

FinVis: Applied Visual Analytics for Personal Financial Planning

Stephen Rudolph, Anya Savikhin, and David S. Ebert, *Member, IEEE*

Abstract—FinVis is a visual analytics tool which allows the non-expert casual user to interpret the return, risk and correlation aspects of financial data and make personal finance decisions. This interactive exploratory tool helps the casual decision-maker quickly choose between various financial portfolio options and view possible outcomes. FinVis allows for exploration of inter-temporal data to analyze outcomes of short or long-term investment decisions. FinVis helps the user overcome cognitive limitations and understand the impact of correlation between financial instruments in order to reap the benefits of portfolio diversification. Because this software is accessible by non-expert users, decision-makers from the general population could benefit greatly from using FinVis in practical applications. We quantify the value of FinVis using experimental economics methods and find that subjects using the FinVis software make better financial portfolio decisions as compared to those using a tabular version with the same information. We also find that FinVis engages the user, which results in greater exploration of the dataset and increased learning as compared to a tabular display. Further, participants using FinVis reported increased confidence in financial decision-making and noted that they were likely to use this tool in practical application.

Index Terms—Casual Information Visualization, personal finance, aggregate risk visualization.

1 INTRODUCTION

Financial planning is a challenging process in which the decision-maker must consider abstract and dynamic concepts such as risk and correlation of financial instruments. Empirical and laboratory research in economics and finance has shown that individuals often make sub-optimal decisions when choosing a financial portfolio due to cognitive limitations [3, 12, 32, 20, 19]. Critical problems faced by individuals are finding the right investment and optimally allocating investments. Optimal asset allocation is often the most important decision for a portfolio's performance [3]. FinVis is a novel tool that can help individuals improve decision-making in financial planning.

American workers are increasingly responsible for their own retirement; therefore, it is important to have a successful financial plan and to select a suitable portfolio of investments. However, a recent study found that less than 25% of individuals over the age of 50 succeeded at saving for retirement, often because they were unaware of fundamental economic concepts such as interest rate or inflation [21]. Researchers have called for the development of decision support systems to aid individuals in understanding financial planning [3]. Systems with textual feedback and graphs can improve decision making in simple tasks [3]. However, visual analytics tools are needed for complex decision-making involving long-term investment planning. Further, these tools must be accessible to the casual user. For this purpose we have developed FinVis, a visual analytics tool which allows the casual untrained user to easily assess and explore information about potential life-long impacts of many financial decisions simultaneously and interactively.

Our key contribution is a visual analytics tool, FinVis, that helps the individual improve decision-making by visualizing aggregate risk alongside a traditional wealth-time plot in a casual context for personal finance data. Since this software is accessible to non-expert users, general users could benefit greatly from using FinVis in practical applications. FinVis is flexible and allows the user to enter as much or as little information as he/she wants and evaluate potential outcomes of multiple decisions over his/her lifetime. The improved decision making which resulted from using FinVis in a laboratory setting suggests that the use of these techniques in practice could improve the ability

of the general public to allocate finances in an appropriate manner. This could have the effect of reducing bankruptcies and foreclosures and increasing the amount of money available at retirement, thereby improving the financial health of households and the economy as a whole.

Our goals were to introduce FinVis as an effective tool and to assess the effectiveness of FinVis to aid individuals in financial decision-making. We designed an economic experiment that allows us to compare decisions made in a financial planning task for a group using FinVis versus a group using a tool with textual representations of the same data. Laboratory experiments such as this one are ideal for explaining decision-making and quantifying the value of visual analytics techniques [26]. We found that FinVis improves decision-making, learning, and exploration of the dataset, and increases user confidence in their understanding of financial planning.

2 RELATED WORK IN VISUALIZATION

FinVis builds on work from several key research areas, including financial visualization [30, 24, 33, 4], casual information visualization [25], temporal visualization [13], and risk visualization [7, 27, 33]. While use of computer-assisted visualization techniques to show financial data goes back many years, these tools were designed for and are ideal for expert users. FinVis is a novel visual analytics tool that is accessible to the casual user. FinVis can also be described as a casual information visualization tool as it provides a simple and powerful information visualization for the non-expert user. The goal of casual information visualization is to spread the benefits of information visualization and visualization to the general public [25]. One of the biggest challenges in making this transition is making the interface and visual metaphors simple and intuitive enough that the average user would require little training.

Some initial work in financial visualization centered on the choice between 2D [16] and 3D [22] representations of the data. While the 3D format allows for more data dimensions to be shown, it adds complexity and can make comparisons between data points difficult, especially for a casual user. The 2D examples also allow for three dimensions of data, but with one axis over-plotted with two dimensions' worth of data. This method works well when one dimension is relatively small. Because the risk dimension can be plotted in terms of wealth and is of much smaller magnitude than net wealth, FinVis opts for the 2D approach. More recently, FINESSE [30], LifeLines [24], Performance Matrix visualization [33], and an adaptation of the Bead system [4] have refined both 2D and 3D techniques to ease some of the drawbacks originally present. These specific solutions were designed for and are ideal for expert users. They use more complicated or denser representations of the data that would not be immediately understood

-
- Stephen Rudolph is with Purdue University, E-mail: srudolph@purdue.edu.
 - Anya Savikhin is with Purdue University, E-mail: anya@purdue.edu.
 - Dr. David Ebert is with Purdue University, E-mail: ebertd@purdue.edu.

Manuscript received 31 March 2008; accepted 1 August 2008; posted online 19 October 2008; mailed on 13 October 2008.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

by the general public.

2.1 Temporal Visualization

Of the many methods available for temporal visualization [1], only a few would make sense to show a casual user. Furthermore, personal financial data as a whole is not cyclic generally, which eliminates some of the remaining options [31, 6]. In most situations, wealth is accumulated over a certain time span and then depleted during retirement and FinVis makes no attempt to predict market cycles. Finally, we desire to show the individual contribution of each individual investment over time along with the aggregation of all investments, which brings us to ThemeRiverTM [13] as the most applicable temporal example.

ThemeRiverTM presents subject magnitude over time. FinVis parallels this; instead of visualizing theme magnitude, FinVis shows risk magnitude for financial decisions. Instead of branching from an axis, however, FinVis will follow the expected wealth curve. In this manner, a user may instantly see the impact of any decision. Another key distinction is that ThemeRiverTM does not include a visualization of uncertainty, which is necessary for FinVis to emphasize the probable outcomes.

2.2 Visualization of Risk

Visualization of financial risk can take many forms. FundExplorer [7], for example, visualizes the diversification of mutual funds. Generally, more diversified mutual funds carry less risk because they are less dependent on any one sector of the economy. In this sense, FundExplorer visualizes risk indirectly. FundExplorer is a useful personal finance tool for choosing carefully among very specific options in the present, but the scope of FinVis is much broader. FinVis is an exploratory tool that allows for short and long-term personal finance decision-making.

In the past, bar charts and confidence intervals have also been used for risk visualization [9]. These techniques would be recognizable to most users, but, by themselves, fail to work for our purposes. Both methods are used to primarily show relative risk between different project requirements. In FinVis, the user must be able to see the aggregate risk at any time, which is not possible with these paradigms.

To solve these issues, FinVis represents probability with partial transparency. The more likely a given outcome is, the more opaque it is. This will help focus the user on the most relevant outcomes. This technique has already been proposed to visualize expected financial wealth and risk simultaneously [27, 33]. As used previously, it can show the performance of one investment well, but does not scale in a useful manner. Combining these visual metaphors with those from ThemeRiverTM by overlaying investments with magnitudes derived from their contribution to the aggregate risk on top of the risk gradient provides a useful visualization basis for FinVis as seen in Figure 1 below. Risk is visualized using data of each investment's standard deviation and the portfolio's standard deviation, taking correlations and benefits from diversification into account.




Fig. 1. Aggregated risk gradient with investment risk contribution boundaries for two sample investments in blue.

3 RELATED WORK IN ECONOMICS AND FINANCE

Modern portfolio selection theory has been a significant topic for research since it was proposed by Markowitz over 50 years ago [23]. This theory suggests that rational investors should use diversification to optimize their portfolios by maximizing return for any given level of risk. Investors following this strategy should take into account risk and correlations in order to choose assets optimally. However, individuals have cognitive limitations and cannot effectively use this strategy for financial planning, and this has been discovered in the field [10, 12] and in laboratory experiments [20, 5]. Computerized decision support systems are needed to improve decision-making in this task [3].

3.1 Decision-Making under Uncertainty

The complexity of information encountered while making financial decisions may overwhelm the decision-maker, who does not have unlimited cognitive abilities and is therefore “boundedly rational” [11, 28]. Judgment and decision-making literature indicates that when available information is too complex and uncertain, individuals use heuristics or “rules of thumb” to make decisions [32, 18]. This kind of decision-making simplifies the cognitive process, but results in sub-optimal choices which are not in line with portfolio selection theory. Visual analytics reduces the complex cognitive work that is needed to perform decision-making tasks [29]. Savikhin et al. [26] recently found that individuals were able to make economic decisions which were significantly closer to the optimal while using an interactive visual analytics tool as compared to having this information in textual form. The implementation of FinVis in financial decision-making is a natural extension of this area of research. Economics experiments can be used to quantify the improvement in decision-making attributed to visualization. The use of laboratory experiments to understand decision-making is widespread in economics, and these have recently been implemented to quantify the value of visual analytics techniques for decision-making [26].

3.2 Overcoming Cognitive Limitations with FinVis

FinVis is a visual analytics tool that can help individuals overcome cognitive limitations and reduce the need for heuristics, improving decision-making. For example, when choosing between n financial instruments, individuals tend to follow the “framing and diversification” heuristic and allocate $1/n$ proportion of their wealth to each instrument, overlooking correlations [3, 12, 14]. We believe that FinVis reduces cognitive costs and eliminates the necessity of using this heuristic. Further, due to “ambiguity aversion”, investors tend to choose an inappropriate default portfolio without considering options [3]. Choosing one’s own portfolio helps the individual to identify with the portfolio and stick to an investment plan [10]. FinVis is an exploratory tool which mitigates these problems by allowing investors to quickly assess outcomes of possible portfolio options and choose an appropriate allocation - this allows individuals to move away from the default portfolio and successfully identify with their choice. Finally, some individuals tend to be overconfident compared to their abilities and after selecting a portfolio would prefer their own portfolio to that of an expert [10]. These “egocentric biases” can lead to serious misinvestments and financial crashes [10]. FinVis can improve the choices of investors by helping them to diversify, thus reducing negative financial impact of untrained and overconfident investors. Other investors are not confident enough and fail to develop any plan [21]. Graphical representations increase confidence and accuracy [8], so that through use of FinVis, overconfident individuals may improve abilities and uncertain individuals may become more confident. Using FinVis can help individuals increase financial knowledge, which is directly correlated with success in investment planning [21].

4 FINANCIAL MODEL IMPLEMENTATION

We use the standard portfolio theory approach to calculate expected returns, standard deviations, and correlations [17]. While the current FinVis program uses fictional investment choices with generic characteristics, the version of FinVis for dissemination to the public will allow users to insert information about investments available to them

in practice. Past returns and standard deviations will be used as the best predictors of future return and risk [17].

4.1 Liquid and Invested Assets

Net wealth can be broken up entirely into liquid assets and invested assets. When an investment is made with a minimum duration, the principal used to make that investment is no longer available for further investment during that period. By tracking this allocation of wealth, we are able to either prevent the user from going into debt, as FinVis currently does, or require the user to take out a loan to cover the debt.

When there is no risk associated with an investment, the principal is deducted from the liquid assets for the duration of the investment and subsequently, the principal and earnings are added back in to the available liquid assets. When risk is a factor, the earnings are no longer a single possible value. We could choose to be as optimistic or pessimistic about the return, but the most probable value will always be the expected earnings, which can be calculated using the interest formulas. This is the measure used in FinVis.

4.2 Interest

We use compound interest when calculating expected returns, using the following formula:

$$I = P(1 + R)^T \quad (1)$$

In this equation and subsequently in this paper, I is the interest or return on investment, P is the principal amount invested, R is the expected annual return rate, and T is the time elapsed in years since the initial investment was made. In the FinVis program, R is labeled as “APR” and T is labeled as “Length.”

4.3 Risk

FinVis displays both individual risk and aggregate risk of investment allocation decisions. Risk is shown as two standard deviations above and below the expected return of the investment, and returns are normally distributed. FinVis displays both individual risk (from each investment choice) and the aggregate risk (from the entire portfolio). This results in approximately a 95% confidence interval, which gives the user a very good idea of the most probable earnings. To account for risk, we modify equations 1 and 2 above by either adding or subtracting two standard deviations (2σ) to the return. The resulting formulas are:

$$I = P(R \pm 2\sigma)T \quad (2)$$

$$I = P(1 + R \pm 2\sigma)^T \quad (3)$$

The aggregate risk of n investments also depends on their correlations. For example, when investments are perfectly correlated, no reduction in risk results from diversification. When those investments are perfectly negatively correlated, one is the complement of the other, and all risk can be eliminated. In order to show the value of investment diversification, we must account for the change in risk due to correlations.

The aggregate risk is defined as:

$$\sigma_{\text{aggregate}} = \sqrt{\sum_{i=0}^n \sum_{j=0}^n w_i w_j \sigma_i \sigma_j c_{i,j}} \quad (4)$$

Where w_i represents the ratio of the i -th investment’s principal to the sum of all principals and $c_{i,j}$ is the correlation between the i -th and j -th investments as a value between -1 (perfectly negatively correlated) and 1 (perfectly correlated).

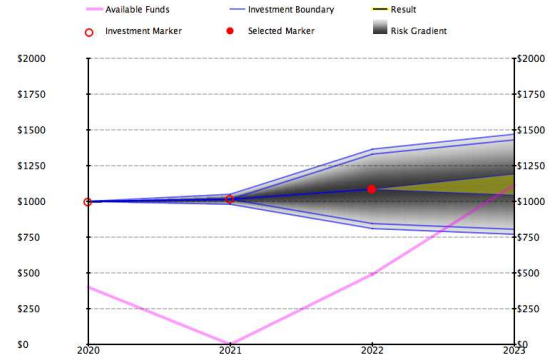
5 VISUAL & INTERACTIVE IMPLEMENTATION

5.1 Investment Display & Modification

Each investment is shown in two ways simultaneously, in the table of all investments (Figure 2(a)) and in the plot window (Figure 2(b)). In the table, the details of each investment are enumerated and each field

	Principal	Year ▼	Type	Length	σ	APR
1	\$600	2020	A	1	3	3
2	\$1018	2021	C	1	12	7
3	\$600	2022	B	1	6	5

(a) Table of Investments, σ Indicates the Standard Deviation of the APR (R).



(b) Visualization of investments listed above.

Fig. 2. Corresponding displays of three investments in three different years.

is sortable. This can allow the user to quickly see which investments have the lowest risk or highest principal for example. A corresponding visualization marks each investment as either a ring when unselected or disc when selected. Each marker is placed at the start year of each investment and at the expected net wealth value for that year.

Once an investment has been selected from the investments table, its properties are shown in the modification panel. Figure 3 shows the panel for the investments in Figure 2 where the selected investment is of Type B. The investment option, principal, and start year are all editable and may be altered as many times as the user wishes. Due to the constraints imposed by the user study, the remaining properties are derived from those three. Changes to investment properties immediately update the visualization to allow for quick exploration.

There are two buttons available between the investment list and modification panel, labeled as “+” and “-”. The former allows the user to add a new investment, which is immediately selected. The latter removes the currently selected investment and selects the next one, if there are any left.

5.2 Visualizing Risk

Risk is shown in two distinct ways. First, the aggregate risk (within two standard deviations) is displayed as a Gaussian gradient where darker parts of the gradient represent more probable outcomes. Because exactly two standard deviations from the mean are always displayed, the range of intensity values can be precomputed and converted to a linear scale.

The risk is also displayed on a per-investment basis. This is done so that the user can ascertain each investment’s contribution to the aggregate risk. Thus, the aggregate risk is divided by linear blue borders as seen in Figure 2(b). Subsequent investments have their risk contribution displayed closer to the center; the most recent contribution is in the middle and the least recent has its optimistic range at the top and its pessimistic on the bottom. Here, optimistic refers to the positive two standard deviations’ worth of risk, and pessimistic the negative.

Unable to find prior art for dividing the aggregate risk, we provide a formula that satisfies two necessary conditions. The sum of all individual risk must equal the aggregate risk and the proportions must

Modify Selected Investment

Investment Length:

Investment Option:

Principal:

Start Year:

End Year:

Interest %:

Standard Deviation %:

Fig. 3. Modification Panel

account for the principal. Larger investments must take up a correspondingly greater share of the aggregate risk. This is the formula we derived to satisfy both of these conditions:

$$\hat{\sigma}_i = \frac{\sqrt{(\sigma_i / \sum_{j=1}^n \sigma_j) * w_i}}{\sigma_{aggregate}} \quad (5)$$

Where σ_i is the unadjusted standard deviation of the investment we are calculating the individual risk for and $\hat{\sigma}_i$ is the effective standard deviation accounting for covariance between investments made in the same year.

5.3 Showing Context

To show the amount remaining to invest, the liquid assets, a partially transparent purple line is plotted on top of the chart. This can be seen in Figure 2(b) for the investment choices listed in Figure 2(a).

In addition to the currently selected investment being highlighted in the table of investments (Figure 2(a)), it is also shown as a yellow highlight overlaid on top of that investment's risk contribution (Figure 2(b)). This allows the user to quickly see the changes as they modify the selected investment.

5.4 Baseline Version

A baseline tabular version of the program was made to be used in the laboratory experiment. The baseline version presents most of the same information, including aggregate risk, expected wealth, and liquid assets in sortable columns with one row per decision year (Figure 4(b)). The only difference between the baseline program and FinVis is that the former also includes visual information of context, risk gradient, and risk boundaries for each individual investment. Both versions have identical controls, capabilities, and investment list displays. A comparison between versions is shown in Figure 4.

6 EXPERIMENTAL DESIGN

6.1 Anticipated Findings

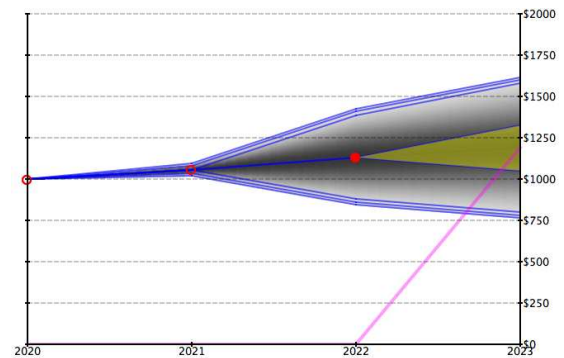
We conjecture that FinVis can help decision-makers in a number of ways, and we directly test each of our predictions using a laboratory experiment. Our laboratory experiment consists of a treatment with subjects making decisions using the FinVis visual analytics application, and a separate baseline treatment with subjects making decisions using an application with the same information in tabular form. First,

	Principal	Year ▼	Type	Length	σ	APR
1	\$500	2020	B	1	6	5
2	\$500	2020	C	1	12	7
3	\$1060	2021	C	1	12	7
4	\$1134	2022	B	1	6	5

(a) Investment Table

Year ▼	Expected Net Wealth	-2 σ	+2 σ	Funds Available
1 2020	\$1000	\$1000	\$1000	\$0
2 2021	\$1060	\$1021	\$1099	\$0
3 2022	\$1134	\$841	\$1427	\$0
4 2023	\$1190	\$761	\$1619	\$1190

(b) Baseline Version



(c) Full FinVis Version

Fig. 4. Both treatments showing the investments listed above.

we posit that FinVis is a useful visual analytics application that can improve decision-making by helping subjects choose optimal allocations which maximize return and minimize risk. Second, while both FinVis and the baseline application allow the same amount of exploration, we anticipate that FinVis will engage subjects and subjects using FinVis will choose to explore the dataset more than their counterparts using the baseline application. This should lead to greater learning over time in the FinVis treatment as compared to the baseline. Third, we predict that FinVis will increase the confidence of participants in decision-making in the experiment and also in practical application.

6.2 Experimental Setup

Inexperienced undergraduate students from Purdue University volunteered to participate in the study - 27 subjects in total participated. Experiments were conducted at the Vernon Smith Experimental Economics Lab using standard laboratory protocols. Subjects earned between \$10 and \$25 in the experiment, where their earnings depended on their choice of investment portfolio. Subjects whose portfolio choices resulted in higher experiment earnings earned a greater amount in the experiment as compared to subjects whose portfolio choices resulted in lower experiment earnings; thus, there was a high incentive to make good decisions which maximized return and minimized risk.

We used a between subjects design, with 13 subjects participating in the treatment using FinVis (FV) and 13 different subjects participating in the baseline treatment (B) using a tabular display with the same information. The information and training was the same in both treatments. Self-reported academic majors of subjects included business, engineering, information technology, science, and health sciences. It is possible that business majors would be likely to perform better in this experiment and that this could account for the improved performance in FV; however, over half of participants in B were business majors, while less than one third of participants in FV were business majors. About 60% of participants were male and 40% were female, with the same proportion in each treatment. The average age of participants was 21, and most participants were between age 18 and 23, with the same approximate ages in both treatments. Note that none of the demographic characteristics had a statistically significant effect on subjects' decisions.

As in previous work [20], we provided subjects with four generic investments, A, B, C, and D with returns and standard deviations of 3%, 5%, 7%, 2% and 3%, 6%, 12%, 0% respectively. The correlation between B and C was -0.8, while all other correlations between investments were 0. Each subject had 3 years (2020, 2021, and 2022) in which to make allocation decisions. We chose years far in the future in an attempt to keep the current financial situation from influencing the controlled laboratory experiment. Three decision years were used because this gives subjects the opportunity to plan for the future. While we could have allowed decisions further into the future in order to make the experiment more akin to a retirement planning task, we feel that this would have complicated the experiment without providing further information about decision-making. In order to allow for learning, the decision task was repeated for 6 periods; thus, subjects made decisions in 3 decision years for a total of 6 periods. After each period, an outcome was randomly selected according to the resulting Gaussian distribution at the end of the final decision year and reported to the subject. At the end of the experiment, random numbers were drawn to determine which period was to be paid. Only 2 periods were paid, which is a common practice in economic experiments. Subjects received a \$1,000 experimental dollar endowment at the start of each period, and were paid for the profit earned after returning the initial endowment to the experimenter. Experimental dollars were converted to U.S. dollars at the rate of \$100 experimental dollars equals \$1 U.S. dollar.

Subjects could make any number of decisions for each year when choosing to allocate the \$1,000 experimental dollar endowment. For example, subjects could either make one decision in each year and allocate the full endowment to one of the four investments, or subjects could make multiple decisions in each year, allocating fractions of the

endowment to any or all of the investment choices. Subjects were told that they could add, remove, or modify investment decisions as many times as they desired, and subjects moved through the experiment at their own pace. The program recorded both the final choices made, the time taken for each period, and the number of times investments were added, removed, or edited. These latter measurements can help us understand to what extent exploration of the dataset occurred.

6.3 Training, Risk and Confidence Elicitation

Before the beginning of the experiment, subjects also received a short training on reading the information screen presented by FinVis. Subjects participating in the baseline treatment received a short training on reading the tabular information screen. The training took approximately 30 minutes for both treatments, and included descriptions of the components of the tool and a tutorial.

A simple quiz was administered in which subjects were given a static information screen similar to the one presented in each treatment and asked to answer 3 questions about the screen to gauge subjects' basic understanding. This quiz was not difficult and did not require exploration of data. Subjects who did not answer all questions correctly were asked to leave and did not participate in the experiment. This minimal requirement ensured that only subjects who paid attention to instructions and knew how to read a table or a graph participated in the experiment. 14 subjects participated in the FinVis treatment, and 13/14 completed the quiz successfully. 13 subjects participated in the baseline treatment, and 13/13 completed the quiz successfully.

Prior to the experiment, we elicited subjects' objective risk preferences using a simple lottery task commonly used in economics experiments [15]. 15 lottery choices were presented and subjects were asked their preference for the safe option versus the gamble. In this task, subjects often choose the safe option in lotteries with a greater chance of loss and choose the gamble in lotteries with a lower chance of loss. Subjects who choose the safe option more often are considered risk averse, while subjects who choose the gamble more often are considered risk loving. At the end of the experiment, one of the lotteries was selected for payment - subjects earned between \$0 and \$3 in this task. We also asked a risk preference question commonly used in behavioral finance to elicit subjective risk tolerance - this question asked subjects to decide whether they would take a gamble to increase or decrease their income [2]. It is useful to compare risk preferences to decisions in the experiment in order to control for risk preferences when analyzing decisions.

At the end of every period before the outcome was known, we also asked subjects how confident they were that their decision was the best one. This is commonly elicited in the finance literature [8] and is important for quantifying the possible increase in confidence from using FinVis. This also helps us to gauge how confidence may change over time in either treatment. Finally, at the end of the experiment subjects were also asked to rate the usefulness of the information screen and the likelihood that they would use this screen to make decisions in practice. Subjects were also asked to report on their confidence in their overall decisions and in their decision-making abilities in financial planning. Confidence and usefulness were rated on an 8-point scale, while likelihood for future use and decision-making abilities were rated on a 3-point scale.

7 RESULTS

Experiment results confirmed our conjecture that FinVis improves decision-making - allocations in the FinVis treatment (FV) resulted in greater expected returns or lower risk as compared to allocations in the baseline treatment (B). Our second prediction was confirmed - we found that the FV group explored the dataset at almost double the rate of the B group, and that significant learning occurred over the 6 period experiment in the FV group but not in the B group. We found support for our third conjecture: post-experiment questionnaires showed that subjects in the FV group were more confident making decisions in the experiment. Also, 100% of the FV group reported that they were likely to use FinVis in practical application, as compared to only 60% of the B group. Note that subjects did not differ statistically significantly in

their risk preferences between treatments, and this was confirmed using an unpaired two sample t-test for both the objective and subjective risk elicitation tasks. This suggests that the two groups are comparable in this regard. Table (WE NEED TO PUT IN THE TABLE) shows the mean investment decisions of subjects for funds A, B, C and D in each decision year for each treatment.

7.1 Anticipated Finding 1: Improved Decision-Making

The optimal decision is one which maximizes the expected return at any level of risk, and this “efficient frontier” was calculated using the theoretical approach developed by Markowitz [23] and commonly used in previous work [20, 19]. Specifically, we solve the following minimization problem:

$$\text{Minimize } P^T P \quad (6)$$

subject to the constraints

$$P^T \varepsilon P = EP^T 1 = 1P^T \geq 0 \quad (7)$$

Where ε is the $n \times n$ variance-covariance matrix, P is a $n \times 1$ vector of investment expected returns, and E is expected return of the portfolio. The efficient frontier is created by varying E , so that each data point on the frontier corresponds to a vector of investment proportions. Figure 5 shows this theoretical benchmark overlaid with a scatter plot of all expected returns and risk properties of actual portfolio choices of subjects in each treatment. (NOTE: THIS IS GOING TO INCLUDE THEORETICAL BENCHMARK BY MON NIGHT). Note that subjects participating in the treatment with the FinVis tool (FV) outperformed subjects participating in the baseline treatment (B) in making decisions which were closer to the optimal. We calculated the expected return and standard deviation of the portfolio allocation choices for each period by calculating the expected return and standard deviation of the chosen portfolio in each year and then taking the weighted average across three years. We analyzed the expected return and risk of portfolio choices in FV versus B using an unpaired independent samples t-test with the null hypothesis that the means are equal, and the results are summarized in Table 1. We report the p-values for the alternative hypothesis of a predicted greater expected return for the FV group versus B group, and p-values for the null hypothesis of equality of riskiness of portfolio between groups. Accounting for learning in the first few periods, subjects in the FV treatment chose portfolio allocations with statistically significantly greater returns as compared to subjects in the B treatment. Further, risk of chosen allocations was not significantly different between treatments.

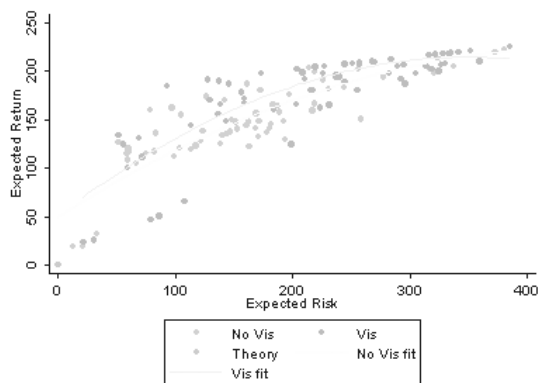


Fig. 5. Theoretical vs. Actual Return & Risk

We can test whether subjects are using the ‘framing and diversification’ heuristic, and allocating $1/n$ of their wealth to each instrument, overlooking risk and correlations. We ran several t-tests for each treatment to determine if the allocation decisions for A, B, C and D were equal. (THIS FINDING NEEDS TO BE PUT IN BY END OF MONDAY)

7.2 Anticipated Finding 2: Greater Exploration and Learning through Visual Analytics

We were also able to test the exploratory nature of FinVis by comparing the amount of time and exploration spent in the FV treatment to the amount of time and exploration spent in the B treatment. We recorded the number of times an investment was added, removed, or modified during each period. Modification of an investment included actions such as changing the type, principal, or year of the investment. In summary, investments were modified an average of 107.5 times per period in FV and 54.9 times per period in B. Investments were removed 6.1 times per period in FV and 3.7 times per period in B, and added 12.9 times per period in FV and 10.6 times per period in B. We used an independent sample t-test for each period and found that subjects in the FV treatment modified investments more than subjects in the B treatment, and this was statistically significant in almost all periods. Further, the FV group removed and added investments at a greater rate, and this was statistically significant in half of the periods. Table 2 summarizes the resulting p-values. It is unclear why results in period 6 reflect reduced exploration in FV, and more laboratory experiments are required to understand this result, but it is possible that subjects have explored sufficiently by this point. Note that while the FV group explored the decision space more, both groups ended with a similar number of investments at the end of each period (a mean of 6.8 investments in FV and 6.9 investments in B, with standard deviations of 3.2 and 3.3, respectively). Note also that in half of the periods, subjects in FV spent statistically significantly more time for each portfolio allocation decision, while in the remaining half of the periods subjects in FV did not spend a statistically significantly different amount of time for each decision as compared to subjects in B. On average, subjects in FV spent 6.9 minutes on each period, and subjects in B spent 4.7 minutes. These results suggest that the FinVis tool improves the amount of exploration, which can help reduce the “ambiguity aversion” bias.

We argue also that this increased exploration may be a contributing factor to increased learning in FV as compared to B. We estimated two random effects generalized least squares regressions for each treatment separately, with expected return and risk characteristics of each subject’s chosen portfolio as dependant variables and time trend dummies, demographic characteristics of subjects (the random effects), and self-reported confidence level as independent variables. We found that only in FV, the time trend dummy variables were statistically significant at the 5% level with positive coefficients in the regression with expected return as the dependent variable. The time trend dummy variables were not statistically significant in B. This suggests that the FV group learned to increase the expected return of their portfolio over time, but the B group did not make this shift. However, we also find that the time trend dummy variables were statistically significant at the 5% level with positive coefficients in the regression with risk as the dependent variable. This suggests that the FV group also began to choose riskier investments over time, while the B group did not experience this change. Further investigation is needed to understand the learning process, and the only conclusion that can be made is that spending more time with the program resulted in a statistically significant change in behavior for the FV group but not for the B group. The other independent variables included in the regressions did not have statistically significant effects on the dependent variables. The overall R-squared value for these regressions was approximately 0.30.

7.3 Anticipated Finding 3: Increased Confidence and Use in Practice

Confidence was elicited at the end of each period using an eight point scale from ‘least confident’ to ‘most confident’. Self reported confidence levels across the experiment were 3.6 for the B group, and 4.1 for the FV group. Subjects were also asked to report their confidence with their choices at the very end of the experiment: the average confidence level of the FV group was 4 while the average confidence level of the B group was 3.5. While t-tests do not find statistically significant differences, a power test suggests that increasing the number of subjects in each treatment would increase the significance of the results. The greater confidence of the RV group is an important result, because

Table 1. Summary of Expected Return/Risk t-test Results.

Period	Expected Return t-test	Direction of Difference	p-value	Expected Risk t-test	p-value
1	FV≠B	FV>B	0.44	FV≠B	0.68
2	FV≠B	FV>B	0.23	FV≠B	0.53
3	FV≠B	FV>B	0.14	FV≠B	0.37
4	FV≠B	FV>B	0.05*	FV≠B	0.38
5	FV≠B	FV>B	0.03*	FV≠B	0.61
6	FV≠B	FV>B	0.02*	FV≠B	0.67

Table 2. P-Values for the Hypothesis, FV > B.

Period	p-value for Number Modifications	p-value for Number Added	p-value for Number Removed
1	0.02*	0.40	0.28
2	0.01*	0.03*	0.02*
3	0.00*	0.03*	0.02*
4	0.01*	0.11	0.08
5	0.05*	0.17	0.14
6	0.40	0.72	0.76

individuals who are not confident would be less likely to be able to develop a financial plan in practical application. We asked subjects if after completing the experiment they were more, less, or equally confident in their abilities to understand financial planning. 85% of subjects in B reported that they were equally confident, while 15% reported that they were less confident. In contrast, 23% of subjects in FV were *more* confident, 62% were equally confident, and 15% were less confident. Thus, FinVis improved the confidence of some subjects. The greater confidence of the FV group is an important result, because individuals who are not confident would be less likely to be able to develop a financial plan in practical application.

We also asked subjects how likely they are to use a visual analytics tool such as the one presented to them to plan investment decisions, if one were made available. 38% of B subjects reported that they were “not very likely” to use the baseline tool, 46% of B subjects were “somewhat likely” to use such a tool, and only 15% of B subjects were “very likely” to use the tool. In contrast, 61% of FV subjects were “somewhat likely” and 39% of FV subjects were “very likely” to use the FinVis tool in practice. This result suggests that FinVis would be useful for the general public, as compared to a baseline version of this tool.

8 CONCLUSIONS AND FUTURE WORK

8.1 Conclusion

We developed a visual analytics tool which allows the non-expert casual user to make improved personal finance decisions through a visualization of return, risk and correlation aspects of financial data. A laboratory experiment was conducted and we found that subjects in the treatment using the FinVis tool outperformed subjects using a tool with the same information in tabular form, and this result was statistically significant. We also found that subjects were able to learn the optimal decision over time. Finally, FinVis helped subjects overcome decision biases that are common in economics. This is a significant result, suggesting that visual analytics lowers the cognitive cost and allows subjects to analyze decisions without resorting to suboptimal ‘rules of thumb’.

8.2 Financial Model Improvements

The existing financial model does not presently account for many possible factors like periodic trends, inflation, transactional costs, and several others. Adding these would make the scenarios more realistic.

It would also be useful to study how different risk measures, such as lower and upper probabilistic means and variances, can improve the effectiveness of FinVis.

8.3 Implementation Improvements

The ultimate aim of FinVis requires that the initial program be extended for general use. This would entail creating a system for easily adding any investment or divestment of any time for any duration. Focus plus context additions could allow for closer inspection while

maintaining a full picture. Additional interactivity, particularly in the area of investment selection, would make the program easier to use. Also, there are noticeable lags when the investment plot is refreshed on older machines; performance could be improved in many ways, possibly by implementing threading throughout the application.

Many academically-focused extensions of FinVis are also possible. If borrowing and lending were implemented, FinVis could be used in the lab to test whether the capital asset pricing model (CAPM) predictions hold. CAPM is generally not supported by previous experimental research [20, 19], but decision-making may improve and better support CAPM predictions when visual analytics tools are utilized.

ACKNOWLEDGEMENTS

The authors wish to thank Roman Sheremeta and Niklas Elmquist for their help and feedback. This work has been funded by the US Department of Homeland Security Regional Visualization and Analytics Center (RVAC) Center of Excellence and the US National Science Foundation (NSF) under Grants 0328984 and 0121288.

REFERENCES

- [1] W. Aigner, S. Miksch, W. Muller, H. Schumann, and C. Tominski. Visual methods for analyzing time-oriented data. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):47–60, Jan.-Feb. 2008.
- [2] R. Barsky, F. Juster, M. Kimball, and M. Shapiro. Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study. *Quarterly Journal of Economics*, 112:537–579, 1997.
- [3] G. Bhandari, K. Hassanein, and R. Deaves. Debiasing investors with decision support systems: An experimental investigation. *Decision Support Systems*, 46:399–410, 2008.
- [4] D. Brodbeck, M. Chalmers, A. Lunzer, and P. Cotture. Domesticating bead: adapting an information visualization system to a financial institution. *Proceedings, IEEE Symposium on Information Visualization, 1997.*, pages 73–80, Oct. 1997.
- [5] C. Cadsby and E. Maynes. Laboratory experiments in corporate and investment finance: a survey. *Managerial and Decision Economics*, 19:277–298, 1998.
- [6] J. V. Carlis and J. A. Konstan. Interactive visualization of serial periodic data. In *UIST '98: Proceedings of the 11th annual ACM symposium on User interface software and technology*, pages 29–38, New York, NY, USA, 1998. ACM.
- [7] C. Csallner, M. Handte, O. Lehmann, and J. Stasko. Fundexplorer: Supporting the diversification of mutual fund portfolios using context treemaps. *IEEE Symposium on Information Visualization, 2003. InfoVis 2003.*, 00:26, 2003.
- [8] G. DeSanctis and S. Jarvenpaa. Graphical presentation of accounting data for financial forecasting: An experimental investigation. *Accounting, Organizations and Society*, 14:509–525, 1989.
- [9] M. Feather, S. Cornford, J. Kiper, and T. Menzies. Experiences using visualization techniques to present requirements, risks to them, and options for risk mitigation. *First International Workshop on Requirements Engineering Visualization, 2006. REV '06.*, pages 10–10, Sept. 2006.

- [10] G. Fellner, W. Guth, and B. Maciejovsky. Illusion of expertise in portfolio decisions: an experimental approach. *Journal of Economic Behavior and Organization*, 55:355–376, 2004.
- [11] G. Gigerenzer and D. Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103:650–669, 1996.
- [12] W. Goetzmann and A. Kumar. Equity portfolio diversification. *Review of Finance*, 2008.
- [13] S. Havre, B. Hetzler, and L. Nowell. Themeriver: visualizing theme changes over time. *IEEE Symposium on Information Visualization, 2000. InfoVis 2000.*, pages 115–123, 2000.
- [14] T. M. Hedesström. *The psychology of diversification: Novice investors' ability to spread risks*. Department of Psychology, Göteborg University, 2006.
- [15] C. Holt and S. Laury. Risk aversion and incentive effects: New data without order effects. *American Economic Review*, pages 902–904, 2005.
- [16] I. Jarett. Using charts to present international trade data. *IEEE Computer Graphics and Applications*, 3(9):55–61, Dec. 1983.
- [17] C. P. Jones. *Investments: Analysis and Management*. John Wiley and Sons, Inc., 9 edition, 2004.
- [18] D. Kahneman and M. Riepe. Aspects of investor psychology. *Journal of Portfolio Management*, 24:52–65, 1998.
- [19] Y. Kroll and H. Levy. Further tests of the separation theorem and the capital asset pricing model. *American Economic Review*, 82:664–670, 1992.
- [20] Y. Kroll, H. Levy, and A. Rapoport. Experimental tests of the separation theorem and the capital asset pricing model. *American Economic Review*, 78:500–519, 1988.
- [21] A. Lusardi and O. S. Mitchell. Financial literacy and planning: Implications for retirement wellbeing. *Michigan Retirement Research Center Research*, (WP 2005-108), 2005.
- [22] M. Lux. Visualization of financial information. In *NPIV '97: Proceedings of the 1997 workshop on New paradigms in information visualization and manipulation*, pages 58–61, New York, NY, USA, 1997. ACM.
- [23] H. Markowitz. Portfolio selection. *Journal of Finance*, pages 77–91, 1952.
- [24] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman. Life-lines: visualizing personal histories. In *CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 221–ff., New York, NY, USA, 1996. ACM.
- [25] Z. Pousman, J. Stasko, and M. Mateas. Casual information visualization: Depictions of data in everyday life. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1145–1152, Nov.-Dec. 2007.
- [26] A. Savikhin, R. Maciejewski, and D. Ebert. Applied visual analytics for economic decision making. *Proceedings of the Symposium on Visual Analytics Science and Technology*, pages 107–114, 2008.
- [27] T. Schreck, T. Tekušová, J. Kohlhammer, and D. Fellner. Trajectory-based visual analysis of large financial time series data. *SIGKDD Explor. Newsl.*, 9(2):30–37, 2007.
- [28] H. Simon. *Models of Bounded Rationality*. MIT Press, 1982.
- [29] J. Thomas and K. Cook. *Illuminating the Path: the research and development agenda for visual analytics*. IEEE, 2005.
- [30] A. Varshney and A. Kaufman. Finesse: a financial information spreadsheet. *Proceedings IEEE Symposium on Information Visualization '96*, pages 70–71, 125, Oct 1996.
- [31] M. Weber, M. Alexa, and W. Müller. Visualizing time-series on spirals. *IEEE Symposium on Information Visualization, 2001. INFOVIS 2001.*, pages 7–13, 2001.
- [32] W. Wright. Cognitive information processing biases: Implications for producers and users of financial information. *Decision Sciences*, 11:284–298, 1980.
- [33] H. Ziegler, T. Nietzsche, and D. Keim. Visual analytics on the financial market: Pixel-based analysis and comparison of long-term investments. *12th International Conference Information Visualisation, 2008. IV '08.*, pages 287–295, July 2008.