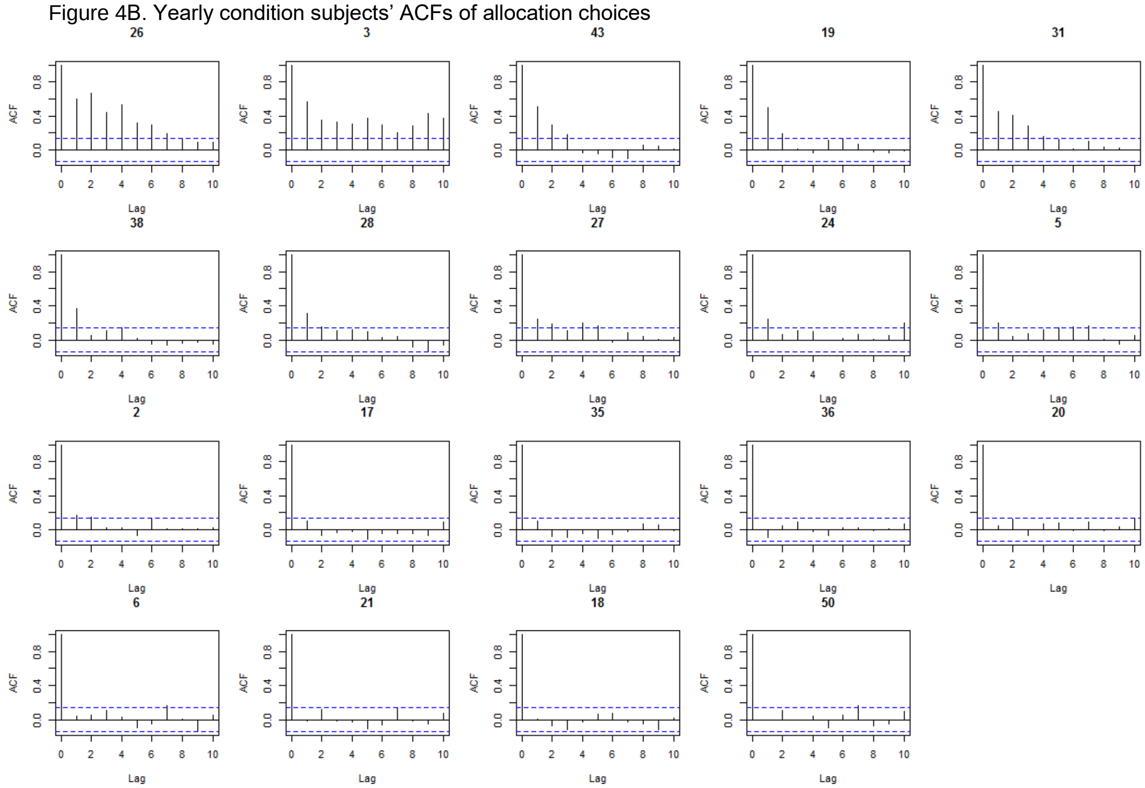

Figure 4C. Inflated monthly condition subjects' ACFs of allocation choices

Next, portfolio allocations over time for each condition were examined. Studying average portfolio allocations over time allows us to examine how subjects learned and adapted to the simulator over time. Figure 5A compares the monthly and yearly conditions. Figure 5B compares the monthly and inflated monthly conditions. At nearly every period, subjects in the monthly condition allocated more of their portfolio to the Bond Fund on average than the subjects in the yearly condition. The same could be said when comparing subjects in the monthly condition with subjects in the inflated monthly condition. Subjects in the monthly condition appear to have their allocations drift towards the Bond Fund over time. Subjects in the yearly and inflated monthly condition appear to have their allocations drift in the opposite direction, with the yearly condition drifting significantly more than the inflated monthly condition.

The coefficient estimates for a pooled-data regression for allocation to the Bond Fund on the trial number (allocation decision) in Table 4 confirm the drifts. Cluster-robust standard errors by subject were used. Only the yearly condition has a statistically significant drift. The yearly condition also has a large estimated drift, a decrease of nearly one percent per trial (every ten periods) compared to the other two conditions with relatively small drifts. This statistically significant negative drift indicates that subjects in the yearly condition became less and less averse to the Bond Fund over time. The same could not be said for the subjects in either the monthly or inflated monthly condition. These results are again nearly identical to Thaler et. al.'s findings except that they find a statistically significant negative drift for their inflated monthly condition as well. However, Thaler et. al. do not use cluster-robust standard errors in their regression which means their standard errors are invalid due to autocorrelation within each subject.

A fixed effects model (varying intercepts only) with regular standard errors and a fixed effects with cluster-robust standard errors model were applied to the data as well. In the fixed effects model with regular standard errors, the trends for all three conditions are statistically significant. However, once cluster-robust standard errors are included, the standard errors return to virtually the same values they were in the first model that did not include fixed effects. This indicates that the fixed effects were unable to pick up the autocorrelation within each subject.

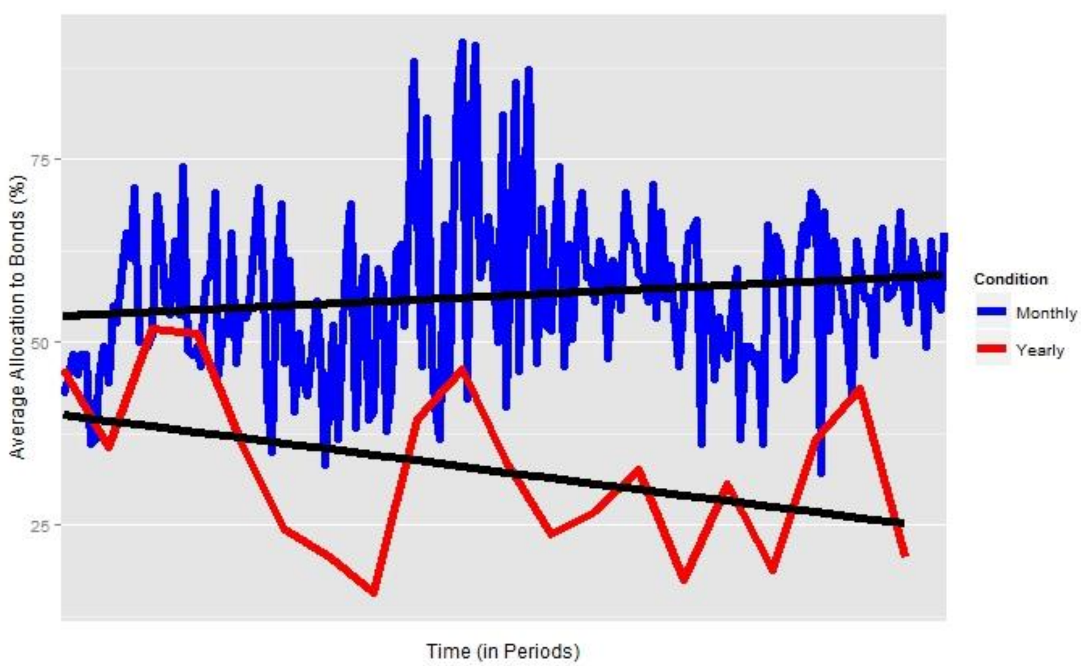

Figure 5A. Comparing average allocations over time (monthly vs. yearly conditions)

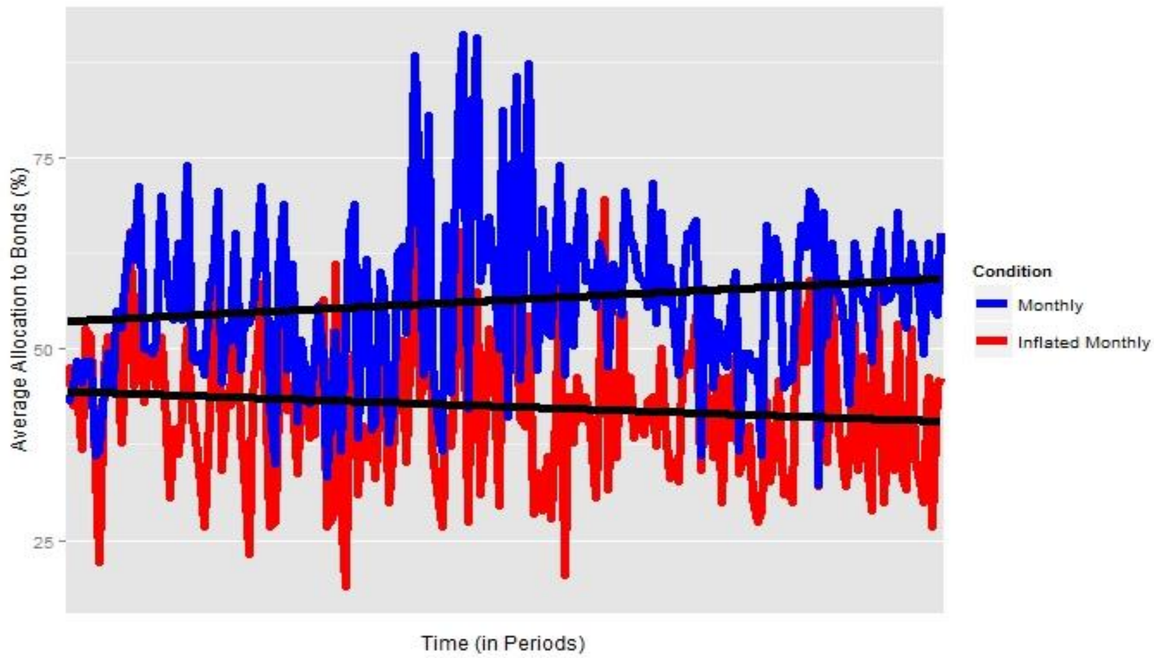

Figure 5B. Comparing average allocations over time (monthly vs. inflated monthly conditions)

Condition	Trial number	Intercept
Monthly	0.028	53.598
Cluster-robust	(0.032) p=0.381	(4.365) p<0.001**
Fixed effects	(0.008) p<0.001**	n/a
Cluster-robust+fixed effects	(0.032) p=0.382	n/a
Yearly	-0.782	40.836
Cluster-robust	(0.160) p<0.001**	(2.385) p<0.001**
Fixed effects	(0.163) p<0.001**	n/a
Cluster-robust+fixed effects	(0.164) p<0.001**	n/a
Inflated Monthly	-0.020	44.458
Cluster-robust	(0.016) p=0.225	(1.987) p<0.001**
Fixed effects	(0.007) p=0.005**	n/a
Cluster-robust+fixed effects	(0.016) p=0.226	n/a

Table 4. Regressions predicting allocation to Bond Fund (%) from trial number
Final allocation choice is excluded. Three standard errors of the trial number coefficient are reported in parenthesis for each of the three different models. * p<0.05 ** p<0.01

Since Thaler et. al.'s experiment performed in 1997, there have been two major stock market crashes. In the post-simulator survey, participants were also asked if they knew what the "dotcom bubble of 2000-2001" and the "financial crisis of 2008-2009" were and if so, name their best guess as to how much value stocks lost over the crashes. Only about half of the subjects (31 out of 53) knew what the Dotcom bubble was, likely due to the young age of the subjects during the time of the crash. Most subjects (47 out of 53) knew what the Financial Crisis was. Final choice allocation was regressed on dummy variables indicating knowledge of the Dotcom bubble and the Financial Crisis. Additionally, for subjects who recalled the crashes, final choice allocation was regressed on best guesses of stock value loss in the Dotcom bubble and also the Financial Crisis to see if subjects who recalled high losses in stocks tended to favor bonds in the simulator. There were no statistically significant results in any of the regressions as shown in Table 5.

Regression	**Dotcom dummy**	**Fin. Crisis dummy**	**Stock value loss (Dotcom)**	**Stock value loss (Fin. Crisis)**
Knowledge of crashes (N=53)	-0.343 (9.685) p=0.972	2.708 (15.062) p=0.858		
Recollection of stock value loss in dotcom bubble (N=31)			-0.124 (0.130) p=0.349	
Recollection of stock value loss in financial crisis (N=47)				0.057 (0.078) p=0.470

Table 5. Regressions of final allocation to Bond Fund (%) on various measurements of finance knowledge and recollection of stock value loss
Standard errors in parenthesis

Thaler et. al. chose an alpha level of 0.01 for all of their statistical tests, a much stricter standard than what is conventional. Although using a more relaxed standard of 0.05 as an alpha level does not change most of my statistical conclusions, there are a few tests that are affected. In Tables 2-3B, comparisons between yearly and inflated-monthly conditions are affected across the board. At an alpha level of 0.01, there is no statistical difference between the two groups in any test, at an alpha level of 0.05, there *is* a statistical difference in *every* test involving the two groups. This would imply that the effectiveness of aggregating returns and forcing subjects to make long term choices is stronger than simply shifting the distribution of returns to avoid losses in addressing myopic loss aversion. This conclusion is supported by the fact that evidence of aggregation and long-term commitment being effective is much stronger than distribution shifting when comparing to the control condition (the monthly condition).

This statistical conclusion diverges from what Thaler et. al. found in any of their tests. One explanation for this is that it is simply due to the difference in alpha level chosen. However, since Thaler et. al. do not provide information on results when the alpha level is 0.05 for their data, this explanation is only speculated.

Conclusion/Discussion

Overall, the results found in this replication were very similar to the results found in Thaler et. al.'s experiment. As measured by several methods, subjects in the monthly condition were overall more likely to allocate a higher percentage their portfolios to the Bond Fund in comparison to subjects in the yearly and inflated monthly conditions. This phenomenon was most likely caused by elimination of perceived losses either by statistical aggregation or distribution shifting. Given the similarity in results, Thaler et. al.'s conclusion that myopic loss aversion exists still stands. Their proposed remedy of forcing people to think and observe long term trends and restricting information on short term trends stands as effective since losses are minimized through statistical aggregation in the yearly condition. However, the evidence is weaker in this replication for the effectiveness of shifting distributions in reducing myopic loss aversion. Thaler et. al. found much stronger evidence that the inflated monthly condition is effective in reducing myopic loss aversion in their experiment, nearly on par with the yearly condition, than in this replication.

This difference in results could be attributed to several explanations. It could simply be a lack of statistical power since the sample size here was not large. The point estimates for the final choice indicate that subjects in the monthly condition tended to favor the Bond Fund significantly more than subjects in the inflated monthly condition (63.33% to 50.53%). There just appeared to be a lack of statistical confidence in rejecting the null hypothesis that there is no difference between the two conditions. Another explanation is that this replication used a significantly lower inflation factor for the inflated monthly condition that what Thaler et. al. used. In discussion with advisors, the 10% inflation factor used by Thaler et. al. was deemed too extreme and lowered for this replication to 5%. This resulted in 3% (6 of 200) of the inflated monthly returns for the Stock Fund still being negative. Thaler et. al. managed to inflate away all losses for their Stock Fund with the higher inflation factor. It is possible that the subjects were so loss-averse that even seeing a few negative returns from the Stock Fund caused them to shy away from it and towards the Bond Fund. In this case, this lowered magnitude of effect coupled with lowered significant evidence of reducing allocation to the Bond Fund serves to bolster the evidence for the existence of myopic loss aversion. In order for the shifting of distributions to be fully effective, one must eliminate all possibility of loss rather than just most. However, further testing is required to fully examine this hypothesis.

If the true effect of the shifting of distributions is indeed strong as found in Thaler et. al.'s experiment, it bodes for an interesting implication. In periods of high inflation, stocks would have to increase their nominal return in order to maintain their real return. This effect is very similar to the shifting of distribution done in the inflated-monthly condition. Though the real return would not have changed for stocks, the nominal returns would include much fewer losses than before. If the inflation is unexpected and investors are subject to the money illusion (inability to distinguish nominal versus real) as mentioned by Thaler et. al., this could result in a temporary increase in the attractiveness of stocks for investors due to the decrease in losses.

Given the resounding effectiveness of the yearly condition in helping subjects make better investment choices, it certainly seems appropriate to consider some practical policy changes. Although it would be unlikely for brokerage companies to completely restrict information and trading ability (that's their entire business model) to amateur investors, they could reduce exposure in order to help their clients make better long-run decisions rather than to encourage day-trading. Instead of presenting daily or even hourly information on how their investments are doing, brokerages could have long-term (ten or more years) aggregated returns displayed initially on their web page or mobile app. Also, they could limit the amount of trades their clients can do on their own per month. To activate additional trading ability beyond that limit, the clients must call the brokerage and speak to an investment professional before receiving that privilege. Indeed, the observation that giving more information and allowing more trading flexibility actually hurts investors is not only found in Thaler et. al.'s experiment and this replication. Malkiel centers the theme of his book "A Random Walk Down Wall Street" around the fact that it is extremely difficult to defeat a long-run passive investing strategy of holding index funds with an active investing (such as day-trading) strategy (2007). This holds for all investors, not just amateurs.

Unfortunately, brokerages make significant profit from trading fees. They would not be interested in policies that encourage patience and passiveness. Myopic loss aversion in amateur investors actually benefits the brokerages, who don't mind major losses in the market because it induces investors to use their services. These policies proposals may be better suited for brokerages such as Vanguard and Fidelity that focus on managing retirement accounts for their clients, which are generally long-run prospects.

It is quite interesting that the field of economics only considered risk-aversion when evaluating decision making for an extremely long time until prospect theory was proposed in 1979 and then expanded into cumulative prospect theory in 1992. In fact, the proposal was so groundbreaking to the field of economics it warranted a Nobel Prize to the authors. Though the underlying equations of prospect theory are extremely complex, the basic idea of the theory was already intuitive to many. Malkiel notes in his book, which was first published in 1973, that most investment professionals considered variance an inadequate measure of the riskiness of an investment. Surely, they thought, no one would consider the possibility of earning a return greater than expected as a risk. However, this is exactly what risk-aversion theory states. CPT alleviates this issue by using loss aversion to weigh losses more appropriately relative to gains.

Though CPT and loss aversion seem intuitive, some criticism emerges when examining the experimental design of Thaler et. al.'s work and this replication. Thaler et. al.'s claim that the yearly condition contains less information than the monthly condition certainly holds for a machine with the ability to record information, but what about humans? Humans on average can only remember about seven numbers at once (Miller, 1956). How then could a subject remember twenty returns (ten per fund) and also be able to compute the average return for both funds all from memory? In the yearly condition, the simulator already does that work for the subject. In this case, subjects in the yearly condition have the critical piece of information (long term trends) necessary for success in the simulator. Subjects in the monthly condition lack this information due to the weakness of human memory and computation abilities. To account for this issue, the simulator could display information about the trailing ten period average returns for both funds in the monthly condition. This modification to the experimental design would be the next step in researching myopic loss aversion in investment behavior.

Bibliography

Egan, M. (2015, March 11). Crisis hangover: Millennials are scared to invest. Retrieved September 22, 2015, from http://money.cnn.com/2015/03/11/investing/investing-millennials-stocks-markets/index.html

Malkiel, B. G. (2007). *A Random Walk Down Wall Street: The Time-Tested Strategy for Successful Investing* (Ninth Edition edition). New York: W. W. Norton & Company.

Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test. *The Quarterly Journal of Economics*, *112*(2), 647–661. http://doi.org/10.1162/003355397555226

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. http://doi.org/10.1007/BF00122574

Zucchi, K. (2014, October 20). Lognormal and Normal Distribution. Retrieved January 10, 2016, from http://www.investopedia.com/articles/investing/102014/lognormal-and-normal-distribution.asp