Can Prospect Theory Explain Market Calendar Effects?

Mayo, Robert

George Mason University

12 December 2013

Can Prospect Theory Explain Market Calendar Effects?

Robert L. Mayo

George Mason University

December 12, 2013

"October. This is one of the peculiarly dangerous months to speculate in stocks. The others are July, January, September, April, November, May, March, June, December, August, and February."

Mark Twain's Pudd'nhead Wilson (1894)

1. Introduction

The purpose of this project is to determine if calendar effects observed in stock markets can be explained by prospect theory. In order to answer this question, I have created an agent based model simulating a stock market.

One of the difficulties that separate economics from the physical sciences is the fact that the subject of our field of study is the product of the choices of human beings. The famous paper by Stigler and Becker, *De Gustibus Non Est Disputandum* [1977] attempts to show that tastes, what are commonly referred to in economics as preferences, are in fact amenable to economic analysis. Although there have been great strides made over the course of the last generation in the understanding of the forces that generate economic decisions, behavioral economics and behavioral finance are still unable to explain many of the departures we see in the real world from what would be predicted by economic models.

One possible source of these anomalies, although by no means the only one, is the fact that while traditional macroeconomics studies the dynamics of groups, i.e. an entire economy, and traditional microeconomics studies the incentives and responses of individuals, either persons or firms, there is no primary sub-discipline devoted to the study of the emergent macro effects of individual micro actions. While it is possible (and frequently done) to simply aggregate the micro to produce a model of the macro, this procedure fails to encompass any aspect of the system as more than the sum of its parts.

Due to the steady advance in computer processing speed and the commensurate reduction in computational cost, agent-based computational simulation provides an opportunity to investigate these emergent phenomena. Calendar effects in stock markets are one of the anomalies that has thus far eluded explanation by traditional economic modeling but may be understandable as an emergent product of interactions between market participants.

1.1 Calendar Effects

Traditional stock market lore holds that there are significant correlations between stock returns and the calendar. Some of the most well known are: The January effect - that returns are abnormally high in the month of January. The Halloween effect - that the market tends to perform better in the six months ending October 31st than in the other half of the year. The weekend effect - that market returns tend to be lower on Monday than on Friday. There is not unanimity in the financial literature about whether these effects are real or myth, let alone what their cause might be if they do exist. A number of studies have found calendar effects of various types in the market histories of many countries and across many decades. Hansen and Lunde [2003] in a study of market indices from Denmark, France, Germany, Hong Kong, Italy, Japan, Norway, Sweden, the UK, and the United States find significant evidence of the January effect and the weekend effect. Similarly, Swinkels, Laurens, et al. [2012] find significant evidence of end-of-year and end-of-month effects. On the other hand, Sullivan, Timmermann, et al. [2001] conclude that evidence of calendar effects is a statistical anomaly produced by overly aggressive data mining.

On a theoretical basis, a January effect could be supported since tax liabilities are frequently calculated on a calendar year rather than a fiscal year. No explanation, however, has been forthcoming about how the other calendar effects discussed above could exist for any significant length of time. The efficient market hypothesis predicts that any perceptible pattern in returns not counterbalanced by another factor external to the metric such as taxation, should be arbitraged away.

1.2 Prospect Theory

In the 1970's, Daniel Kahneman and Amos Tversky began to investigate the Allais paradox.

This is the observation made by French economist Maurice Allais that people will routinely make different choices when presented with mathematically identical gambles depending on how they are presented. The foundational paper of Kahneman and Tversky's work, *Prospect theory: An analysis of decision under risk* [1979] began to model this phenomenon. They characterize it as a process of decision making by heuristics. People judge uncertain outcomes based on decision rules that are tractable, sacrificing precision for reliability. The commonly used rule of thumb is to seek gains and avoid losses. This simple rule, however, while quite reliable sacrifices accuracy in two ways: First, it ranks outcomes lexiographically, giving an absolute preference for any perceived gain over any perceived loss without factoring the associated probabilities of those outcomes. Second, it defines the outcomes as relative to the current position. That is, it consideres an outcome as profit or loss compared to current wealth, rather than on an absolute scale.

An example of this phenomenon in action is as follows: If you offer a subject a choice between a certainty of receiving $100 or a 50% chance of winning $210, they will more likely than not choose the certain $100, even though the expected value of the second choice is $105. This is consistent with the subject being risk averse. They are willing to 'pay' the loss of a potential $5 to 'buy' the avoidance of the uncertainty in the second choice. However, if the same subject is offered a certainty of a $100 loss or an 50% probability of a $210 loss, they will be more likely to choose the second alternative than the first. The first choice has an expected value of - $100 and the second has an expected value of - $105. Given that the subject was the same individual with the same risk preference, he should be willing to again 'pay' the same amount of $5 to avoid the same amount of uncertainty of the second choice. And yet, that expectation is consistently violated in the real world.

Kahneman and Tversky proposed that the key to understanding this effect was the fact that whether the prospect from the point of view of the subject was a gain or a loss was not included in the calculation of the expected value of each choice. In essence, people first ranked the outcomes by whether they were a positive or negative deviation from present wealth, and only after that considered the expected value of the outcome. In practice, this results in differing preferences for risk for gains and losses. People want to win, therefore they tend to take the gamble that maximizes the probability of winning. Similarly, people want to avoid losing so they will tend to take gambles that reduce the probability of a loss. The resultant differeing risk profiles based on deviations from current wealth can be seen in the asymetrical pattern shown in figure 1[1].

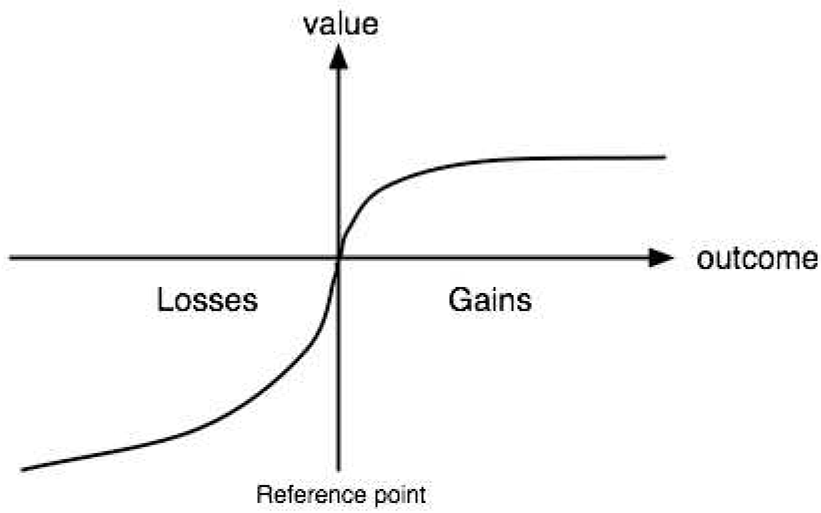

1.3 Hypothesis

The possible relevance of prospect theory to stock market calendar effects is in their respective temporality. Calendar effects are by definition temporal phenomena; changes in effective risk preference due to imperfect heuristics are not. However, if we go further and view the reference point of wealth as changing over time, then non-optimizing risk preferences could acquire a temporal component as well. When weighing two uncertain alternatives, if the subject must incorporate his current wealth into the decision function, a choice must be made of 'wealth as of when?' A reference point in time must be included as part of the definition of the reference point in wealth.

When would the temporal reference point most likely be? If someone has to define their wealth as of a relatively recent date in order to make a current economic decision, how do they choose among

[1] "Value function in Prospect Theory", drawing by Marc Oliver Rieger. Released under the GNU Free Documentation License.

an infinite set of possibilities? This is a version of the "where to meet in New York problem." If two people wanted to meet in New York City at a particular time but had not agreed on a location, where would they go? They would fall back on common cultural geographic reference points and choose Times Square or perhaps the Empire State building. In the same way, someone who needs to choose a recent point in time, absent any other considerations, will tend to fall back on common cultural temporal reference points and choose the start of the day, the start of the week, or the start of the month.

On this basis, I propose that the interaction of commonly held time reference points input into prospect theoretic decision functions could generate highly correlated, boundedly rational economic decisions in large groups. My hypothesis is that calendar effects are caused by the inconsistent risk preferences described by prospect theory.

2. The Model

2.1 Overview

This agent based stock market simulation is at core an interaction between traders and stocks. As in a real world stock market, this market begins with human beings making choices in a complex and uncertain environment. A variable number of traders use one of ten different strategies to evaluate each stock based on current price, current dividend, price history, dividend history, cash on hand, and interest rate to decide how to rebalance their portfolio by issuing an order to buy or sell a number of shares. Real world stock markets typically allow both market orders and limit orders. A market order is a request from a trader to buy or sell a specified number of shares at the best price immediately available. A limit order is a request to buy or sell a specific number of shares when possible at a specific price or better. Market orders sacrifice price predictability for guaranteed quantity execution, while limit orders sacrifice quantity predictability for a guaranteed price. Since market orders always execute and limit orders do not, market orders increase liquidity while limit orders reduce it. Faced with this difference, I chose to implement market orders only. Including limit orders in the simulation would have required me to also simulate a class of market makers (entities that provide market liquidity by buying or selling in unlimited amounts) or allow trades to not clear, which would be equivalent to allowing the market to not achieve price equilibrium.

After researching artificial market designs, [LeBaron, 2001], [Benhammadaa and Chikhib, 2012], [Samanidou1, Zschischang, et al., 2007], it became apparent that a central issue of agent-based market design is the choice of structure for order execution and price discovery. Two basic approaches are common: order book and Walrasian auctioneer. In an order book system traders place limit orders which include price boundaries for their desired trades. These are entered into a queue called (unsurprisingly) an order book. Buy orders and sell orders with overlapping quantities and price boundaries are executed, usually at the price of whichever order was placed first. These overlapping limit orders function indistinguishably from market orders. Limit orders that do not overlap are not immediately executed and are stored in the order book. An example may be helpful at this point. Assume a market with 4 traders, each placing a limit order for a given stock with a prior price of $15 per share. The orders are as seen in table 1.

Table 1. Example Limit Orders

Trader number	Trade direction	Quantity	Limit price
1	Buy	100	$15
2	Sell	150	$14
3	Sell	75	$15
4	Buy	75	$16

Trader 4 wants to buy 75 shares and is willing to pay up to $16 per share. Trader 3 wants to sell 75 shares and is willing to take accept as low as $15 per share. These trades overlap in both price and quantity, so they can be matched and executed. Assuming the orders were entered into the order book in order of trader number, then trader 3's order was placed before trader 4's order and therefore the trade will execute at trader 3's limit price of $15 and are both removed from the order book. Traders 1 and 2 have orders that overlap in price (trader 1 offers $15, trader 2 will accept $14) but they are not fully overlapping in quantity. Trader 2 can sell 100 shares to trader1 at $15. Trader 1's order has then been executed in full and it is removed from the book, but trader 2's order has only been partially filled. Trader 2 has sold 100 shares leaving an open order in the book of an offer to sell 50 shares at $14.

Since the outstanding order is to sell, inducing other traders to alter their orders to increase the quantity demanded to match requires the price to decrease. In other words, the previous price of $15 is now too high. The price adjustment equation used in the first version of the Santa Fe Institute artificial stock market was:

$$p_{(t+1)} = \alpha p_t \sum_{i=1}^{n} (B_i - S_i) \tag{1}$$

Where p is price, t is time, B is shared ordered to be bought, S is shared ordered to be sold, and i increments over n traders. α is the amount to adjust the price for a given quantity of unfilled net orders.

As described in *"Building the Santa Fe Artificial Stock Market"* [LeBaron, 2002], this procedure caused serious problems. The market showed unacceptably high sensitivity to the value of α. If set slightly too low, the market would be unable to adjust quickly enough to approach equilibrium and the number of unfilled orders would accumulate. If set a bit too high, market volatility would explode causing wild swings between unexecuted buy orders and sell orders. Preliminary versions of my simulation used this system and exhibited the same behavior.

To alleviate this problem, I redesigned the price discovery mechanism to pattern a Walrasian auctioneer. The Walrasian auctioneer is a simplified economic model of the process equilibrating quantity supplied and demanded. Assume an omniscient auctioneer. At every point in time this auctioneer will determine the supply and demand functions of all current and potential market participants and will adjust the market price to equate quantity offered for sale with quantity desired to be purchased, thus clearing the market. Although this mechanism is simpler than an order book system in theory, it proved intractably more complex in implementation. I resolved this dilemma by designing a hybrid between the two systems.

I created a price discovery system with nearly the simplicity of an order book, but with the advantage of a Walrasian auctioneer in adjusting price prior to executing trade orders. In this system, an auctioneer proposes a price and queries the traders for preliminary orders. If the quantity supplied does not equal quantity demanded, the preliminary orders are discarded and the price is altered a fixed amount in the direction of equilibrium indicated by the sign of the net order quantity. This new candidate price is then presented to the traders and they are again asked to submit preliminary orders. This process iterates until a price is found that returns a zero net demand order set. This process has the dual advantages of not requiring the auctioneer to know the supply and demand functions of the traders while simultaneously eliminating α from equation 1. This solution proved to both stable and acceptably fast running on the available less than state of the art hardware.

With the infrastructure of a reliable market clearing mechanism built, broadening the heterogeneity of the traders was a relatively simple task. I sequentially added multiple types of traders, each utilizing one basic trading strategy. The first trader variant are called "technical traders", "chartists", or sometimes "noise traders". This type of trader uses a strategy predicated on the assumption that markets follow discernible repeating patterns. I implemented this in its simplest form; these agents viewed the recent history of the price of the stock and assumed the trend direction would continue. If the price has been going up, it will continue to go up, therefore buy. If the price has been going down, it will continue to go down, therefore sell. The second trading strategy is called "fundamental" trading. Traders using this strategy ignore the patterns of price movement of stock and instead focus on the quality of the firm the stock represents. In this model, the "fundamental" quality of the firm is known from its earnings expressed as profits paid out in dividends. This definition has the great advantage of taking the staggering complexity of judging a corporation's economic strength and reducing it to a single number. The third type of trader added was the "contrarian". Contrarian trading is simply to do the opposite of what other traders are seen to do.

After creating these three types of traders, I multiplied them by three levels of aggressiveness yielding nine trading strategies: technical passive, technical mid, technical aggressive, contrarian passive, contrarian mid, contrarian aggressive, fundamental passive, fundamental mid, and fundamental aggressive. The degree of aggression determined what percentage of a trader's cash or portfolio to commit to each trade, from 5% for passive traders up to 25% for aggressive traders.

During the presentation of this model to the class, Dr. Tian/you raised a question that I interpreted as a concern that the complexity and multitude of trading strategies might end up hard coding in behaviors that could not be distinguished from emergent features of the market. After additional research, I discovered that your suggestion of zero intelligence traders was a both reasonable and potentially interesting option. On that basis, I added a second independent system of traders composed exclusively of zero intelligence agents. One of the two parallel populations of traders, either zero intelligence or optimizing, can be chosen to be used for each run while the other is disabled. This option is available through a simple True/False parameter setting[2].

2.2 Design details

Model variables are listed in table 1.

Table 1. Model variables

Variable	Type	Description
NUMBER_OF_RUNS	INTEGER $[1,\infty)$	Number of times the simulation will be repeated
NUMBER_OF_STOCKS	INTEGER $[1,\infty)$	Number of stocks in each market
NUMBER_OF_TRADERS	INTEGER $[1,\infty)$	Number of traders in each market
NUMBER_OF_DAYS	INTEGER $[1,\infty)$	Number of days of trading to be simulated

[2] As an aside, I discovered that the market dynamics were every bit as complex and interesting using zero intelligence agents as using complicated optimizers, and system performance seemed more stable as an added benefit. In short, you were right.

CASH	FLOAT (0,∞)	Starting cash balance of each trader
ZERO_INTEL	BOOLEAN [T/F]	Use zero intelligence traders instead of optimizing agents
PROSPECT	BOOLEAN [T/F]	Include prospect irrationality in trades
PROSPECT_RATE	FLOAT (0,∞)	Magnitude of prospect irrationality
AVG_LOOKBACK_DAYS	INTEGER [1,∞)	Number of days traders remember
INTEREST_RATE	FLOAT (0,∞)	Rate of return on cash per day
MARKET_TREND	FLOAT (0,∞)	Change in average stock price per day
HISTORY	INTEGER [1,∞)	Number of days of initial market history
wealth	FLOAT (0,∞)	Cash plus value of stocks owned for each trader
price	FLOAT (0,∞)	Current stock price
dividend	FLOAT (0,∞)	Current stock dividend
marketindex	FLOAT (0,∞)	Average price of all stocks in a market
quantity	INTEGER [0,∞)	Number of shares to be bought or sold by a trader

2.1.1 Model parameters

The parameters in this model are listed in all capital letters both in table 1 and in the model code. For the simulation to run reliably, there must be more than a small handful of traders. As the number is reduced below roughly a dozen, the market will increasingly exhibit erratic behavior. This is not a flaw in the model, but rather an accurate reflection of the reality of real work market function. In the extreme case of only two traders, their supply and demand schedules may not make any mutually beneficial trades possible. A large number of traders essentially guarantees that a sufficient number of trades will occur for meaningful results to be generated.

Interest rate determines the desirability of holding cash for a given market environment. As interest rate is reduced, traders using all types of optimizing strategies will tend to shift resources out of cash and into stocks causing the market index to rise. Conversely, a very high rate of interest will tend to shift resources out of the stock market to be held as interest bearing cash. If sufficiently high, the advantageous return on cash will cause the market index to approach zero.

Market trend is included as an optional setting to simulate trader behavior within an artificially induced bull or bear market. It was always set to zero for this experiment, but may be employed in future research.

History is calculated as two times average lookback days plus one. It is included because all optimizing traders take past stock and/or dividend behavior into account when deciding on a stock order. A synthetic history must be created and available for traders to reference on the first day of trading or their decision functions would fail.

2.1.2 Agents

There are four classes in the model: traders, stocks, markets, and simulation.

2.1.2.1 Traders

As discussed previously in section two, traders may be chosen to act as zero intelligence agents, with their trading behavior limited only by a liquidity constraint, or by setting ZERO_INTEL to false traders will utilize various optimizing trading strategies. Optimizing traders have a memory of their prior trades, portfolio, cash, and calculated wealth extending back on average the number of days specified in AVG_LOOKBACK_DAYS. A random deviation from this number is set for each trader at its creation. Traders also have memory for the same length of time of the price and dividend history for any stock they are considering trading. After a trader creates a trade order, it will be modified by a prospect function if PROSPECT is set to True. This is the same for both zero intelligence and optimizing agents.

2.1.2.2 Stocks

Stocks exist in this simulation as passive recipients of actions by traders and by the effects of both market wide and stock specific news. Stocks have a price, dividend, and unlimited history of the same. Although an unlimited number of stocks can be used, only one was created for the purposes of this experiment. This made a binary choice for traders between a generic market index and cash.

2.1.2.3 Market

The simulation can accommodate multiple markets running simultaneously but this feature is not used in this experiment. Only one market is used. Within the market class are functions to simulate the effect of exogenous market relevant news entering the system. The first of these randomly generates news effecting the entire market. This effect may be positive or negative and is implemented such that the average stock will move in a particular direction by a particular amount, but there will be variation between stocks with a small number moving in the opposite direction as the market as a whole. The second function randomly introduces news relevant to a single stock only. This is introduced as an alteration in the profitability of the underlying company expressed as a change in the dividend.

2.1.2.4 Simulation

The single instance of the simulation class creates and runs the model. It includes functions for price discovery, trade execution, account reconciliation, and output. Output takes three forms. If the NUMBER_OF_RUNS parameter is set to one, a graph will be created showing the level of the market index over the course of the run. A comma delimited TXT file will be created containing day number and index price for each day of the run. If the NUMBER_OF_RUNS parameter is set to two or more, the graph is suppressed and a comma delimited TXT file is created containing run number and final index price for each run.

3. Testing

3.1 Verification

Various diagnostic runs of the model were conducted printing each variable in turn. Only one major problem was found. I observed that changing the interest rate was not having the expected effect on stock prices. When I increased the interest rate, stock prices went up instead of down. If cash was returning a high yield, it should draw resources away from stocks and make the price decline. After substantial investigation, I realized that the money paid to traders as interest on their cash holdings was being created de novo and this additional money was being used to purchase more stock, even if at a lower rate. In essence, I had accidentally simulated the inflationary fiat money creation of an irresponsible monetary authority. To remedy the situation, I had a choice to fully simulate and control a central bank, or stop injecting money into the system. Although I hope to introduce a central bank into the simulation at a later date, for this experiment I prevented money creation by not crediting interest payments to traders accounts.

3.2 Sensitivity

Sensitivity of the model was tested with respect to four variables: interest rate, average look back days, number of traders, and prospect rate. The model was run 100 times with all four variables set at their reference values. The model was then run an additional 100 times with a single variable increased by 5% while all others were held constant. This procedure was repeated for each of the four variables. Sensitivity was calculated as percent change in market index value divided by percent change in the variable, averaged over 100 simulation runs. Results are listed in table 2.

Table 2. Local Sensitivity Analysis

Variable	Reference value	Sensitivity
Interest rate	0.01	-0.0231
Average look back days	20	0.0109
Number of traders	100	0.0172
Prospect rate	0.1	0.0024

3.3 Robustness

Robustness of the model was tested with respect to four variables: number of stocks, average look back days, number of traders, and cash. For each variable, the simulation was run 100 times with the variable moving through all plausible values while all other variables were held constant. In all four cases the results were within expected ranges with no OLS regression line showing any substantial unexpected trend. Results are listed in figures 3, 4, 5, and 6 respectively.

Figure 3.

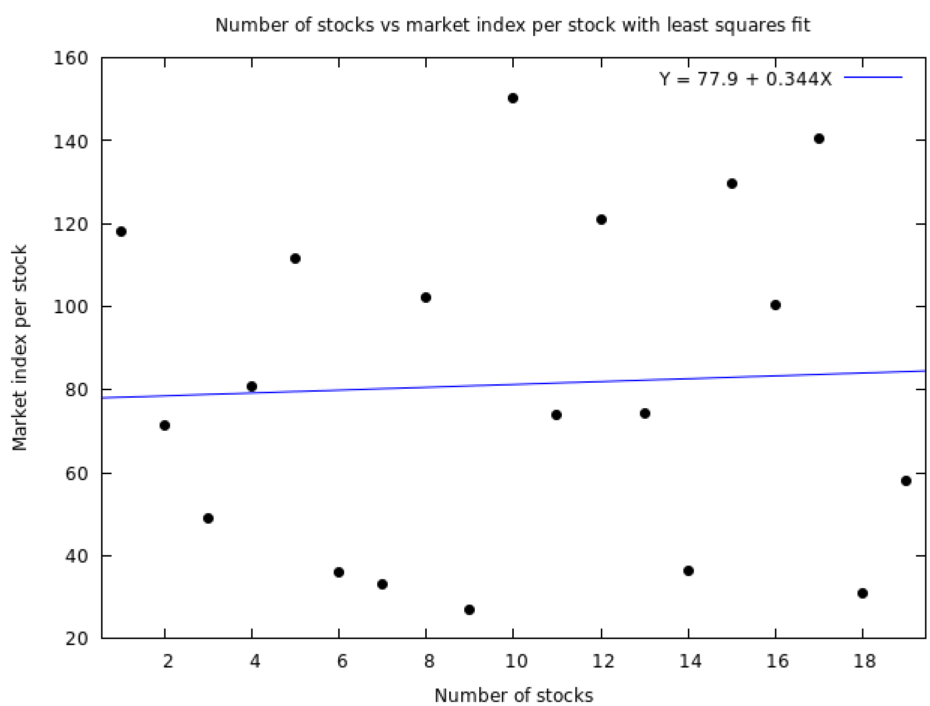

Figure 4.

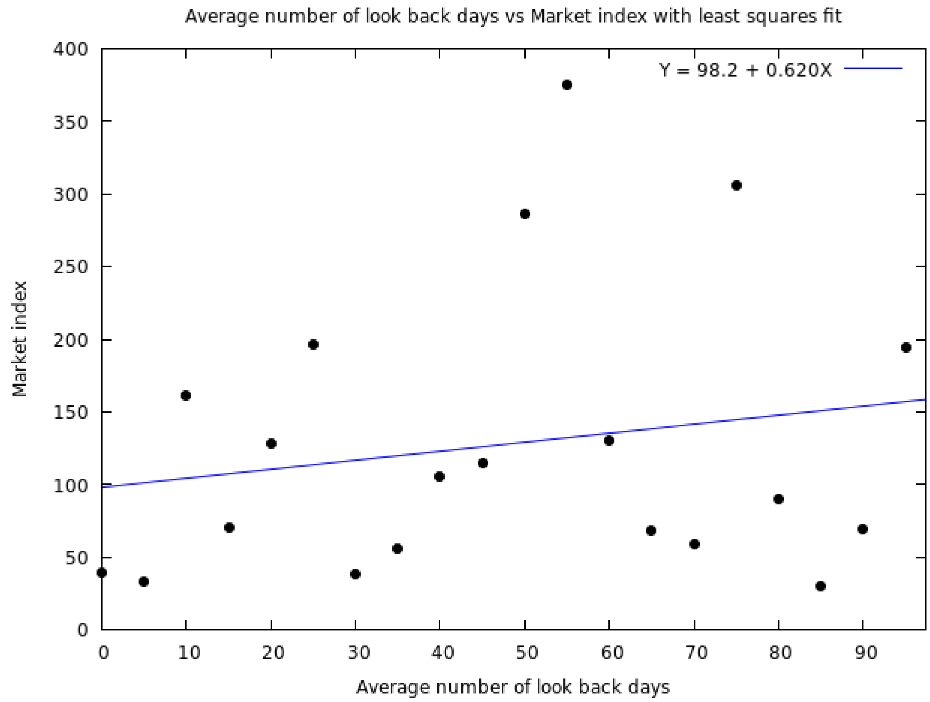

Figure 5.

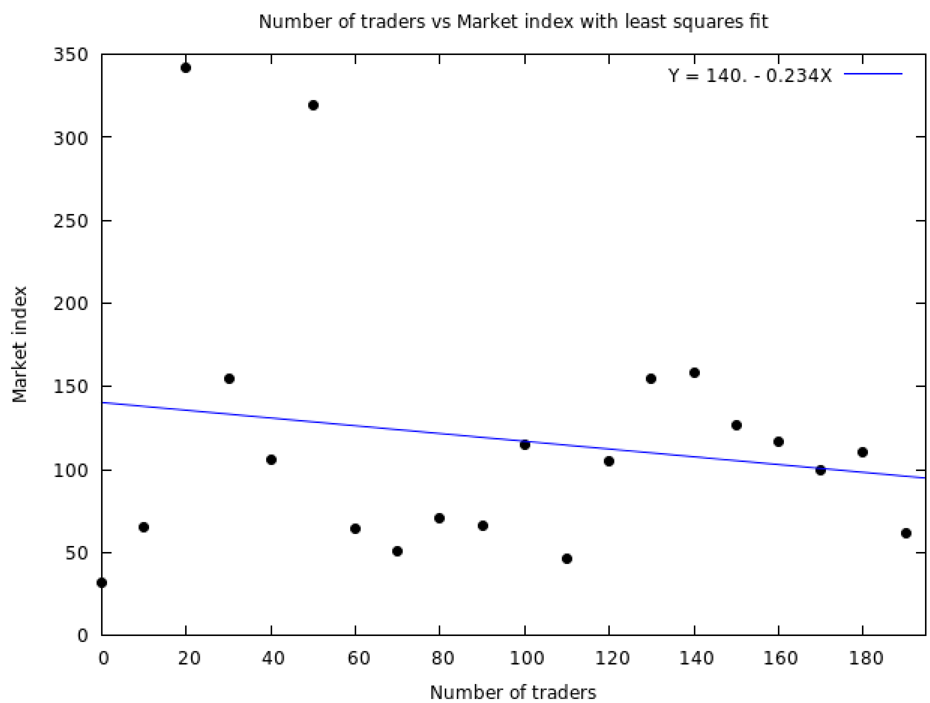

Figure 6.

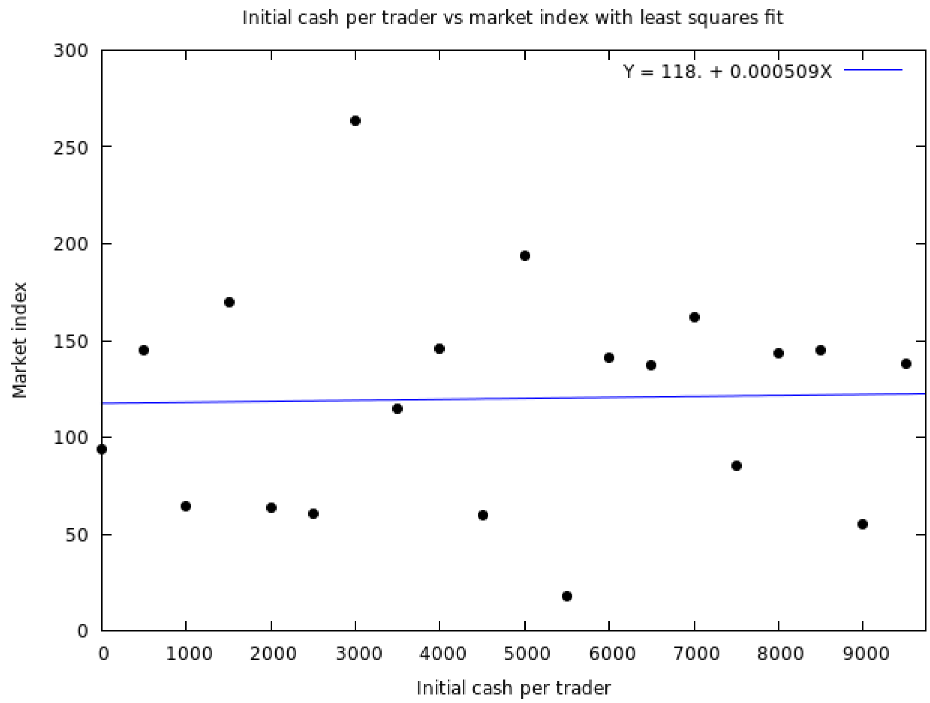

4. Results

The simulation was configured with the following parameter values:

NUMBER_OF_RUNS=1

NUMBER_OF_STOCKS=1

NUMBER_OF_TRADERS=100

NUMBER_OF_DAYS=365

CASH=10000

AVG_LOOKBACK_DAYS=20

INTEREST_RATE=1.00

MARKET_TREND=0.0

PROSPECT_RATE=0.5

It was then run four times, once each using the following settings:

Zero intelligence agents with no prospect irrationality

Zero intelligence agents with prospect irrationality

Optimizing agents with no prospect irrationality

Optimizing agents with prospect irrationality

For each run data was grouped by day of the week to determine if market returns showed a correlation with a particular day. Summary statistics are shown in table 3.

Table 3. Summary statistics by day of the week

Group	Statistic	Monday	Tuesday	Wednesday	Thursday	Friday
Zero intelligence No prospect	Mean	59.77	60.17	60.02	59.60	59.39
	Standard Deviation	14.07	14.06	13.46	12.93	12.54
Zero intelligence Prospect	Mean	282.68	284.27	283.28	282.93	283.40
	Standard Deviation	78.76	79.05	76.58	75.85	75.68
Optimizing No prospect	Mean	97.48	97.60	97.66	97.61	97.30
	Standard Deviation	19.52	19.24	20.38	20.17	20.70
Optimizing Prospect	Mean	55.62	55.94	55.94	55.82	55.39
	Standard Deviation	21.21	21.97	22.52	22.17	21.95

Analysis of variance was conducted for each group by day of the week to determine if returns were significantly different on any day of the week within any group. In all four cases P-values were greater than 0.99 indicating no evidence of difference between the days of the week. Results are shown in tables 4.

Table 4. ANOVA by day of the week

Group	F-statistic	P-value
Zero intelligence No prospect	0.0397	0.9969
Zero intelligence Prospect	0.0045	0.9999
Optimizing No prospect	0.0038	0.9999
Optimizing Prospect	0.0084	0.9998

5. Conclusion

There was no sign of any calendar effect in any of the configurations tested to a very high degree of confidence. In addition, there was no obvious difference in the market results generated between optimizing traders following multiple complex strategies and simple zero intelligence traders constrained only by their respective liquidity.

This leads me to conclude one of three things must be true. Either 1.) my simulation failed to generate calendar effects because it is an inadequate model of a real world stock market, or 2.) it failed because it is an adequate model, but day of the week calendar effects found in other research are statistical anomalies created by excessive data mining, or 3.) calendar effects do exist in real world markets, but they are not caused by prospect theory irrationality. On the basis of this experiment, I have no factual grounds to differentiate between these possibilities.

Stepping away from the data, I do have a gut feeling that my model was implemented adequately to test my theory. My suspicion is that weekday based abnormal returns in large efficient markets are statistical ghosts rather than real phenomena. If economists can find patterns in market data, the famous (or infamous) quants on Wall Street with access to nearly unlimited computational resources could surely find them as well in pursuit of profits through arbitrage.

A next possible step for my stock market simulation may be to rebuild it from the start, but this time in Java rather than Python. I expect to be learning Java as part of future CSS courses and recreating the simulation in a new language would provide an opportunity for comparison as well as good practice to develop my programming skills.

References

Benhammadaa, Sadek and Salim Chikhib . "A semi-formal specification for a generic model of artificial stock markets ." *Procedia Computer Science* 1 (2012): 1465-1474.

Hnasen, Peter, Lunde Asger and Nason James. "Testing the Significance of Calendar Effects ." *Federal Reserve Bank of Atlanta Working Paper 2005-2*(2005)

Kahneman, Daniel and Amos Tversky. "Prospect theory: An analysis of decision under risk." *Econometrica*47 (1979): 263-291.

LeBaron, Blake. "A builder's guide to agent-based financial markets ." *Quantitative Finance* 1 (2001): 254-276.

LeBaron, Blake. "Building the Santa Fe Artificial Stock Market ." *Working Paper* , http://www2.econ.iastate.edu/tesfatsi/blake.sfisum.pdf (accessed December 12, 2013).

Samanidoul, E, E Zschischang, D Stauffer and T Lux. "Agent-based models of financial markets ."*Reports on Progress in Physics* 70 (2007): 409-450.

Stigler, George and Gary Becker. "De Gustibus Non Est Disputandum ." *The American Economic Review* 67 (1977): 76-90.

Sullivan, Ryan, Allen Timmermann and Halbert White. "Dangers of data mining: The case of calendar effects in stock returns." *Journal of Econometrics*105 (1979): 249-286.

Swinkels, Laurens and Pim van Vliet. "An Anatomy of Calendar Effects." *Journal of Asset Management*13(4) (2010): 271-286.