Time-Series Earnings Expectations and Post-Earnings-Announcment-Drift: A Computational Model

By Eric Weisbrod

AML 591 Final Project

Dr. Marco Janssen

December 13, 2010

1. Introduction

Dickhaut and Xin (2009) (hereafter, DX), note that "For 40 years, accountants using archival data have sought an understanding of the relationship between stock price behavior and information (and, in particular, accounting information" (DX, p. 1805). Specifically, one line of research seeks to determine whether stock prices fully reflect the implications of current earnings for future earnings. Findings using archival stock market and earnings data suggest that stock prices behave *as if* investors do not fully understand the quarterly earnings-generating process. However, since individual investor behavior is unobservable, such inferences about investor behavior are based solely upon aggregate market data. This paper uses an individual-based computational model (also referred to as an agent-based model, or ABM) to examine whether the failure of individual investors to correctly incorporate the time series properties of earnings into their forecasts is a plausible explanation for observed stock price behavior. The findings from the simulation are consistent with observed archival stock price behavior, and suggest that only a subset of market participants need behave irrationally to generate the observed aggregate price behavior. The remainder of the paper proceeds as follows: Section 2 elaborates on the research question and discusses relevant literature, Section 3 describes the model, Section 4 presents the model analysis, and Section 5 concludes.

2. Research Question and Relevant Literature

As cited above, archival research in accounting and finance has investigated the relationship between stock prices and earnings information for over 40 years. The efficient market hypothesis states, in its strongest form, that security prices fully reflect all available information (Fama 1991). However, researchers acknowledge that the strong form of this hypothesis is unlikely to hold in the presence of market frictions such as trading costs and information asymmetry. A weaker form of the efficient market hypothesis, *semi-strong* efficiency predicts that, given the incentives for arbitrage in the marketplace, no *profitable* (e.g., after trading costs) trades should be achievable subsequent to the release of *public* information. Beginning with Ball and Brown (1968), accounting research has examined the efficiency of market prices with respect to public earnings announcements. While Ball and Brown (1968) predict and find a positive association between annual earnings changes and annual stock returns, their results also suggest that stock prices tend to "drift" in the same direction as earnings news subsequent to the public earnings release. This result represents an anomaly to the efficient market hypothesis known as post-earnings-announcement-drift (PEAD). Namely, given the findings in Ball and Brown (1968), a rational investor could execute a profitable trading strategy by purchasing (shorting) stocks subsequent to public good (bad) news earnings announcements and holding them for a few months as prices continue to drift upwards (downwards). Furthermore, despite being documented in Ball and Brown (1968), as well as in numerous subsequent published studies (See Kothari 2001 for a review), PEAD continues to be observed in recent market data (Nichols and Wahlen 2004).

Finding an explanation for market inefficiencies such as PEAD is a relevant research question for both market participants and researchers. Market regulators are concerned with providing fair and efficient capital markets for investors. As mentioned above, market inefficiencies are also relevant to investors because they represent profitable arbitrage opportunities. Finally, tests of market efficiency are important to researchers attempting to understand financial market behavior.

While numerous accounting studies have attempted to explain the presence/persistence of PEAD (Again, see Kothari 2001), most studies make inferences about micro-level investor behavior from macro-level archival data. DX describe the issue, "…[N]amely, that there is no 'individualistic' explanation of how prices in archival data are formed. Accountants show semi-strong form efficiency or lack thereof by using sequences of prices from Center for Research in Security Prices (CRSP) and Compustat data, but without a model of how these prices arise from individual decisions (DX, p. 1806). In an attempt to begin to bridge this gap, DX present a general, individual-based, computational model of price formation with risky assets to give insight into archival market-pricing results, such as PEAD. DX show that, under their double-auction framework, post-announcement drift, such as that reported by Ball and Brown (1968), can occur when some individuals in the marketplace possess private information, and this private information is correlated, but not perfectly, with public information. While the results in DX represent one possible individual-based explanation for PEAD, the authors note that their computational model offers "an entire framework to examine theoretical predicions across a wide set of parametric specifications and contexts" (DX, p. 1806).

Thus, I extend the framework introduced by DX to examine a more specific explanation for PEAD from the accounting literature. Beginning with Bernard and Thomas (1990) studies attempting to explain PEAD have provided evidence that the abnormal returns following earnings announcements appear to be consistent with investors naively forecasting earnings as following a seasonal random-walk (SRW), while, in reality, the earnings-generating process more closely resembles a seasonal martingale process (Brown and Rozeff 1979). In support of such archival evidence, Maines and Hand (1996) find, in an experimental setting, that subjects appear to underestimate the autoregressive and fourth-quarter moving average components of

Brown and Rozeff (1979) model. Brown and Han (2000) further extend this stream of literature by examining archival market data and showing that, even for firms whose historical earnings-generating process closely resembles a simple AR(1) autoregressive pattern (as opposed to the more complex Brown-Rozeff pattern), stock prices appear to behave as if investors incorrectly assume that earnings follow a SRW pattern.

By extending the DX framework to model a market in which investors trade equity shares for a firm with earnings generated following an AR(1) pattern, I bridge the gap between the experimental evidence provided by Maines and Hand (1996) and the archival data examined in Brown and Han (2000). Specifically, while Maines and Hand (1996) find that individual investors often fail to incorporate the historical time-series properties of earnings into their future earnings *forecasts*, they do not examine how these forecasts affect market *prices*. On the other hand, Brown and Han (2000) examine market *prices*, and find abnormal returns consistent with investors basing their valuations on incorrect forecasts, but their findings are based on a number of assumptions due to the fact that actual investor forecasts are *unobservable* in their archival setting.[1] The purpose of the this study is to demonstrate an individual-based model of price-formation whereby individual time-series forecast errors might translate into aggregate pricing behavior approximating that observed in archival data.

3. Description of the Model

[1] Brown and Han (2000) acknowledge that their hypothesis assumes that "the researcher knows (and the market reflects) the true quarterly earnings-generating process; the researcher and the market obtain identical estimates of the time-series model's parameters; the researcher correctly measures [cumulative abnormal returns]; all the valuation implications of [the current-period earnings surprise] are captured by the [measured cumulative abnormal return]; [the earnings surprise] is a permanent shock; and the only information available to the market at time t is the 'earnings surprise…'" (Brown and Han 2000, footnote 5).

My model follows the double-auction framework presented in DX, modified to examine a situation in which heterogeneous investors trade in the equity of a firm whose earnings follow an AR(1) pattern. The substantive features of my model, as well as the differences between my model and that presented in DX are as follows (a more complete description of the model, following the ODD protocol, appears in the appendix):

1. DX model a double-auction market for a single risky asset, which will eventually pay out to investors one of two discrete amounts (high or low) denominated in a riskless asset. I also model a single risky asset, however, the risky asset represents a firm that generates earnings and is expected to pay dividends to its investors at some point in the future. Thus firm value is determined as the present value of expected dividends (PVED), as in the savings account analysis in Ohlson (2009). Unlike DX, the expected dividend payout is not constrained to be one of two discrete values.
2. In DX, all investors value the risky asset using the same CARA utility function, with individual investors possessing heterogeneous risk coefficients. I also implicitly assume that investors exhibit constant absolute risk aversion, but I abstract investors' heterogeneous risk preferences by modeling investors as valuing the firm using the same PVED valuation model, with heterogeneous discount rates.
3. While DX model a single information release, where the information signal represents the probability that the risky asset will pay the high payoff (as opposed to the low payoff), I model a periodic information release, where each information signal represent a quarterly earnings amount, and quarterly earnings are generated following an AR(1) process with random-normal stochastic earnings shocks.

4. In DX, investors have varying expectations about the eventual payoff of the risky asset because some investors observe private information signals, while others possess only public information signals. In my model, all agents observe the same publicly available information. However, some investors, given the publicly-observable historical earnings series, forecast future earnings using an AR(1) expectation model with the same autoregressive coefficients used in the true earnings-generation process, while other investors naively forecast future earnings using a SRW expectation model.

All other aspects of the model are substantively identical to those in DX. Following DX, there are four basic steps in the model. In step 1, each agent determines his or her valuation for the traded shares. In step 2, the agent uses market information to assess the likelihood of being able to trade. In step 3, the agent takes an optimal trading action, and step 4 describes the stochastic arrival of agents to the market. I will briefly elaborate on each step.

Step 1: Valuation

At each time step, each trader arrives at his or her assessment of the reservation value of the traded asset according to the following valuation formula:

$$V_{i,t} = \frac{E_i[x_{t+1} + x_{t+2} + x_{t+3} + x_{t+4}]}{r_i} \tag{0.1}$$

where $V_{i,t}$ is the value of the asset to agent i at time t, x_{t+j} is the forecasted earnings signal at quarter $t+j$ and r_i is the rate used by agent i to discount expected future dividends. Note that each trader comes to the market with an individualized risk attitude in the form of a discount rate. Furthermore, depending on the parameters of the simulation run, some traders will form their expectations of future earnings using an AR(1) forecast, and some will form their expectations using a SRW forecast.

Step 2: Agent Computations of Likelihoods

At any point in time *t* there is a history of bids, asks, and trades in the marketplace. In this step, traders use this market history to determine the likelihood that a particular bid or ask will be accepted in the marketplace. I assume that traders have limited memory. Using his or her valuation computed in the previous step, combined with this estimated likelihood, each agent determines the action that will maximize his or her expected surplus. As noted by DX:

> While a trader knows his or her value of the asset, he or she knows nothing about the risk distribution of other trader types, nor how the most recent market observations reflect the consequence of differential information or risk-sharing. Thus, every trader bases his or her assessment on the likelihood of a particular bid (ask) being accepted on recent observed information in the market about whether that bid (ask) will be successful or unsuccessful. We assume, because the trader simply has such limited information and no basis for as- signing priors, that the only information on which the trader can base his or her action regarding a bid or ask is the recent prices, bids, and asks. (DX, p. 1812)

For bids and asks previously observed in the marketplace, traders use a counting system to determine the likelihoods of that bid or ask being accepted at time *t*, and for previously unobserved bids and asks, agents determine likelihoods based on linear interpolation.

Step 3: Determination of Optimal Action

In step 3, each trader combines the valuation from step 1 with the estimated likelihoods from step 2 to determine the expected profit for each possible action. The expected profit for particular agent i, with a valuation V_i, for bid B, will be $P(B)(V_i - B)$ where $P(B)$ is the estimated probability that bid B will be accepted in the marketplace. Likewise, the expected profit for agent i, with a valuation V_i, for ask A, will be $P(A)(A - Vi)$ where $P(A)$ is the estimated likelihood that ask A will be accepted in the marketplace.

Following DX, Choice calculation proceeds as follows: the agent behaves as if doing four calculations:

$$\underset{B}{Max}\, P(B)(V_i - B),$$

$$\underset{A}{Max}\, P(A)(A - V_i),$$

$$V_i - OB,$$

$$OA - V_i,$$

where OB (OA) is the existing outstanding bid (ask). Determining the maximum profit from these calculations leads to the selection of an action that will be a particular B, A, OB, or OA, where the maximum profit is greater than 0. Each subject can make a bid, make an ask, or take an outstanding bid or ask. Thus, he or she may be a buyer or a seller. If none of these actions would be profitable, the trader may do nothing.

Step 4: Whose Bid/Ask is Posted First – A Decentralized Process

The double auction market mechanism acknowledges only one action of only one player. Once this action is incorporated into the auction all traders are notified, and the trading process begins again. Thus, the determination of who moves in the auction is important in the simulation. Following DX:

> It is crucial that the auction be decentralized if prices are a consequence of individual choice. This means that we need a behavioral theory of how such decentralization takes

> place. No single person (or mechanism) knows all profits or valuations of each actor. We assume that on average, the more profit a single player makes, the more likely he or she will enter the auction; that is, higher profits to a player induce that player to act faster. Interpreting this proposition strictly would imply that the person with the highest expected profit would always move first. However, we assume some noise in the response, which could be attributed to the circumstances of individual players.

Thus, the determination of who acts first in the market is made using a “roulette wheel” mechanism where the probability that a traders action will be accepted in the marketplace is proportional to the trader’s expected profit from the action. For further discussion of the behavioral motivation and related literature which form the foundation for this selection mechanism, see DX.

4. Model Analysis

I simulate the market described above using the Netlogo software package (Wilensky 1999). In every run of the simulation, there are a fixed number of 10 total traders, and a new earnings signal is announced every 500 timesteps. Each run of the model is simulated for a total of 4500 timesteps, representing a series of 8 consecutive earnings announcements (the outstanding earnings value at the beginning of the simulation is not analyzed, to give time for an equilibrium to develop in the marketplace). The parameter of interest, which is varied from simulation to simulation, is *N_SEASONAL*, which represents the subset of the ten total traders who make forecasts using a SRW forecast model (as opposed to an AR(1) model). *N_SEASONAL* is analyzed at 0, 3, 5, 7, and 10, out of 10 total traders in the marketplace. Figure 1 depicts the pricing patterns from runs of the simulation at each value of *N_SEASONAL*, using the same earnings time-series.[2]

I perform 100 runs of the simulation for each examined value of *N_SEASONAL*. Following Brown and Han (2000), I analyze the cumulative abnormal returns around each simulated earnings announcement. To motivate their regression analysis, Brown and Han (2000) present two possible scenarios of price formation. In scenario 1, stock prices fully reflect the implications of current quarterly earnings for future quarterly earnings, and investors are aware of the correct autoregressive parameters underlying the earnings time series. In this scenario, the market's expectation of quarter t's earnings as of quarter $t-1$ is:

$$E(X_t \mid X_{t-1})^M = \phi Q_{t-1} \tag{0.2}$$

Where E is the expectation operator and E^m designates the capital market's expectation. Following Brown and Han (2000), in scenario 1 cumulative abnormal returns relate to the contemporaneous and lagged errors of the AR(1) model as follows:

$$\begin{aligned} CAR_t &= \lambda[X_t - E(X_t \mid X_{t-1})^m] \\ &= \lambda[\phi X_{t-1} + e_t - \phi X_{t-1}] \\ &= \lambda e_t \end{aligned} \tag{0.3}$$

where CAR_t is the abnormal return at the time of the quarterly earnings announcement t, λ is a positive multiplier, and e_t is the random-normal shock at time t from the true earnings-generating process.

[2] While Figure 1 depicts numerous runs of the simulation using the same earnings time-series to provide a visual anchor, the earnings time-series is allowed to vary stochastically, following an independent AR(1) process, over each run of the simulation used in the regression analysis that follows.

In scenario 2, stock prices do not fully reflect the implications of current earnings for future earnings of AR(1) firms. Brown and Han (2000), illustrate price formation process if quarterly earnings follow an AR(1) process, but investors believe the quarterly earnings-generating process is a SRW. In this case, the market's expectation of quarter t's earnings as of quarter $t\text{-}1$ is $E(X_t|X_{t-1})^M = X_{t-4}$. Thus, the abnormal return at quarter t relates to the contemporaneous error and the first four lagged errors of the AR(1) model as follows:

$$\begin{aligned} CAR_t &= \lambda[X_t - E(X_t \mid X_{t-1})^m] \\ &= \lambda[\phi X_{t-1} + e_t - X_{t-4}] \\ &= \lambda[e_t + \phi e_{t-1} + \phi^2 e_{t-2} + \phi^3 e_{t-3} + (\phi^4 - 1) e_{t-4}] \end{aligned} \tag{0.4}$$

As noted by Brown and Han (2000), "Consistent with both scenarios CAR_t is positively related to the contemporaneous error of the AR1 model. However, consistent with scenario 2 only, e_{t-1}, e_{t-2}, and e_{t-3} have positive, but decreasing, coefficients; and e_{t-4} has a negative coefficient. The coefficients of e_{t-2} and e_{t-3} approach zero from above as the ϕ parameter is squared and cubed, respectively." (Brown and Han 2000, footnote 8).

Accordingly, Brown and Han (2000) estimate the following regression model to analyze the associations between CAR_t and the current and lagged earnings forecast errors:

$$CAR_t = \alpha_0 + \beta_0 e_t + \beta_1 e_{t-1} + \beta_2 e_{t-2} + \beta_3 e_{t-3} + \beta_4 e_{t-4} + \varepsilon_t \tag{0.5}$$

Where:

CAR_t = three-day (-2,0) size-adjusted abnormal return, and day 0 is the quarter t earnings announcement date,

α_0 = intercept,

β_j = the multiplier on the earnings forecast error for quarter $t-j$ based on the AR(1) model, for j = 0 to 4,

e_{t-j} = the earnings forecast error for quarter $t-j$ based on the AR(1) model, for j = 0 to 4, and

ε_t = the residual error of the regression model.

Thus, for example, if stock prices are based on the erroneous assumption that the quarterly earnings-generating process of AR(1) firms is a seasonal random walk, the coefficient of the lag four error term, β_4, will be negative.

At each examined value of *N_SEASONAL*, I estimate the regression in equation 1.5 on the simulated data by calculating CAR_t as follows: First, I assume that a trading day can be approximated by 100 timesteps in the model. Second, I calculate a daily price by taking the mean price over 100 timesteps. Then, I calculate a (-2, 0) return by taking the difference between the average price for the 100 timesteps following a simulated earnings announcement and the average price over time -200 to -100 relative to the earnings announcement, and dividing by the earlier average price. For example, when the first earnings announcement occurs at t = 500, I calculate CAR500 as: $(AVG_{500,600}(Price) - AVG_{300,400}(Price))/ AVG_{300,400}(Price)$. While I label this return as a cumulative *abnormal* return for comparison with Brown and Han (2000), my calculation actually represents a simple raw return. This is appropriate because, in my simulated setting, there is only one firm in the marketplace and the sole source of information in the economy is the earnings time-series. Thus, no benchmark return is necessary for extracting the firm-specific "abnormal" portion of the return. Each run of the simulation produces eight earnings announcement observations, and the simulation is run 100 times for each value of *N_SEASONAL*, resulting in a total of 800 observations in each regression.

Table 1 presents the results of my regression analysis alongside the original results from Brown and Han (2000). When estimating equation 1.5 for their full sample of firms, Brown and Han (2000) find a positive and significant coefficient on β_0, as well as a negative and significant coefficient on β_4, consistent with scenario 2, where investors irrationally forecast earnings as following a SRW pattern. In my simulation, for values of *N_SEASONAL* less that 5, I find a positive and significant coefficient on β_0, as predicted by Scenario 1 of Brown and Han (2000). However, inconsistent with the predictions for *either* scenario 1 or scenario 2 in Brown and Han (2000), I find significant *negative* coefficients for $\beta_1, \beta_2, \beta_3$ and β_4 (with the slight exception of *N_SEASONAL* = 3, where β_3 is insignificant). I am unable to explain this anomalous result, but may be able to find a suitable explanation after more thought. In contrast to the results for values of N_SEASONAL less that 5, I find results consistent with those predicted by scenario 2 of Brown and Han (2000) for values of *N_SEASONAL* greater than or equal to five. Specifically, I find a positive and significant coefficient on β_0, positive, significiant, and decreasing coefficients on β_1, β_2, and β_3, and a significant negative coefficient on β_4.

These results indicate that the findings from the archival data in Brown and Han (2000) could, plausibly, be the result of investors making trades based on irrational time-series earnings forecasts. Importantly, my results provide evidence that only a subset, and not a majority, of investors in the marketplace need behave irrationally for the PEAD pricing behavior to occur.

5. Conclusion

While much of parameter space of my model has not yet been analyzed for sensitivity, the preliminary results from the model lead to useful conclusions. Proponents of the efficient market hypothesis are often skeptical of results that suggest that market prices behave irrationally, because they believe that, even if some market participants behave irrationally at any point in time, it is unlikely that the majority of investors would exhibit coordinated irrational behavior, and, that rational traders, in any proportion of the population, will quickly arbitrage away any irrational pricing in the marketplace. The results of my simulation suggest that, under a plausible set of assumptions governing a decentralized market scenario, irrational pricing behavior can result even from a subset, and not a majority, of traders behaving irrationally. Furthermore, in this setting, the rational traders in the marketplace are not able to fully arbitrage away the irrational pricing. My findings extend those in Brown and Han (2000) by linking aggregate pricing data to individual investor behavior. My findings also extend those reported by DX, along numerous dimensions.

First, I extend the double-action framework from DX to analyze multiple periodic information signals, and modify the information signals and valuation functions from DX to resemble a more realistic setting in which investors trade in the equity of a firm, based on the firm's earnings information. Second, I demonstrate how return data can be calculated and analyzed in the DX market framework. Finally, DX demonstrate that pricing drift, similar to that observed in archival data, may result from fully rational traders possessing asymmetric information. My findings demonstrate an alternative explanation for pricing drift, under which all investors possess identical public information, but some investors fail to fully incorporate the time series implications of current earnings for future earnings as expected under fully rational expectations. In this regard, my findings are consistent with those from behavior research indicating that human investors are only boundedly rational (e.g. Maines and Hand 1996).

My study can be extended in a number of ways. First, investors can be modeled using a more diverse set of assumptions. For example, future research might model investors with relative, as opposed to absolute, risk aversion, such that investors' endowment plays a role in their investment decisions. Investors could also be modeled using heterogeneous valuation functions, mixing the simple, capitalized forward earnings valuation used in this study with residual income valuation and abnormal earnings valuation models discussed in Ohlson (2009). Another interesting extension would be to model investors with the potential for learning. That is, investors could each form their own dynamic, fully or boundedly rational, expectations about future earnings from the historical earnings time series, using artificial intelligence or genetic algorithms. Future research could also examine pricing behavior in a setting in which investors can choose an investment portfolio among multiple firms in the marketplace, as well as pricing behavior in a market setting incorporating realistic information costs, trading costs, and short selling constraints. As noted by DX, their framework can be extended "across a wide set of parametric specifications and contexts" (DX, p. 1806).

References

Ball, R., & Brown, P. (1968). An empirical evaluation of accounting income numbers. *Journal of Accounting Research, 6*(2), pp. 159-178.

Bernard, V. L., & Thomas, J. K. (1990). Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics, 13*(4), 305. doi:DOI: 10.1016/0165-4101(90)90008-R"

Brown, L. D., & Han, J. C. Y. (2000). Do stock prices fully reflect the implications of current earnings for future earnings for AR1 firms? *Journal of Accounting Research, 38*(1), 149-164. doi:http://www.blackwellpublishing.com/journal.asp?ref=0021-8456&site=1

Brown, L. D., & Rozeff, M. S. (1979). Univariate time-series models of qtrly accounting earnings per share: A proposed model. *Journal of Accounting Research, 17*(1), 179-189. doi:http://www.blackwellpublishing.com/journal.asp?ref=0021-8456&site=1

Dickhaut, J., & Xin, B. (2009). Market efficiencies and drift: A computational model. *The Accounting Review, 84*(6), 1805.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance, 25*(2, Papers and Proceedings of the Twenty-Eighth Annual Meeting of the American Finance Association New York, N.Y. December, 28-30, 1969), pp. 383-417.

Fama, E. F. (1991). Efficient capital markets: II. *Journal of Finance, 46*(5), 1575-1617. doi:http://www.blackwellpublishing.com/journal.asp?ref=0022-1082

Grimm, Volker, Uta Berger, Finn Bastiansen, Sigrunn Eliassen, Vincent Ginot, Jarl Giske, John Goss-Custard, Tamara Grand, Simone K. Heinz, Geir Huse, Andreas Huth, Jane U. Jepsen, Christian Jorgensen, Wolf M. Mooij, Birgit Muller, Guy Pe'er, Cyril Piou, Steven F. Railsback, Andrew M. Robbins, Martha M. Robbins, Eva Rossmanith, Nadja Ruger, Espen Strand, Sami Souissi, Richard A. Stillman, Rune Vabo, Ute Visser, Donald L. DeAngelis. (2006). A standard protocol for describing individual-based and agent-based models, *Ecological Modelling*, 198(1-2) 115-126

Kothari, S. P. (2001). Capital markets research in accounting. *Journal of Accounting and Economics, 31*(1-3), 105. doi:DOI: 10.1016/S0165-4101(01)00030-1"

Maines, L. A., & Hand, J. R. M. (1996). Individuals' perceptions and misperceptions of time series properties of quarterly earnings. *The Accounting Review, 71*(3), pp. 317-336.

Nichols, D. C., & Wahlen, J. M. (2004). How do earnings numbers relate to stock returns? A review of classic accounting research with updated evidence. *Accounting Horizons, 18*(4), 263.

Ohlson, J. A. (2009). Accounting data and value: The basic results. *Contemporary Accounting Research/Recherche Comptable Contemporaine, 26*(1), 231-259. doi:http://caaa.metapress.com.ezproxy1.lib.asu.edu/link.asp?id=300305

Wilensky, U. 1999. NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected

Learning and Computer-Based Modeling, Northwestern University. Evanston, IL.

Appendix

This model description follows the ODD (Overview, Design concepts, Details) protocol for describing individual- and agent-based models (Grimm et al. 2006, 20XX[3]).

1. Purpose

The purpose of the model is to understand how observed price paths from public stock markets might arise from individual trading decisions. Specifically, the model examines cases in which heterogeneous individual traders trade in the equity of a firm based on forecasts of future earnings, using either a rational or irrational earnings forecast methodology. The pricing behavior of the model is compared to observed archival pricing behavior to determine whether irrational earnings forecasts are a plausible explanation for the observed archival behavior.

2. Entities, state variables, and scales

The model is comprised of two types of entities, traders and a market institution. There are two types of assets in the model, a riskless asset (denominated as $) and a risky asset. The risky asset is modeled as the equity of a hypothetical firm. The model is also defined by state variables representing information (earnings) signals about the expected payoff (dividends) of the risky asset (traded firm).

Market Variables:

- Observed History – A record of bids, asks, and takes which have occurred in the market.

Risky Asset:

- Earnings History – The earnings history is generated from a quarterly earnings seed of $5 plus a random floating-point amount of noise generated between 0 and $2.50. Subsequent to the initial seed value, quarterly earnings are generated as $X_{t+1} = 2.5 + .8X_t + \varepsilon_t$, where X is earnings for quarter t, and ε_t is a random-normal shock distributed N(0,1). During simulation setup, an earnings history of length 20 is generated to initialize the market. Subsequently, the earnings history is updated with a new quarterly earnings observation every 500 timesteps.

Trader Variables:

- Wealth – A number representing units of the riskless asset

- Risk Parameter – A discount rate randomly selected from a normal distribution with the parameters N(.08, .005). For the regression analysis in this study, ten risk parameters were initially drawn from this normal distribution, and this pool of 10

[3] This is a reference to the ODD "first update" manuscript.

parameters was held fixed throughout the simulations. However, the 10 values were randomly assigned among the 10 traders in each simulation run. The 10 values used in my analysis are: 0.0835 0.0829 0.0718 0.0869 0.0846 0.0815 0.0855 0.076 0.0918 0.0797 0.0834.

- Forecast methodology – A binary variable representing the trader's forecast methodology. The variable indicates that the trader uses either an AR(1) or a seasonal random walk (SRW) forecast methodology.
    - Under the AR(1) forecast methodology, investors forecast quarterly earnings for quarter t as $E(X_t|X_{t-1}) = 2.5 + .8X_{t-1}$
    - Under the SRW forecast methodology, investors forecast quarterly earnings for quarter t as $E(X_t|X_{t-1}) = X_{t-4}$
- Memory Length – The number of preceding trades/offers that an agent can remember.

Time is denoted by iterations of the model, where one time step denotes one trading action occurring in the market, such as a bid or ask being offered or accepted. Thus, time is not absolute in the sense of minutes or hours, such that, when making analogies to real-world markets, one step of the model in a liquid, electronic market, might occur within microseconds of real time, while in an illiquid market, each trading step might be separated by days, months, or even years. For purposes of the regression analysis, one trading day is approximated as 100 timesteps.

3. Process overview and scheduling

In each iteration of the market, agents each perform three procedures independently (and concurrently), before any communication occurs:

Step 1: Valuation

Step 2: Agent Computations of Likelihoods

Step 3: Determination of Optimal Action

At this point, all agents submit their optimal action to the market institution. The market institution acknowledges only one action of only one trader in each iteration of the market. Whose bid/ask is posted by the market is a stochastic process (described more fully below) based on each trader's expected profit relative to the total expected profits of all traders in the market, such that, higher expected profits induce a trader to act faster.

Once the market mechanism has selected an action to be acknowledged, this action is incorporated into the observed history of the market. All traders are notified of the selected action, and bidding begins again from step 1.

4. Design concepts

Basic principles: In the model, agents follow expected utility theory given the forecasts that they make using publicly available information. However, some agents may make irrational forecasts. All agents are assumed to have the same constant absolute risk aversion utility function, but there is no conceptual restriction on how risk is incorporated. The double auction market mechanism also follows a basic principle that bids and asks are subject to an improvement rule; that is, if there is an outstanding bid and ask at time t, agents may only submit a t+1 bid that is higher or a t+1 ask that is lower than the outstanding bid or ask. Finally, the timing over which bids and asks are accepted by the market mechanism is governed by a general theory of reaction times stating that higher rewards stimulate faster reaction times.

Emergence: Price pattern and volatility of the market.

Adaptation: Individuals calculate the likelihood of their bid/ask being accepted based on the history of bids/asks/prices in the market, such that agents' current actions are adapted to the history of past actions.

Objectives: Agents seek to maximize their expected utility.

Prediction: Agents incorporate signals about the probability of the payoff of the risky asset into their expectations.

Sensing: Some agents rationally "sense" the implications of the historical earnings pattern for future earnings, while other agents irrationally believe that future earnings will follow a SRW, failing to "sense" the auto-regressive properties of the earnings time-series.

Interaction: All interaction occurs indirectly through the market mechanism. Agents communicate messages to the market mechanism in the form of an ordered triplet of values indicating the identity of the agent, the type of message (bid/ask) and an amount. The market history is a list of accepted messages, such that, for each iteration, one message from one agent is added to the market history, and the history is re-broadcast to all agents.

Stochasticity: Risk parameters are stochastically assigned to traders, and the process determining which bid/ask is accepted by the market mechanism at time *t* is stochastic.

Observation: Outstanding bids and asks, as well as accepted trade prices, are collected for observation. The earnings history is also collected for analysis.

5. Initialization

In all scenarios, the economy is populated with 10 agents. Agents are randomly assigned different risk coefficients, but across economies, the collection of risk coefficients is held constant, as described above. The earnings history is also initialized as described in Section 2 above.

6. Input data

"The model does not use input data to represent time-varying processes."

7. Submodels

Remember-history

- This process takes the market history as an input, and annotates the market history to reflect the limited memory of the investors. For example, if the memory length is set to 5, then, if the market history is longer than five trades, this process counts backwards through the past 6 trades in the market, and deletes all information prior-to and including, the 6^{th} previous trade. The process also deletes the outstanding bid and ask from the market history, as it is unknown whether these bids and asks will be accepted or rejected.

Fill-dlist

- This process takes the remembered history as an input, and tallies number of bids, asks, and rejected/accepted bids and asks for each monetary value observed in the remembered history. This process also counts the number of asks (bids) and rejected asks (bids) less (greater) than the each value, and the number of taken asks (bids) greater (less) than each value.

Find-ask & find-bid

- These processes use the counts collected by fill-dlist to calculate the likelihood that a given bid or ask will be accepted in the marketplace. Likelihoods are calculated as defined by Dickhaut and Xin (2009) using the counting process described by Dickhaut and Xin (2009) for previously observed values and linear interpolation for previously unobserved values. Probabilities are calculated for all possible bids and asks between the outstanding bid and ask, in $0.01 increments. Agents then iteratively evaluate the expected profit at each $0.01 trading increment and find-ask/find-bid report the ask and bid with the highest expected trading profits, respectively. This iterative solution of evaluating every $0.01 "ticksize" may result in different model behavior than if agents were allowing to maximize their expected profits using linear optimization over a continuous trading space, but continuous linear optimization has not been evaluated.

Build-expectations

- This process is executed by agents and calls find-ask and find-bid as sub-processes. Agents then use the expected profits from find-ask and find-bid, as well as the certain profits that would result for accepting the outstanding bid or ask. Agents then choose the profit maximizing action, or choose to take no action at all. If agents take an action, this process formats the action into a market message triplet and submits the action to the market mechanism.

Send-message

- This process is called by the market mechanism. It collects messages from each trader, along with each trader's expected profit from the message being accepted by the market. The process then uses a "roulette wheel" algorithm, where the probability that a message will be accepted is proportional to the expected profit to trader who generated the message, to randomly select one message to add to the market history. The process then appends the message to the market history, which will be visible to all traders in the next time step, as the form their valuations and estimate likelihoods.

Figure 1: Pricing History over time, for fixed simulated earnings time-series simulated at various values of *N_SEASONAL*

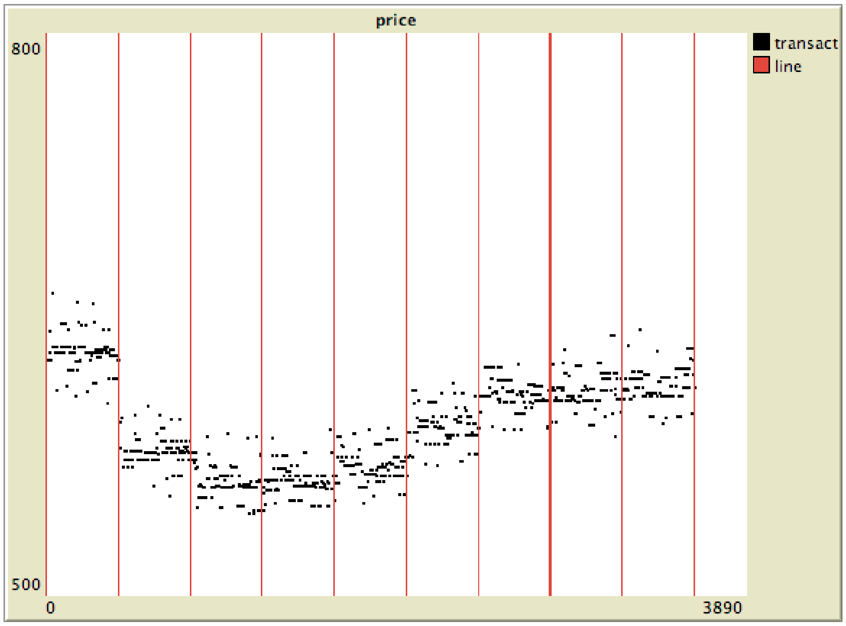

n=10

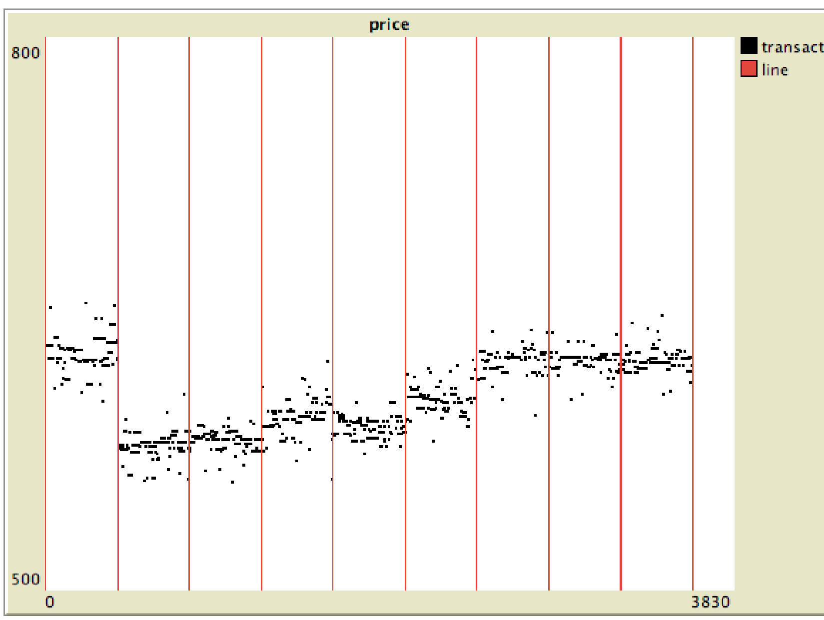

n=7

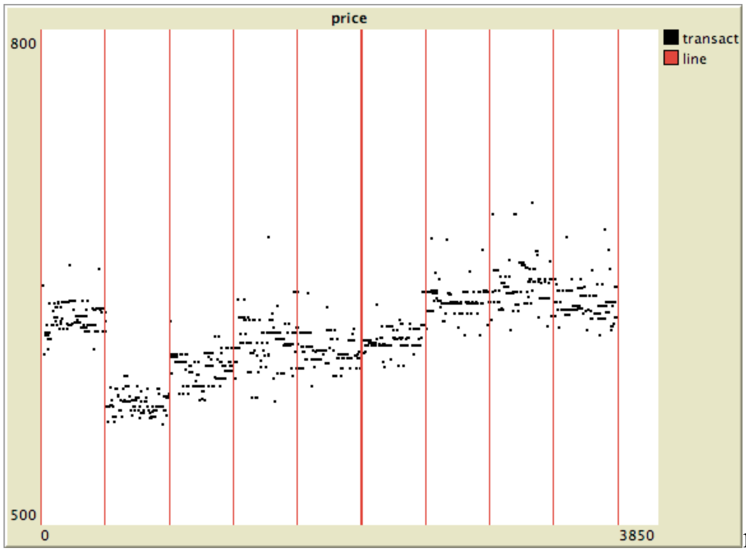
price
800
500
0
3850
transact
line
n=5

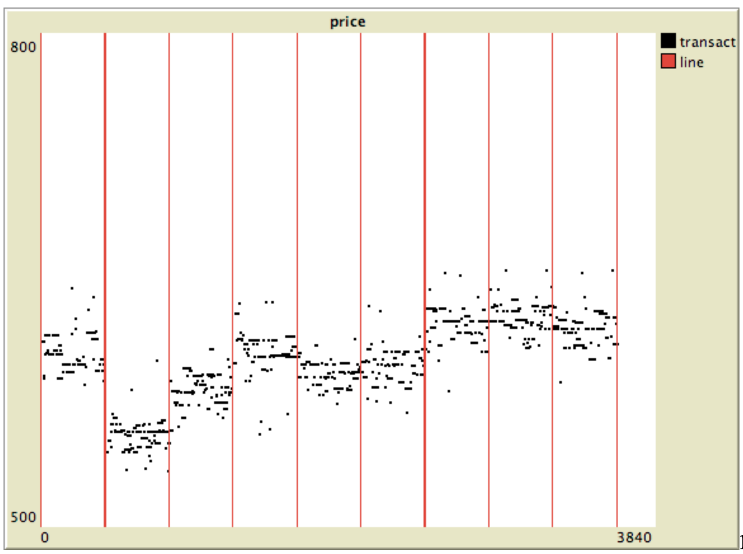
price
800
500
0
3840
transact
line
n=3

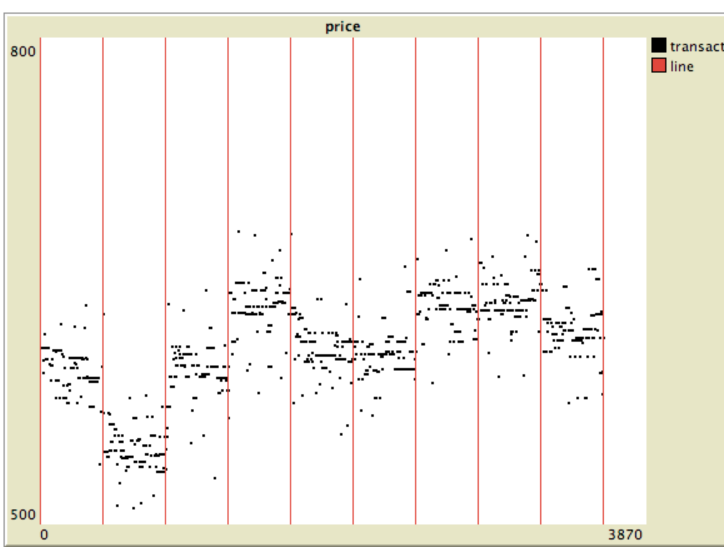
price
800
500
0
3870
transact
line
n=0

Table 1:
Regression Results: OLS Regressions of
Cumulative Abnormal Returns on Current and Previous Earnings Surprises

$$CAR_t = \alpha_0 + \beta_0 e_t + \beta_1 e_{t-1} + \beta_2 e_{t-2} + \beta_3 e_{t-3} + \beta_4 e_{t-4} + \varepsilon_t$$

	Predicted Sign	*BH2000* Estimate (t-stat)	Number of Seasonal Random Walk (SRW) Investors (out of 10 total investors): *0* Estimate (t-stat)	*3* Estimate (t-stat)	*5* Estimate (t-stat)	*7* Estimate (t-stat)	*10* Estimate (t-stat)
α_0		0.002 ***	0.002 ***	0.004 ***	0.003 ***	0.005 ***	0.005 ***
		(2.71)	(5.06)	(8.57)	(5.31)	(8.54)	(7.95)
β_0	+	0.012 ***	0.535 ***	0.436 ***	0.381 ***	0.339 ***	0.230 ***
		(12.11)	(109.43)	(75.66)	(57.64)	(53.09)	(31.07)
β_1	+	0.001	-0.114 ***	-0.021 ***	0.027 ***	0.086 ***	0.178 ***
		(1.50)	(-23.58)	(-3.66)	(4.04)	(13.57)	(24.44)
β_2	+	0.000	-0.095 ***	-0.025 ***	0.022 ***	0.075 ***	0.141 ***
		(-0.03)	(-19.26)	(-4.20)	(3.35)	(11.74)	(19.03)
β_3	-	-0.001	-0.073 ***	-0.008	0.019 ***	0.058 ***	0.110 ***
		(-1.38)	(-14.84)	(-1.32)	(2.98)	(8.97)	(14.51)
β_4	-	-0.004 ***	-0.058 ***	-0.076 ***	-0.105 ***	-0.115 ***	-0.148 ***
		(-3.91)	(-11.98)	(-13.49)	(-16.34)	(-17.88)	(-19.89)
Number of Observations		**5,281**	**800**	**800**	**800**	**800**	**800**
Adjusted R^2		**0.028**	**0.943**	**0.881**	**0.819**	**0.823**	**0.771**

Variables are defined in Section 4 of the text.

*, **, and *** indicate statistical significance at the 10%, 5% and 1% levels, respectively, based on two-tailed t-tests.