



What Do You Do If You've Been at the Poker Table for Twenty Minutes and Still Can't Spot the Sucker? Implications for Individual Investors

Greg Samsa, PhD

Duke Department of Biostatistics and Bioinformatics,
11084 Hock Plaza, Durham NC 27510, USA.

ABSTRACT

In comparison with other market participants, individual investors face an informational deficit. Recognizing this, we ask whether strategies exist which might nevertheless hold out hope for them outperforming the market. We proceed by simulating the behavior of mean-regressing stocks, and estimate the excess returns of a strategy of writing a short-term covered call option and, if the option is not exercised, selling the stock the next time it exceeds the purchase price. Excess returns in the range of 4-6% were observed. Obtaining these returns did not require accepting additional risk, either in the economic or statistical sense, as mean-reverting stocks tend to have both less business risk and lower price volatility than others. Nor did they require that the investors purchase stocks below their true value, only that the investor could identify mean-reverting stocks, and also identify a lower bound on the value around which prices will regress. One of the advantages possessed by individual investors is that their trading behavior is unlikely to affect prices, and thus more strategies are realistically available to them. This is an example of a strategy which couldn't be easily implemented by an institution but might be appropriate for an individual investor.

Key words: Stock market, investment strategies, individual investors, regression toward the mean

INTRODUCTION

Unlike most of the contents of my inbox, this particular email piqued my attention. A hedge fund manager wanted a statistician's opinion about itching, promised that our conversation would be limited to publicly available information, and offered a generous hourly rate. I assumed that some drug trial was being performed which would be pivotal to a company's future, that the purpose of the drug was to reduce itching, and that by considering the study design – for example, its sample size, choice of measures, etc. – I could help refine their prediction about whether or not the trial would be positive. Calling in a statistician seemed like a great idea, and is the sort of resource that few individual investors would have, and that's one of the many structural disadvantages which individual investors face.

Although I can't vouch for a direct relationship between cause and effect, soon after responding to the email I began to daydream. In my dream some of my colleagues had invited me to a poker game, the stakes exceeded my comfort zone but, social pressures being what they are, I nevertheless found myself sitting behind a stack of chips and a hand of No Limit Texas Hold'em was being played. My hole cards were solid, I put on some mirror shades just like the poker professionals on television, made a bet, and was immediately raised. Then I thought to myself: "Almost always, within 20 minutes at the poker table its clear to me who's the sucker, and the best strategy is to pick on them. But I've been here half an hour and still don't see one." And, of course, I recalled the second part of the saying, too: "If you can't spot the sucker at the table, it's you."

I'm an individual investor, and even though the world has advanced to the point where a competent trading platform contains fairly comprehensive financial information, assessments from multiple analysts, multiple predictions about corporate performance, etc., and although the nature of the informational advantage in question has changed over time, it's still the case that larger investors retain a significant informational advantage over me. In other words, when it comes to predicting corporate earnings (among many others), not even to mention deconstructing those earnings from an accounting perspective, I'm the sucker at the table.

I'd recently been reading some of the academic literature in the hope of finding articles describing strategies that are especially suitable for individual investors, and came upon a thoughtful effort (Coval, Hirshleifer & Shumway, 2005) with the intriguing title "Can individual investors beat the market?". The authors used actual trading data from individual investors within two time periods, identified those with superior performance during the first time period, asked whether those same investors outperformed the others during the second time period, and found that they did. Out of an abundance of caution the authors even tested the hypothesis that this excellent performance was based on inside information and concluded that it wasn't. To extend the analogy, and recalling that we're the suckers at the table, the literature on individual investors points less to them having inside information than having an inside straight and making the mistake of drawing to it -- some manifestations of this phenomenon being over-trading, selling winning stocks too soon, and holding losing stocks too long.

All of this suggests the question: "What did the individual investors in Coval's study do to outperform their colleagues?", and also the more general question "Recognizing that they begin with an informational deficit, is there nevertheless anything individual investors might do to try to outperform the market?".

METHODS

As this work is primarily intended to be conceptual, we proceed via simulation. More specifically, we simulate monthly stock prices for the subset of companies whose stock price tends to undergo regression toward the mean (RTM). We then apply a strategy which is intended to take advantage of the noisiness of stock prices while recognizing the limitations of individual investors, and estimate its excess returns. Finally, we discuss the implications of our findings.

CASE STUDY

Specifications

We simulate stock prices over time. As specifications, we assume that the stock in question has a baseline price of $Y_0 = \$100$ per share, and also that $T=100$ is the point around which stock prices will regress. During each of 240 subsequent months, the new price Y_i is derived from the following formula:

$$Y_i = Y_{i-1} + \beta(T - Y_{i-1}) + \varepsilon_i.$$

Here:

- Y_{i-1} is the price during the previous month.
- So long as it is positive, β (where $0 < \beta < 1$) induces regression toward $T=100$. If $\beta=0$ there is no RTM and the model is consistent with the Efficient Market Hypothesis. As β increases, so does the impact of RTM. We use values of β of 0.01, 0.05, 0.10 and 0.20 – 0.10 is probably a reasonable default.
- The random perturbation ε_i is drawn from a standard normal distribution (i.e., with mean 0 and standard deviation 1), independent of the previous perturbations and prices.

Thus, the sequence of prices $\{Y_i\}$ is a realization of a stochastic process centered around 100 with a modest tendency toward mean-reversion. The rationale for the simulation specifications is discussed later.

We assume that the investor in question purchases the stock for \$100 and then immediately writes a covered call option, with a monthly expiration date, a premium of \$3 and a strike price of \$100. In other words, the investor will earn the call premium of \$3 no matter what. If the price exceeds \$100 the stock will be called away – that is, the purchaser of the call will exercise the right to buy the stock at \$100. If the call is exercised, the investor finds an identical stock with a price of \$100 and repeats. Otherwise, the investor holds the stock and sells it during the first month in which the price exceeds \$100. At the end of the holding period the stock is sold, despite the potential loss.

Table 1: Illustration of the logic of the simulation (first 12 months for one replication, $\beta=0.10$)

Month	ϵ	Previous price	RTM	Possible new price	Actual new price	B: call premium	B: contribution to return	B: cumulative return
1	1.80	100	0	101.80	100	3.00	0	3.00
2	-0.08	100	0	99.92	99.92	3.00	0	6.00
3	0.40	99.92	.01	100.32	100	0	0.32	6.32
4	-1.08	100	0	98.92	98.92	3.00	0	9.32
5	2.24	98.92	.11	101.26	100	0	1.26	10.62
6	-0.62	100	0	99.38	99.38	0	0	10.62
7	0.51	99.38	.06	99.95	99.95	0	0	10.62
8	-0.09	99.95	.01	99.87	99.87	0	0	10.62
9	-0.59	99.87	.01	99.29	99.29	0	0	10.62
10	.03	99.29	.07	99.39	99.39	0	0	10.62
11	-0.73	99.39	.06	98.71	98.71	0	0	10.62
12	-0.25	98.71	.13	98.59	98.59	0	0	10.62

Table 1 illustrates the logic of the simulation. The price at the end of the first month is $100 + 0.10(100-100) + 1.80$, which equals 101.80, and so the stock would be called away, with the investor earning the call premium of 3.00 (but not the difference between 101.80 and 100). In month 2 the investor earns a new call premium, and at the end of the month owns the stock without encumbrance (i.e., because $Y_2 < 100$ the call option expires worthless). In month 3, when the price exceeds 100 and thus the stock is sold, the investor earns an additional $(100.32-100) = 0.32$. Returns accumulate, and if the simulation were to end at 12 months the absolute return would be $10.62 + (98.59-100) = 9.19$. More generally, the absolute return for this strategy is $AR = \Sigma(\text{call premiums}) + \Sigma(Y_i-100) + (100-Y_{240})$, where the second summation is taken over the months during which the stock is sold, recognizing that if the stock is sold during the first month after purchase the contribution to the absolute return will be 0.

To annualize the above absolute returns, take $R+1 = ((100+AR)/100)^{(1/T)}$, where T is the follow-up period, in years.

All results are based on 100 iterations of the simulation. Appendix 1 provides the SAS code.

To comment on the simulation specifications:

- Because the value of T is assumed to be constant, any observed returns are above and beyond the general increases in stock prices over time, and thus represent excess returns.
- The base price is taken to be 100 for convenience. The intention is that the investor buys stocks with an identical total cost – if the price differs from \$100 per share then the number of shares purchased can be adjusted to achieve this goal.

- It is assumed that new stocks with the same characteristics can be found – in particular, it is assumed that during those times during which the overall market is overvalued at least one mean-regressing stock can still be found which isn't. Often, this can be achieved by purchasing stocks in out-of-favor sectors of the market.
- It is assumed that the investor is able to identify those stocks which are likely to undergo mean-reversion. Indeed, such stocks tend to have less risk from both an economic sense (i.e., stable business prospects) and a statistical sense (i.e., lower price volatility), and thus are particularly appropriate for individual investors.
- It is assumed that the investor is able to place a lower bound on T. In practice, and because of the knowledge deficit that the individual investor faces, the criteria used to define this lower bound will typically be simple and conservative – for example, a dividend which exceeds bond yields and is well-covered.
- Assuming that a stock price will fluctuate around the same value of T for the entire holding period is unrealistic. However, it turns out that most stocks will be held for significantly shorter periods than this, thus limiting the impact of this assumption.
- Assuming that a monthly call option can be sold for a premium of at least 3% is usually realistic, so long as the stock in question exhibits a non-trivial amount of volatility in its price.
- The assumption that the strike price of the call option will be identical to the current price isn't realistic in the literal sense, but the functional equivalent can be achieved by selling a call option at the closest strike price below the current stock price. The call premium would recover the maximum amount of time value plus any difference between the strike price and the current stock price. (Here, the intention is to extract the maximum time value per unit time period, and this is typically achieved through an at-the-money call option and a short time interval.)
- Similarly, because of how their expiration dates are timed, the assumption that the investor can purchase an option that expires in precisely one month isn't realistic in the literal sense, but can be approximated (especially for those stocks with weekly options).
- In practice, the random perturbation isn't necessarily Gaussian, with a particular concern being the presence of negative extreme values. Here, some level of protection against negative extreme values is provided by the conservative assumptions the investor makes during the purchasing decision. Negative extreme values might still occur, but are relatively likely to be corrected within the short- to intermediate term.
- This strategy ignores trading costs (which, nowadays, are relatively trivial) and taxes (which aren't).
- Importantly, it isn't assumed that the investor is able to spot a bargain. Temporarily interpreting T as the stock's "true value", it isn't required that the investor purchase stocks below that true value, but instead that a weaker condition holds: namely, that the investor doesn't purchase stocks above true value. (Of course, if the investor can successfully identify bargains returns will exceed those reported here.)
- It isn't necessary to assume that T equals a stock's "true economic value", nor even that this construct exists and can be precisely estimated, but to the extent that T is based on sound economic logic so much the better.

RESULTS

Table 2 provides some illustrative results, as a function of β .

Table 2: Returns as a function of β

Target	Strike price	Call premium	Beta	Mean (s.d.)	Minimum return	Maximum return	# trades
100	100	3	0.01	2.45 (1.27)	-0.17	5.80	21.16 (12.44)
100	100	3	0.05	4.09 (0.88)	2.24	6.38	38.11 (11.81)
100	100	3	0.10	4.98 (0.70)	3.53	6.87	50.03 (11.13)
100	100	3	0.20	5.98 (0.53)	4.80	7.47	66.11 (10.14)

The excess returns are positive in every case. So long as β exceeds the trivially weak value of 0.01, the mean excess returns fall in the 4-6% range. The greater the value of β the more trades take place, and also the greater the return. For example, when $\beta=0.10$, approximately 2.5 trades per year (i.e., 50.03 trades / 20 years) are made, on average, and the mean excess return is 4.98%. Approximately 90% of the return is derived from call premiums (data not shown).

Table 3 presents the holding periods observed during the simulation when $\beta=0.10$. As might be expected, the call option was exercised in approximately half of the cases. In round numbers, approximately 4 out of 5 times the stock was sold within the first 6 months, approximately 9 out of 10 times the stock was sold within the first year, and the maximum holding time was 6 years.

Table 3: Holding times, $\beta=0.10$

Month	Relative percentage	Cumulative percentage
1	50.4%	50.4%
2	13.2%	63.6%
3	7.5%	71.1%
4	4.5%	75.6%
5	3.5%	79.0%
6	2.4%	81.5%
7	2.1%	83.6%
8	1.7%	85.3%
9	1.5%	86.8%
10	1.3%	88.1%
11	1.0%	89.1%
12	0.9%	90.0%
13-24	6.6%	96.6%
25-36	2.1%	98.7%
37-72 (max)	1.3%	100.0%

DISCUSSION

The simulations demonstrate a result which might initially seem counterintuitive: namely, that under a seemingly innocent (albeit not general) set of assumptions excess returns can be obtained without having to purchase stocks below their true value. Moreover, these excess returns did not require accepting additional risk, either in the economic or statistical sense, as mean-reverting stocks tend to have both less business risk and lower price volatility than others, and so our results are likely to be even stronger when using the criteria of risk-adjusted returns.

Returning to the poker game, once the player in question has the self-awareness to spot the sucker at the table there are only two reasonable courses of action. The first option is to leave the game, which in the context of investments is to stop trying to outperform the market and purchase an index fund instead. The second option is to play a different game, which in the context of poker would involve selecting an unusual strategy which takes advantage of the

sucker's relative lack of skill. For example, the sucker might play fewer hands and go all in on those hands, thus placing others in a difficult position given that the last thing superior players want to do when tangling with a sucker is to increase the volatility of their expected returns.

The covered call writing strategy illustrated here is an investment version of playing a different game. In other words, the individual investor chooses not to play the games of predicting earnings or predicting short-term price fluctuations, but instead attempts to take advantage of the noise in the system. In essence, any price above 100 is considered to be "positive noise" and the stock is sold, either intentionally (i.e., after the first month) or by being called away (i.e., during the first month), and any price below 100 is considered to be "negative noise" and thus the stock is held.

Interestingly, changing the game allows the individual investor to tune out much of the financial media, at least the parts which are concerned with short-term price changes and also those parts which attempt to predict earnings – that is, those parts of the financial media which assume that the individual investor is (and should be) trying to play the same game as other market participants. Individual investors aren't hedge fund managers, and to the extent that the financial media encourage them to behave like ones they are doing a disservice.

Part of the conceptual justification for the variation on a traditional covered call writing strategy described here pertains to how the option premium in question should best be interpreted. For longer-dated options a significant portion of the premium represents disagreement about future earnings, a discussion about which we have argued the individual investor shouldn't enter. On the other hand, the premiums for very short-term options almost entirely represent speculation about upcoming near-term events which are unlikely to affect the long-term value of the company, short-term price patterns as anticipated through technical analysis, etc. These short-term price fluctuations shouldn't be of interest to an individual investor (especially if unfavorable fluctuations are likely to be eventually reversed by RTM), and thus it makes good sense to monetize their expected value.

We've focused on the things the individual investor can't do particularly well. It becomes natural to ask, then, whether there is anything that they might be able to do better than larger and better-informed market participants. Our answer to this question is a qualified "yes". For example:

- Individual investors can trade without affecting prices, thus making more strategies realistically available to them.
- Individual investors don't have customers. For example, individual investors don't have an incentive to chase short-term performance in order to attract customers, they aren't tempted to engage in "window dressing" where losing stocks are culled before the composition of their portfolios is reported, etc. – more generally, individual investors aren't subject to the seemingly perverse incentives of some of the other market participants.

The strategy described here isn't the only way to potentially take advantage of noise among mean-regressing stocks. Elsewhere, we've demonstrated that the strategy of selling such stocks whenever they achieve a modest gain has similar expected returns (Samsa, 2014). That version of the overall idea takes advantage of the fact that the investor isn't encumbered by having written an option, and holds out the possibility that many of the profits will be taken in a matter of hours to days, although the bird in the hand associated with the option premium is lacking. Both strategies differ from the way in which investment strategies are usually evaluated in historical databases – namely, by allowing the holding period to vary rather than

staying fixed. Indeed, we would argue that the more usual approach of keeping the holding period fixed embeds additional noise into the returns estimated from historical databases, thus inducing a bias toward null findings, and moreover that our approach is more consistent with how individual investors actually behave.

From an aesthetic perspective, the proposed strategy is truly and profoundly ugly, as it not only embraces the stereotype that small investors trade too much, sell winners too soon, and keep losers too long, but also hopes to pay taxes on short-term rather than long-term gains. Not even to mention that any snapshot of the portfolio will predominately contain unrecognized losses, and that can be psychologically difficult for those of us who have been trained to salivate at the sight of green and do quite the opposite when encountering red. But it also has a certain appeal: it takes advantage of the level of volatility in the market – a level of volatility caused by computerized trading, hedge funds and other speculators, and all those other factors which imply that most of the short-term changes in stock prices are divorced from, or at least out of proportion to, changes in economic value of the underlying companies which those prices represent. (And it probably does so rather safely – as the companies which would pass the individual investor's rudimentary and conservative screen would also resemble the contents of a typical "value portfolio".)

Another way to state all of this is that in the stock market it's devilishly hard to find the signal in the presence of all the noise. For the small investor whose goals are limited to saving for retirement, college tuition, or other laudable goal, and find their efforts complicated by the high level of volatility stock prices and the sleep deprivation that this can imply, there's a certain justice in having a way for short-term speculators to pay them instead of the other way around.

Of course, if this strategy is to be used to actually trade, it should be recognized that it is "short volatility", with all the risks (and potential remedies) which this implies. It should be used with caution.

References

Coval JD, Hirshleifer D, Shumway T. Can individual investors beat the market? September 2005. Social Science Research Network Electronic Collection. [Http://ssrn.com/abstract=364000](http://ssrn.com/abstract=364000).

Samsa G. A new critique of the traditional method for empirically estimating the returns associated with strategies of stock investment: impact of fixing the investment horizon. *Bentley University Case Studies in Business Industry & Government Statistics*. 2014; 5(2): 102-110.

APPENDIX 1: SAS CODE FOR THE SIMULATION

```

*****;
* This program generates realizations of a stochastic      *;
* process of monthly stock prices, under the assumption *;
* of a modest amount of mean-reversion. These prices   *;
* are then processed through the application of a      *;
* strategy involving selling covered call options.     *;
* When the stock is sold we buy another with identical *;
* Finally, total returns are calculated.                *;
*;
*****;

*****;
* Two datasets are created. T_end contains the data    *;
* from the final month of the simulation, including    *;
* various accumulated variables, which in turn are used *;
* in the calculation of returns. Holding_period counts *;
* the number of months that a stock was held, at the  *;
* time of the sale.                                    *;
*****;

data t_end holding_period (keep=month_counter) test;

*****;
* The outer do loop is over replications of the      *;
* simulation, the inner do loop is over months.      *;
*                                                    *;
* max_rep is the number of replications.            *;
* max_month is the number of months.                 *;
* beta is the impact of regression toward the mean.  *;
* price_target is the price target.                  *;
*****;

max_rep=100;
max_month=240;
years=max_month/12;
beta=0.10;
price_target=100;

do rep=1 to max_rep;

do month=0 to max_month;

*****;
* At baseline (month 0), perform initializations.    *;
* This includes the previous price, taken to be 100, *;
* and also various counters which, for example,     *;
* accumulate the various components of total return. *;
* Call_counter counts the number of written calls,  *;
* call_premium quantifies the call premium, d       *;
* call_month keeps track of whether a call is currently *;
* in force, and premium_counter accumulates the call *;
* premiums.                                          *;
*****;

if month=0 then do;
prev_price=100;
price=100;
new_price=100;
sales_counter=0;
extra=0;
call_counter=0;
call_premium=3;

```



```
    call_month=1;
    premium_counter=0;
    month_counter=0;
end;

*****;
* During the actual months of the simulation (i.e., not *;
* month 0), generate the random perturbation (i.e., *;
* from a standard normal distribution), create a *;
* a potential new price, and then determine whether *;
* potential new price should be overridden by selling *;
* the stock and starting over. *;
*****;

*****;
* Create the random perturbation;
*****;

    delta=rannor(1);

*****;
* Increment the month counter, which keeps track of how *;
* long the stock has been held. *;
*****;

    month_counter=month_counter+1;

*****;
* This assigns the value of price, taken from the *;
* previous month and thus not yet updated, to the *;
* variable name price_prev. *;
*****;

    prev_price=price;

*****;
* Then new price is the old price, plus the impact of *;
* regression toward the mean, plus the perturbation. *;
*****;

    rtm=beta*(100-prev_price);

    price=prev_price+rtm+delta;
    new_price=price;

    if call_month=1 then do;
        call_counter=call_counter+1;
        premium_counter=premium_counter+call_premium;
    end;

*****;
* Ask whether the current price exceeds the target. *;
* If so, update the various counters and set the price *;
* to 100. Extra is the contribution to the total *;
* return. Sales_counter keeps track of the number of *;
* times the stock is sold. Now, the logic has to keep *;
* track of whether a call option is in force. *;
*****;
```

```

if ((price>price_target) and (call_month=1)) then do;
  sales_counter=sales_counter+1;
  price=100;
  call_month=1;
  output holding_period;
  month_counter=0;
end;

if ((price>price_target) and (call_month=0)) then do;
  extra=extra+(price-100);
  sales_counter=sales_counter+1;
  price=100;
  call_month=1;
  output holding_period;
  month_counter=0;
end;

if (price_target>price) then do;
  call_month=0;
end;

*****;
* Only output the results to the file during at the end *;
* of the follow-up period. *;
*****;

if month=max_month then output t_end;
if rep=1 then output test;
end;

end;

run;

*****;
* Using the data from the final months of each *;
* replication, the loss is any deficit at the end of *;
* the holding period, the gain has been accumulated *;
* within the variable extra, and these are the inputs *;
* to the calculation of absolute and annualized return. *;
*****;

data returns;
  set t_end;
  loss=price-100;
  gain = premium_counter + extra;
  absolute_return = loss + gain;
  annual_return=(absolute_return+100)/100**(1/years);
  ar=(annual_return-1)*100;
run;

proc means data=returns n mean std min max sum maxdec=2;
  var beta price_target annual_return price loss extra
  premium_counter gain absolute_return sales_counter ar;
  title 'return summary';
run;

proc univariate freq data=holding_period;
  var month_counter;
  title 'holding period';
run;

```