Modeling Decision Making Biases in the Context of a Market

Kawika Pierson
MIT Sloan School of Business PhD Student
305 Memorial Drive #207
Cambridge MA, 20139
617-838-2888
kawika@mit.edu

Keywords: Decision making biases, Market, Overconfidence, Availability, Trading

Abstract

The effects of two behavioral decision making biases are evaluated within the context of a system dynamics model of a market for a commodity, overconfidence and availability. Overconfidence is modeled as an increase in the percent of a trader's capital they are willing to commit to any trade and is found to have the effect of increasing profits for traders with good information relative to traders with poor information, as well as increasing the volatility of the returns for traders with good information more than for traders with poor information. The Availability Bias is modeled as a overweighting of information easily available to a trader and is found to have the effect of increasing the returns of traders with good information easily available to them and decreasing the returns of traders with poor information easily available.

Introduction

Quantifying the effects of decision making biases on trader's profits is a difficult problem. Complexities arise from several angles, including difficulties inherent in running controlled experiments in continuously shifting market conditions, difficulties in finding trading firms willing to provide access to their traders and also from difficulties isolating and quantifying when, where and how these biases intervene in trading decisions. Because of these difficulties, system dynamics lends itself naturally as a tool for accomplishing this goal.

With this goal in mind, this paper starts by constructing a system dynamics model of a market with functionalities inherent in it that will allow us to run detailed tests of several hypotheses about the effect these biases might have on traders. The model allows us to run repeated, controlled experiments over underlying market conditions that are identical in each case, as well as isolate the exact mechanism and size of the biases in each case. We then lay out the experimental framework for how we will test these hypotheses and display and analyze our results.

Model Structure

Supply, Demand and the Market

The basic engine behind the model is a pair of supply and demand curves that solve for the correct price with twin delays. These two table functions that lookup price and read out what the values of supply and demand ought to be function exactly like supply and demand curves familiar with economists, demand decreases with price increases and production reacts in the other direction. Supply and demand can not instantly respond to the signals price is sending them however. As in the real world, supply and demand in the model adjust to their price indicated values with some delay. In general the delay time for price effecting supply will be different than the delay time for price effecting demand, therefore the model allows for these constants to be set separately. The values for supply and demand above feed back into the model's determination of price.

While none of the tests conducted in this paper utilize this functionality, the model has implemented parameterization of these table functions that allow the user to test more complicated supply and demand relationships such as sudden shocks to producers cost functions or gradual decreases in consumers consumption habits due to viable alternatives coming to market. This implementation was created along the lines of that found in Repenning's "Dynamics of Implementation" paper (Repenning 2001).

Each unit of the commodity that is produced must be hedged by selling a contract on the market and every unit of the commodity that is demanded must be hedged by buying a contract on the exchange. The model then compares the number of contracts long (buying) to the number of contracts short (selling) in the market at any one time and based on which number is higher and the size of the difference adjusts price at the next time step within some limited range. The exact mechanism of this price adjustment is accomplished with a table function. The ratio of the residual contracts [abs(long-short)] to the total number of contracts traded is the input to the table function which then outputs the magnitude of the price change resulting from that trading. In the absence of speculative trading these contracts would be the only ones traded on the exchange, and would provide the mechanism through which production and consumption decisions could feedback through price onto future decisions. Thus, when demand is much larger than supply more contracts will come to the market long, price will move up and over time demand will slack, production will grow and the market will come into balance. A symmetric description is true for the case of production being higher than demand. Figure one shows this base case result of the model. Price, supply and demand are initially in disequilibrium, they quickly adjust, but oscillate around their goal, coming into equilibrium by the end of the model run.

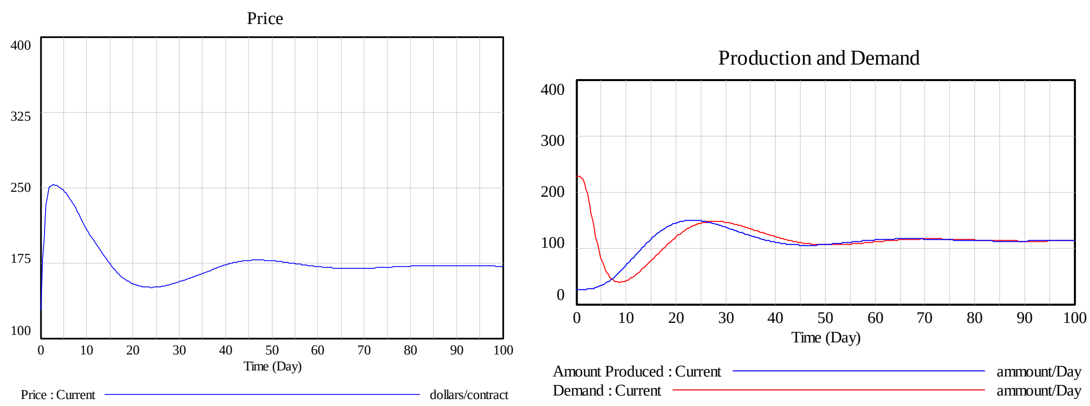

Figure 1. Price, Supply and Demand in the base case of the model with no trading

This mechanism for price discovery is substantially different from the actual process of discrete bids and offers within a market, but a case can be made for it approximating that process. On an exchange, individual bids and offers may respond independently of the total order flow in any one direction; however the bids and offers given by floor traders are very sensitive to their interpretation of how much outside order flow is coming from the long side as opposed to the short. If there is a great imbalance in one particular direction the price that floor traders offer moves to correct that imbalance. This is the process that the price discovery mechanism in the model mirrors. This implementation has the advantage of capturing the underlying dynamics of how a market discovers price without having to model the sometimes messy discreteness found in the actual bid matching process.

<u>Speculation</u>

There are two strategies available to traders in the model, fundamental strategies and technical strategies. These two classes correlate well with the major classifications of trading strategies practicing traders would identify. Technical traders in practice follow a range of strategies which all posit that past prices contain some information about what future prices will be. Some are essentially trend followers, while others trade a variety of other indicators. Within the model, one of the most popularly used indicators, the moving average cross rule, is used as the decision rule of technical traders. Essentially, this rule works by making each trader compute two different smoothes of price, one of which uses a shorter averaging time than the other. When the value of the short moving average rises above the value of the long moving average the decision rule says that the market situation is bullish and the trader should get long, when the value of the short moving average is below the value of the long moving average, the decision rule says that the situation is bearish and the trader should get short. This decision rule works because the short moving average reacts more quickly to the trend of price then the long moving average, and thus their relative values will signal the short term direction of the trend in price. In this model, the values that the traders use for the long and short averaging time cascade in order to represent more fully the wide spectrum of trend following techniques employees in modern markets.

The Finance community has largely shunned technical analysis due in part to the influence of the Efficient Markets hypothesis' random walk theory; see for instance Fama's survey. (Fama 1970) One of the seminal revelations of the hypothesis was that markets fully discount all relevant information, which means that forward prices can not be projected from past price data and thus the primary approach used by technical analysts could not possibly work. More recent work has suggested flaws in this theory however, Lo and MacKinlay published one example, in which the authors find that the random walk theory can not be upheld, since stock prices have a small autocorrelation. (Lo and MacKinlay 1998) Despite this debate however, there is ample evidence that the existence of technical trading is a behaviorally realistic fact about modern capital markets, and thus there is a compelling argument for including its effects in a model of a market.

Fundamental traders look to supply and demand to make their determination of price direction. They cannot obtain this information themselves in the detail they need, therefore they rely on a network of opinion makers to provide them with opinions of the market direction based on fundamental information. Each opinion maker samples a noisy value of the carry out and the demand every “observation interval” days. The size of the noise component of the signal they receive is easily controllable within the model and can be varied independently for each opinion maker creating a situation where certain opinion makers will in general be more correct about the actual value of supply and demand than others. The noise added to these observations is pink in spectrum, and is implemented with a standard pink noise generator. These observations are then used by the opinion maker to compute the forward value of the carry out to use ratio, the trend of which becomes the opinion maker’s opinion on the direction of price. If the trend of the carry out to use ratio is growing then stocks are building and price should go down, whereas if the carry out to use ratio is shrinking then stocks are being used and price should rise. This mimics almost exactly the method employed by commodity analysts, although it neglects the effect of the absolute level of the carry out. Adding in the effect of the absolute level of the carry out is a goal of future extensions of our model.

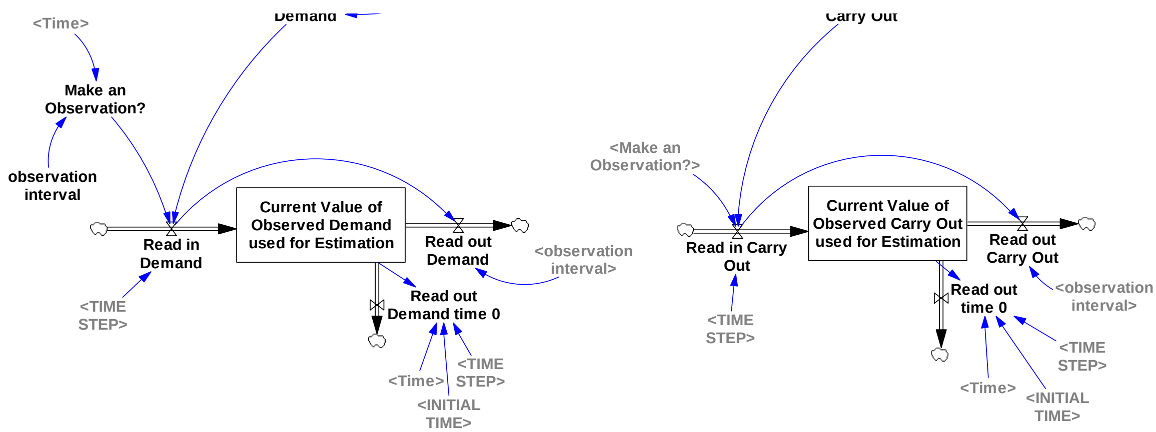

Figure 2. Opinion Makers’ sampling of demand, a non-standard formulation

The sampling mechanism employed to approximate the fact that real world opinion makers can only observe static pictures of the demand and consumption of consumers is the only non-standard formulation in this portion of the model. From a spectral analysis standpoint the harder cutoff represented by this sampling differs substantially from the softer low pass filtering of a smooth. Since a smooth is also used in the formulation to represent the time lag inherent in opinion makers updating their views of supply and demand one might suggest that these two effects be combined into one longer smooth. However this modification significantly alters the spectrum of noise that is passed from the “noisy observations” on through the system to price and greatly reduces the plausibility of the price series' created in the model. For this reason, and also because the sampling mechanism is behaviorally plausible from the author's viewpoint there is significant evidence to warrant using this non-standard structure in the model.

The opinions created by these fundamental analysts are then given to the traders, after passing through an awareness screen which allows information to be known by some traders and not others. The traders then weight each of the market direction opinions representing the degree to which the trader trusts that opinion maker to correctly analyze the market, and then each trader takes the weighted average of these opinions. The sign of this number is the direction the trader thinks the market is moving and the size of the number if how confident the trader ought to be if they were completely rational and actually weighed each opinion in the proportions implied by the model. The confidence of each trader affects the fraction of their total capital committed to each bet they make, so a confidence of 1 will cause the trader to bet 100% of his available capital.

Once traders make a decision using one of these two strategies the model will notice that the trader wishes to make a trade and record the value of price that the trader traded. These trading flags are then used along with the size of the positions each trader took to calculate the number of contracts coming to the market long or short in any given time period. This number then is used to compute price in the same way as described above for the base case except that now the flows of contracts to the market come not only from producers and users but also from speculators.

There is also significant model structure dedicated to tracking the profit or loss over time from each of the trader's trades and adjusting the level of capital each trader has due to this profit or loss. The data collected in this way will allow us to evaluate the effect of biases on the profits enjoyed by traders.

Decision Making Biases

There is ample evidence within the decision making literature that people are generally overconfident in the correctness of their views. In their paper on Naïve Realism, Ross and Ward note that "[In] prediction, for example, failure to make sufficient allowance for the possibility that the situation facing the actor actually will be quite different than the way we are construing it (and/or the possibility that the relevant actors construal of it will be quite different from our own) breeds unrealistically high levels of confidence, and ill advised gambles." This tenant of naïve realism is undoubtedly a bias which is present within the context of a market. (Ross and Ward 1996) With the many sources of often contradictory information available to traders, a surprising number of them will have highly confident views on the direction of the market at any given time.

Within the model this tendency towards overconfidence can be analyzed by measuring the effect of the constant additive modifier of confidence that explicitly makes each trader more confident than they would have been otherwise. This overconfidence then translates within the model into the trader committing a higher percentage of their total capital to each trade than they otherwise would have. The base case for a world where everyone is exactly as confident as they ought to be given the information available to them is when this variable is set to zero, high values of the variable correspond to greater levels of overconfidence by the trader.

Hypothesis 1 – The natural tendency of traders to be overconfident will produce higher profits for traders with relatively good information and less of an increase in profits for traders with relatively poor information.

This hypothesis is plausible because over the long term, traders with better information will probably be paid by the market to be overconfident since their confidence will let them profit more from their information. Their overconfidence will simply make their already good bets larger and so more profitable. This is likely not the case for traders using comparatively poor information about supply and demand, since increasing the size of their trades will on balance cause them to make less money than their better informed counterparts.

Kahneman and Tversky present several decision making biases in there seminal paper. One explored in this model is "Availability," the thrust of which is that people tend to be biased towards information that is more available to them. In the model, we can evaluate the effect of this bias by separating the traders into three groups, ten traders in firm 1, ten in firm 2 and ten independent. We will then compare the size of the trader's cumulative profits in the case of all traders valuing all information equally to the case of traders overweighting information available to their firm. (Kahneman and Tversky 1988)

Hypothesis 2 – The availability bias will have a net negative effect on trader's profit unless the information easily available to them is better than the information that would be possible, but difficult to observe.

This is likely to be the case since overweighting information from particular opinion makers will cause traders to rely too heavily on one set of opinions which on balance are more likely to be incorrect than a more evenly weighted average of all available opinions. If it were the case that the information easily available to the trader provided a much better picture of price's likely direction than the other information the trader could access with more effort, then it would be likely that this bias would work in favor of the trader, but that is the only case where this would be the effect. The trader would have better information, and the good practice of heavily overweighting that information. .

Results

In order to test the effect of overconfidence on traders profits, the traders were split into two groups, one that relied on opinions with noise standard deviations between 30 and 50 added to the observations of supply and demand, and one that relied on more accurate opinions ranging from 10 to 30 in noise standard deviation. Two hundred simulations were run for each value of the overconfidence factor, with the noise seed varied for each simulation and the results were averaged and recorded in the table below:

	Average of Profits						
Overconfidence Factor	0	0.05	0.1	0.2	0.3	0.4	0.5
Good Information	64871	90194	113592	181808	256134	302064	262960
Poor Information	36508	33076	66071	111957	169173	204953	191190
Difference	28363	41626	47521	69850	86961	97111	71770
	Standard Deviation of Profits						
Overconfidence Factor	0	0.05	0.1	0.2	0.3	0.4	0.5
Good Information	19657	33076	38566	62492	102136	144680	119785
Poor Information	12850	19081	27779	47539	85108	103996	103930
SD as % of AVG	30.30%	36.67%	33.95%	34.37%	39.87%	47.90%	45.55%
	35.20%	39.29%	42.04%	42.46%	50.31%	50.74%	54.36%

Figure 3. Results for the Monte Carlo simulation of the Overconfidence Bias

The results confirm hypothesis one, but show a slightly more nuanced picture. For one thing, the profits of both traders with poor information and good information increased for increases in the overconfidence factor up to and including 0.4. This is due primarily to the fact that the sums bet by both groups of traders were larger with larger levels of overconfidence and on average the trades recommended by fundamental analysis were profitable. Thus larger sums of money risked on profitable trades will result in more money made. One reason why the profits of both groups actually dropped for the change from 40-50% overconfidence is that by the time the model was running at 40% overconfidence only very few trades were being executed at less than 100% of capital committed. These were the trades that were affected by the increase in confidence, and consequently the increase in capital committed actually decreased the profits of traders because the increase capital was placed on losing trades.

As hypothesis one would suggest, the overconfidence bias helped traders with good information more than traders with poor information for all increases in the overconfidence factor up to and including 0.4. This can be seen from the increase in the value of "Difference" in figure three. For higher levels of overconfidence, this difference shrank, as more and more of the less profitable bets for each class of trader were undertaken with large amounts of capital.

The standard deviation of the profits for each class of trader also tells an interesting story. As can be seen in the lower half of figure three the standard deviation of the profits for traders with both good and bad information, as a percent of their average profits, increased substantially from the base case to the cases with high overconfidence. This data suggests that in the real world overconfidence may increase the volatility of returns for traders, suggesting an interesting area for research into the effects of these biases that could serve as a test of the model's results.

In testing the effect of the availability bias on traders profits a similar set of tests were conducted. In these simulations the opinion makers and traders were each separated into three groups. The first three opinion makers had noise with standard deviations ranging from 5 to 15, the second group ranged from 70 to 90 and the last four opinion makers had noise standard deviations ranging from 20 to 50. These groups were designated as firm one's proprietary information, firm two's and public information respectively. The traders were then divided into three groups, with the first ten belonging to firm 1, and heavily overweighting their proprietary information, the second ten belonging to firm 2 and heavily overweighting their proprietary information and the last ten trading only on the information available publicly. In the cases where a trader had access to proprietary information, the public information was also a part of their decision making process, but was underweighted compared to the information easily available to them from within their firm.

Using this setup, two hundred simulations were run varying the noise seed for the opinion makers' observations and the average profit for each group of traders was recorded. The base case is the case where each trader is approximately as likely to weight proprietary information heavily as they are to weight publicly available information heavily. The results are shown in this table:

Average Profits	**Base Case**	**Biased Case**	**Percent Change**
Firm 1 (Good)	$38352	$39621	3.31%
Firm 2 (Poor)	$22345	$16536	-26.00%
Independent	$38529	$26517	-32.18%

Figure 4. Average Profits of traders for the test of the Availability Bias

The results above confirm hypothesis two, since the profits for traders with good information rose, while the profits for traders with poor information fell. Two things about the data stand out however. One is the fact that the increase in profits for firm 1 is a good deal smaller than what one might have guessed would be the case given the superiority of the information they were using to make their trades. Results very similar to this were observed for repeated runs with different noise seeds, so the increase is likely not a statistical artifact, but its size sheds some doubt on the hypothesis that the availability bias has a measurable positive effect on traders with relatively good information in the real world. If there is any such positive effect it is likely very small.

The negative effect of the availability bias for traders with poor information is evident beyond any doubt looking at the data, what is surprising though is the extent to which the independent third of traders had their profits reduced by the other firms trading with the bias. Since the weights each independent trader placed on each information source did not change between the two runs one might be tempted to say that there must be some flaw in the model's formulation. However the fact that these trader's profits change even though their weightings did not is a plausible outcome of the other traders changing their weights. Since the weights used by the other traders influence their trading decisions which in turn influence the price traded on the exchange, which in turn influences supply and demand decisions which in turn further influence prices as well as the signals being sent by the opinion makers in the model we would expect some change in the profits of the independent traders given changes in the weights of the other twenty.

In fact, this insight gives us another angle to consider the data from. If we were willing to say that the Independent traders were a "control" group in that the weights that they placed on the incoming sources of information did not change, then we might be willing to say that the market returns for each of the groups should thought of as competing against a benchmark change of -32.18 percent in the control. This thinking would increase the returns exhibited by the group one traders, and edge the returns of the group two traders to become slightly positive. However the dynamics of the situation are complex enough that this sort of linear thinking is rarely the correct heuristic to apply to these situations, and so for the purposes of this paper the unadjusted results will be taken as final. Efforts to study more in depth the mechanisms and effects of the availability bias in markets are needed, and potential avenues for future research.

Conclusions and Ideas for Future Research

The results from our simulations uphold and extend the hypotheses outlined in the decision making biases section. The effect of the overconfidence bias within the model was to amplify the effects of the relative information quality of the trader up to a limit. At this limit, the trader's overconfidence caused them to enter into marginally poor trades with such a large percentage of their capital that the positive effect of overconfidence for traders with good information was undercut. Further, the volatility of profits for all traders, as a percentage of their average profits, increased with increases in overconfidence. This result suggests a potentially fruitful avenue for research into the effect of these biases in actual trading settings, though the usual difficulties with access and quantification of parameters will arise.

The effect of the availability bias was negative for all groups, except for the group with relatively good information. Traders who were fortunate enough to have high quality information about the markets easily available to them were slightly more profitable, but in general overweighting information that was easily obtainable had a strong negative effect on profits.

Some ideas for further applications of the model include fleshing out the supply and demand sectors of the model, potentially by incorporating it with the Sterman's Commodity Cycle model. (Sterman 2000) This would allow for a richer picture of how speculators effect decisions about what to produce and consume through the feedbacks from price onto supply and demand, as well as increase how realistic the market conditions faced by speculators in the model are. Additionally, the price setting mechanism in the model is only an approximation of the actual process that takes place within a market. An extension of this mechanism to mirror the discrete bid ask process of a modern market could be helpful for increasing the realism of the results, although at the time scales considered by the model, the approximation used is adequate. Also, extending the model to capture the effect that the absolute level of the carry out to use ratio has on fundamental traders' opinions of price direction will help to bring the model closer to the actual decision processes employed in the market.

Cited Works

Fama, Eugene F. 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance*: 383-417

Lo, and MacKinlay. 1988. Stock Market Prices do not follow random walks: evidence from a simple specification test. *The Review of Financial Studies*: 41-66

Repenning, Nelson. 2003. A Simulation-Based Approach to Understanding the Dynamics of Innovation Implementation. *Organization Science* 13: 109

Ross, L. and Ward, A. 1996. Naïve realism in everyday life: Implications for social conflict and misunderstanding. In *Values and Knowledge,* ed. E.S. Reed and E. Turiel, 103-35. Mahwah N.J: L. Erlbaum Associates.

Sterman, John. 2000. The Invisible Hand Sometimes Shakes: Commodity Cycles. In *Business Dynamics*: *Systems Thinking and Modeling for a Complex World,* Chapter 20. McGraw-Hill/Irwin.

Tversky, and Kahneman. 1974. Judgment under uncertainty: Heuristics and Biases. *Science* 185: 1124-31

Appendix A – Weights given by traders in Tests of Hypothesis 1

1,3,1,1,1,0,0,0,0,0; 1,2,1,1,2,0,0,0,0,0; 2,1,1,1,3,0,0,0,0,0; 3,1,2,1,2,0,0,0,0,0; 3,1,3,1,2,0,0,0,0,0; 3,1,2,2,1,0,0,0,0,0; 2,1,1,1,2,0,0,0,0,0; 2,1,3,1,2,0,0,0,0,0; 2,2,3,3,3,0,0,0,0,0; 1,2,1,3,1,0,0,0,0,0; 2,3,2,3,1,0,0,0,0,0; 2,3,2,1,1,0,0,0,0,0; 1,1,1,2,3,0,0,0,0,0; 1,1,2,2,2,0,0,0,0,0; 3,1,2,1,2,0,0,0,0,0; 0,0,0,0,0,2,3,1,2,1; 0,0,0,0,0,1,1,2,1,3; 0,0,0,0,0,2,1,3,1,2; 0,0,0,0,0,1,1,1,2,2; 0,0,0,0,0,1,2,2,2,1; 0,0,0,0,0,1,1,1,1,1; 0,0,0,0,0,3,1,2,1,2; 0,0,0,0,0,3,3,3,1,1; 0,0,0,0,0,1,2,1,1,2; 0,0,0,0,0,2,1,2,1,2; 0,0,0,0,0,1,1,1,2,3; 0,0,0,0,0,2,1,2,2,1; 0,0,0,0,0,3,1,3,3,1; 0,0,0,0,0,3,2,2,1,2; 0,0,0,0,0,3,3,2,1,2;

Appendix B – Weights given by traders in Tests of Hypothesis 3

9,4,6,0,0,0,1,1,1,1; 7,3,8,0,0,0,1,1,1,1; 9,8,9,0,0,0,1,1,1,1; 4,7,6,0,0,0,1,1,1,1; 9,5,5,0,0,0,1,1,1,1; 4,8,4,0,0,0,1,1,1,1; 9,3,4,0,0,0,1,1,1,1; 7,6,2,0,0,0,1,1,1,1; 1,2,9,0,0,0,1,1,1,1; 9,4,7,0,0,0,1,1,1,1; 0,0,0,9,6,3,1,1,1,1; 0,0,0,5,9,8,1,1,1,1; 0,0,0,2,7,9,1,1,1,1; 0,0,0,7,6,3,1,1,1,1; 0,0,0,9,5,3,1,1,1,1; 0,0,0,2,4,6,1,1,1,1; 0,0,0,8,3,2,1,1,1,1; 0,0,0,5,3,2,1,1,1,1; 0,0,0,1,2,3,1,1,1,1; 0,0,0,4,6,1,1,1,1,1; 0,0,0,0,0,0,9,7,5,3; 0,0,0,0,0,0,7,5,3,1; 0,0,0,0,0,0,5,7,5,3; 0,0,0,0,0,0,5,9,5,3; 0,0,0,0,0,0,1,3,7,3; 0,0,0,0,0,0,2,5,9,5; 0,0,0,0,0,0,1,3,5,7; 0,0,0,0,0,0,3,5,7,9; 0,0,0,0,0,0,1,1,1,1; 0,0,0,0,0,0,4,8,4,8;

9,4,6,0,0,0,7,3,8,1; 7,3,8,0,0,0,9,4,6,1; 9,8,9,0,0,0,1,7,3,8; 4,7,6,0,0,0,4,6,1,1; 9,5,5,0,0,0,1,4,6,1; 4,8,4,0,0,0,2,9,5,3; 9,3,4,0,0,0,7,6,2,1; 7,6,2,0,0,0,3,2,9,5; 1,2,9,0,0,0,9,4,9,5; 9,4,7,0,0,0,9,8,9,1; 0,0,0,9,6,3,3,9,4,6; 0,0,0,5,9,8,6,1,5,3; 0,0,0,2,7,9,1,9,5,3; 0,0,0,7,6,3,1,9,8,9; 0,0,0,9,5,3,1,7,6,2; 0,0,0,2,4,6,1,5,3,2; 0,0,0,8,3,2,6,3,1,9; 0,0,0,5,3,2,9,5,3,1; 0,0,0,1,2,3,5,3,2,1; 0,0,0,4,6,1,9,6,3,3; 0,0,0,0,0,0,9,7,5,3; 0,0,0,0,0,0,7,5,3,1; 0,0,0,0,0,0,5,7,5,3; 0,0,0,0,0,0,5,9,5,3; 0,0,0,0,0,0,1,3,7,3; 0,0,0,0,0,0,2,5,9,5; 0,0,0,0,0,0,1,3,5,7; 0,0,0,0,0,0,3,5,7,9; 0,0,0,0,0,0,1,1,1,1; 0,0,0,0,0,0,4,8,4,8;