Microconsistency in Simple Empirical Agent-based Financial Models

Blake LeBaron *

International Business School

Brandeis University

February, 2016

Abstract

Models with small numbers of agents have recently been simplified for direct empirical estimation. Parameters are estimated at the macro level to get a best fit to the data, but usually little analysis is done at the micro level to examine the choices made by agents for forecasting rules. This paper explores one of these recent models from the standpoint of micro agent behavior. It is shown that at the fitted forecasting rules, agents would prefer deviating to other nearby rules. The simple two type model is then compared with several multi-type models allowing for agents to use a broader set of rules. This can impact the dynamics of the generated time series, but it also may not if one takes the parameter estimates of the original model as an exogenous restriction on a reasonable support for the forecasting rules. This emphasizes that these models may be imposing some hidden micro assumptions about agent behavior.

Keywords: Learning, Heterogeneous agent models, Asset Pricing, Financial Time Series, Adaptive behavior

---

*415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. This research is partially supported by funding from the Institute for New Economic Thinking, INET.

1 Introduction

Modeling asset markets with heterogeneous agents and beliefs still is a field in relative infancy with many approaches, and many results coming from researchers in a wide variety of disciplines. While it is accepted that the world is populated with agents possessing a variety of beliefs about the future, it is not clear what impact, if any, this will have on asset pricing. Most heterogeneous agent modeling suggests that the impact is large, and may be an important driver of what has appeared to be irrational levels of market volatility and trading volume. However, heterogeneous agent models, by design, imply large numbers of parameters for which some are poorly identified. This "too many degrees of freedom" problem has made it difficult to take the models to the data directly. Recently, highly stylized models with small numbers of traders have been built and directly estimated on economic time series. This paper looks at one of the models in detail and compares it to techniques and features which are used in more complex computational models.

Agent-based models in economics and finance focus on relatively simple adaptive behavior where agents follow behavioral rules, but stand ready to adjust their rules in the presence of empirical evidence suggesting that other strategies might be better. This dynamic shifting may be the key component in distinguishing these models. The population of agents is not predetermined to be heterogeneous as in some other heterogeneous agent approaches. The constant shifting of agents makes these models richer, but also more intractable. One style of modeling was to leave the strategies very "free form" and let actual strategies in use emerge from a somewhat nonparametric soup of strategies.[1] Models of this type can be appealing in that most anything that is possible is doable. Two crucial drawbacks for these models are that they are relatively complex, involving many unknown parameters, and require relatively complex computer code to operate. While fitting many features of financial data, they are often viewed with some skepticism.

A much simpler class of models also emerged at about the same time.[2] They are referred to as "few type" models. Their appeal is obvious. Working with stripped down strategy spaces they employ as few as two forecasting rules to describe financial markets. One can often implement these models with a small set of equations, and get a handle on some of the analytics driving their dynamics. At first these models offered only analytical tractability, and lighter demands for computer time. However, in an important paper Boswijk, Hommes & Manzan (2007) show that a simple two agent model could yield an equivalent, and relatively tractable, time series model which could be directly estimated from data. This changed the model

[1] In LeBaron (2006) these are refered to as "many type" strategies and many examples of this modeling style are given.

[2] For example, Brock & Hommes (1998), Kirman (1991), Lux (1997). Also, see surveys such as Chiarella, Dieci & He (2009), Hommes (2006), Hommes & Wagener (2009), and Lux (2009), .

building and testing space in a very big way and has been followed by several other papers.[3]

These models appear to be providing a good fit to the data, and a kind of heterogeneous agent benchmark for financial markets. This paper looks at them a little more deeply from a theory perspective. They are built off of macro data estimation, but are relatively quiet about some aspects of individual agent decisions and adaptation. Also, they make some big assumptions themselves to get their tractability. Specifically, agents are assumed to follow the two forecasting rules prescribed by the model, and estimated to best fit the data. This ignores the issue, fundamental to the "many type" world, of letting the agents try to search around the forecasting space. I will explore the space of forecasting models to test the microconsistency of the forecasts.[4]

This paper explores these models from the perspective of understanding what would happen when agents are offered different strategies. Would they want to move? Also, do the dynamics of the model change if the range of strategies available to the agents is changed? The answer to both these questions turns out to be yes. However, there are still some sharp assumptions which can be made that justify the 2-type empirical approach, and these will be discussed.

Section 2 performs some initial simulation and testing for the candidate model which is taken from Hommes & in 't Veld (2014). Section 3 estimates local learning gradients or objective utility functions in local regions around the fitted parameters. It also explores the impact of a multi-agent model operating within the support of the estimated parameters. Section 4 explores the dynamics of models where the two strategies are invaded by a third strategy which is not necessarily local. It also performs some initial time series analysis to see which time series features remain robust to these changes. Section 5 concludes.

2 Model testing and stability

This paper follows Hommes & in 't Veld (2014) and Boswijk et al. (2007) in modeling the movements of asset prices and fundamentals as a two agent system of beliefs. As previously mentioned, this simplification is powerful in that it allows direct estimation of the agent-based parameters.

The price of the risky asset is given by P_t, and it also pays a risky dividend, Y_t each period. The excess

[3] Several examples are Chiarella, He & Zwinkels (2014), Lof (2012), and ter Ellen & Zwinkels (2010).

[4] This question is a old and deep one in agent-based models. The origin is probably Schelling (1978) which stressed that micro and macro optimality may be far apart in many models. There are also some models in the heterogeneous agent world of monetary policy which do stress a form of optimality at the micro level. An example of this is Branch & Evans (2011) which has agents using under parameterized forecasting models, but they are still required to estimate these in a statistically optimal manner.

payout of the risky asset is defined by,

$$R_{t+1} = P_{t+1} + Y_{t+1} - (1+r)P_t \tag{1}$$

where r is a risk free asset. Demands for shares of the asset are determined through standard mean variance preferences and are given by,

$$z_{h,t} = \frac{E_{h,t}R_{t+1}}{a\sigma^2} \tag{2}$$

where a is the coefficient of absolute risk aversion, and σ^2 is the variance of the asset which is assumed to be constant. The fraction of each type in the population is given by $n_{h,t}$. Assuming the asset is available in zero net supply and summing the demands gives,

$$\sum_{h=1}^{H} n_{h,t} \frac{E_{h,t}(P_{t+1} + Y_{t+1}) - (1+r)P_t}{a\sigma^2} = 0, \tag{3}$$

and the corresponding pricing equation,

$$P_t = \frac{1}{1+r} \sum_{h=1}^{H} n_{h,t} E_{h,t}(P_{t+1} + Y_{t+1}). \tag{4}$$

The stochastic process for dividends is common knowledge and is given by a geometric random walk,

$$\log Y_{t+1} = \mu + \log Y_t + \nu_{t+1} \qquad \nu_{t+1} \sim N(0, \sigma_\nu^2) \tag{5}$$

where

$$\frac{Y_{t+1}}{Y_t} = e^{\mu + \nu_{t+1}} \tag{6}$$

and

$$\frac{Y_{t+1}}{Y_t} = (1+g)\epsilon_{t+1}, \qquad (1+g) = e^{\mu + (1/2)\sigma_\nu^2}, \quad E(\epsilon_t) = 1, \tag{7}$$

and given common knowledge, $E_{h,t}Y_{t+1} = (1+g)Y_t$ for all agents h.

The model dynamics are greatly simplified by expressing most pricing relationships in terms of the price/dividend ratio, $\delta_t = P_t/Y_t$. The authors assume that the dividend growth rate is conditionally independent of δ_{t+1} which allows them to write,

$$E_{h,t}\frac{P_{t+1}}{Y_t} = E_{h,t}\{\delta_{t+1}\}E_{h,t}\{Y_{t+1}/Y_t\} = (1+g)E_{h,t}\{\delta_{t+1}\}. \tag{8}$$

This now allows rewriting the pricing equation (4) as,

$$\delta_t = \frac{1}{R^*}\{1 + \sum_{h=1}^{H} n_{h,t} E_{h,t} \delta_{t+1}\}, \qquad R^* = \frac{1+r}{1+g}. \tag{9}$$

Finally, pricing is expressed in differences from the price/dividend ratio determined as present values from the Gordon growth model as in,

$$P_t^* = \frac{1+g}{r-g} Y_t \tag{10}$$

$$\delta^* = \frac{P_t^*}{Y_t} = \frac{1+g}{r-g} \tag{11}$$

$$x_t = \delta_t - \delta^*. \tag{12}$$

In Hommes & in 't Veld (2014) the static Gordon model is replaced with a conditional model allowing r and g to change over time. This is passed through the data to develop a conditional value of for $\delta^* = \delta_t^*$. Since the purpose here is to simply regenerate basic data, and since the Gordon model renormalization can be viewed as separate from the dynamics of x_t there is no need to be concerned with this here. The model will generate x_t, and this can be adjusted to price/dividend ratios by adding back δ^*. Using this adjustment converts equation 9 to,

$$x_t = \frac{1}{R^*}\{\sum_{h=1}^{H} n_{h,t} E_{h,t} x_{t+1}\}. \tag{13}$$

Now that the overall structure of the model is set, more details will be given. [5] As mentioned earlier, the model considers only two types. One will follow a stabilizing, or mean reverting style of forecast, and the other is a form of destabilizing trend following type of strategy. We will refer to these as type 1 (reverting), and type 2 (trending) strategies. The fraction of type 1 agents at time t is given by $n_{1,t}$, and type 2 is given by $n_{2,t} = (1 - n_{1,t})$.

Also, adding noise is necessary both for estimation and realistic simulation. This now gives a final equation for the dynamics of x_t,

$$x_t = \frac{1}{R^*}(n_{1,t} E_{1,t} x_{t+1} + (1 - n_{1,t}) E_{2,t} x_{t+1}) + \epsilon_t, \tag{14}$$

where ϵ_t is normally distributed $N(0, \sigma_\epsilon^2)$ supply noise. Closing the model requires only two additional

[5] Timing is critical in these models, so time subscripts are designed to help readers develop their own codes to replicate. They may not exactly align with the original authors' notation.

components. The expectations of the two types, and the agent adaptations which feed into $n_{1,t}$ The expectations are written as functions of x_t,

$$E_{1,t}x_{t+1} = \phi_1 x_{t-1} \tag{15}$$

$$E_{2,t}x_{t+1} = \phi_2 x_{t-1}, \tag{16}$$

giving a simple pricing relationship in equation 14 once $n_{1,t}$ is determined. Notice that the timing in the model requires that expectations for time period $t+1$ to be determined at $t-1$. This is common in many types of models, and helps to make price dynamics more tractable.

The key feature underlying these models is the adaptive behavior of the agent types. Agents in the population are assumed to endogenously chose between the two types of strategies based on past performance. In Hommes & in 't Veld (2014) they are assumed to use past realized profits to determine their future rule choice. The profitability for each type at the start of period t (before x_t is determined) is given by,

$$\pi_{h,t-1} = R_{t-1}z_{h,t-2} = R_{t-1}\frac{E_{h,t-2}R_{t-1}}{a\sigma^2}. \tag{17}$$

Boswijk et al. (2007) show that the single period realized profitability for a strategy can be written as

$$\pi_{h,t-1} = \frac{(1+g)^2}{a\eta^2}(E_{h,t-2}x_{t-1} - R^* x_{t-2})(x_{t-1} - R^* x_{t-2}) \tag{18}$$

$$\pi_{h,t-1} = \frac{(1+g)^2}{a\eta^2}(\phi_h x_{t-3} - R^* x_{t-2})(x_{t-1} - R^* x_{t-2}) \tag{19}$$

$$\pi_{h,t-1} = C(\phi_h x_{t-3} - R^* x_{t-2})(x_{t-1} - R^* x_{t-2}) \tag{20}$$

$$\tag{21}$$

with $\eta^2 = (1+\delta^*)^2(1+g)^2\sigma_\epsilon^2$. The reason for writing the last equation is to emphasize that eventually those parameters will be grouped into a single constant C, and combined with the intensity of choice parameter, β. This greatly reduces the number parameters, but unfortunately none of these can be identified in the estimation. Single period profits are converted into a smoothed longer term utility or fitness measure using,

$$U_{h,t-1} = (1-\omega)\pi_{h,t-1} + \omega U_{h,t-2} \tag{22}$$

which smoothes noise by generating an exponentially weighted average moving into the past for each strategy h. This strategy fitness measure determines the fraction of trader types using a multinomial logit as in Brock & Hommes (1997),

$$n_{1,t} = \frac{e^{\beta U_{1,t-1}}}{e^{\beta U_{1,t-1}} + e^{\beta U_{2,t-1}}}. \tag{23}$$

It should now be clear how C is not identified and is absorbed into the intensity of choice parameter, β. The computer simulations in this paper will always use the algebraically equivalent, but often numerically more stable,

$$n_{1,t} = \frac{1}{1 + e^{\beta(U_{2,t-1} - U_{1,t-1})}}.$$

R^* is calibrated to the data which is quarterly, and set to 1.008.[6] The remaining parameters, $(\phi_1, \phi_2, \beta, \omega, \sigma_\epsilon^2)$ need to be estimated. The system of equations can be matched to the data using nonlinear least squares. The authors do this and give several sets of reasonable parameters. In their estimates ϕ_1 is slightly less than 1 (stabilizing), and ϕ_2 is slightly greater than 1 with explosive expectations. The model fit is not sensitive to the value of β, and because of this, several of the estimated parameter vectors simply fix β to a reasonable guess. The dynamics of the entire system are similar to a threshold or exponential autoregressive model, but more complicated given the dynamics for $n_{t,1}$. This makes it impossible to derive analytics for model stationarity. This will be done here through simulations.

To get a quick feel for how well the model works visually the first two figures compare actual price dividend ratios with one of the model simulations. Figure 1 uses a long time series build by merging the CRSP series (through 1926) with older data from William Schwert back to the mid 1880's. The raw series is the standard annual price/dividend ratio. Recently, many stocks have been using share repurchases as a second vehicle for getting cash back in the hands of shareholders. The second (red) line displays a price/dividend ratio adjusted for share repurchases. Both display the typical patterns of long erratic cycles, but the latter appears more stationary by ameliorating the recent behavior in the series.[7] Figure 2 plots a similar length time series from the model. It is using parameter set (B) from table 1. It is visually similar to the actual data. These figures are simply motivation, since it is not the purpose of this paper to redo the estimation results performed in Hommes & in 't Veld (2014). They show that this model fits a similar price/dividend series well, and perform many diagnostics on their estimated models.

To decide on the appropriate model benchmark for the rest of the paper the three sets of estimated parameters will be simulated for a very long time series of 10,000,000 time steps, corresponding to quarters

[6] See Hommes & in 't Veld (2014) for details. In their data they estimate annual $g = 1.3$, and annual $r = 4.69$. This is then converted to quarterly values. The paper also contains full information on the data used, and model estimation.

[7] See LeBaron (2013) for details on these series.

of U.S. stock market data. The parameters are taken from table 1 in Hommes & in 't Veld (2014), and displayed here in table 1. Parameter sets (A), (B), and (C) correspond to their parameter sets. The first columns display the parameters for the various models, and the last column shows the fraction of 250 runs which explode. Two of the 3 parameter sets, (A) and (C), are explosive, but (B) is stable, with no explosive trajectories. Most of the experiments in the paper will use (B). A good conjecture is that instability is tied to the magnitude of ϕ_2. This is tested in the last two rows of table 1 where the parameters from (B) are used, but increasing ϕ_2 first to 1.02, and then to 1.03. For the first of these experiments 23 percent of the runs explode, and for the larger value it is 100 percent. It would appear that the range of ϕ_2 moving from 1.017 to 1.026 is where the model becomes explosive. This issue will be dealt with later in the paper. For the moment this justifies the use of parameter set (B) with very long simulation runs.

The long runs are necessary to get very precise estimates of several moments in the data.[8] Table 3 shows the properties of several key estimated moments from the simulation. They are listed across the columns, and the table reports the mean and standard deviation of these estimated moments across 250 simulations with sample sizes of $10,000,000$ each. In all cases the standard deviations are very small relative to the mean except for the case of $E(x_t)$ where the mean is zero. This shows that at these long sample sizes the model is able to give estimates which are very close to the true values. It is also important to see that the table tests internal information from the model such as $E(n_{1,t})$, and $E(U_{1,t})$ since these will be used extensively in later sections for estimating the shape of the objective function at the fixed agent parameters.

3 Fitness gradients

The main objective of this paper is to explore the micro consistency of the agent-based model which is represented by the estimated macro parameters. They are estimated to be a best fit to the overall data, but they do not test whether they are consistent with the underlying agent-based decision making which is part of all agent-based models. Agents are forced to stay withing the fixed forecasting rules because these parameters best fit the data. The key question is "If agents were given the chance to change their forecasts a little would they take this opportunity?"

To do this several measures will be estimated. First, the value of changing the strategy by a small amount on overall expected utility is estimated. This is not a completely straightforward estimation, so several measures will be used. The first estimates assume that an individual agent is moving alone, and

[8]For this it is obvious that ergodicity is being assumed for the models and parameters.

not taking into account the fact that others might move as well. This would seem to be the appropriate way to run this test for an individual. First, the strategy 2, which is given in ϕ_2 will be modified slightly to ϕ_2'. Profits for the strategy in each period would go to

$$\pi_{2,t}' = (\phi_2' x_{t-2} - R^* x_{t-1})(x_t - R^* x_{t-1}) \tag{24}$$

which will cause a corresponding change in $U_{2,t}$ to $U_{2,t}'$. To estimate overall utility of this strategy the agent uses $n_{2,t} = (1 - n_{1,t})$ as a probability that they are using strategy 2 in a given period,

$$\bar{U} = \frac{1}{T} \sum_{t=1}^{T} (n_{1,t} U_{1,t} + (1 - n_{1,t}) U_{2,t}) \tag{25}$$

and comparing with this the new utility,

$$\bar{U}' = \frac{1}{T} \sum_{t=1}^{T} (n_{1,t} U_{1,t} + (1 - n_{1,t}) U_{2,t}'). \tag{26}$$

However, a slightly more rational agent could now assess the switching probabilities using the new utility, $U_{2,t}'$, and use this to estimate new fractions, $n_{1,t}'$, and then estimate,

$$\hat{U}' = \frac{1}{T} \sum_{t=1}^{T} (n_{1,t}' U_{1,t} + (1 - n_{1,t}') U_{2,t}'). \tag{27}$$

Finally, it is useful to estimate utilities based on the entire population changing. This is not an individual agent experiment, but it is a useful calculation to do, to explore the entire social impact on the system of the change in strategy. For this case the model is completely rerun at the new value for $\phi_2 = \phi_2'$ with all utilities, fractions (n_t) and corresponding pricing (x_t) allowed to change.

All three of these utility changes are shown in figure 3. The lines labeled "n fixed" and "n adjust" correspond to the first two experiments where agents only consider changing their own strategies. Both schedules are similar and show that the gradient is not zero at the estimated value of ϕ_2. This is important since it means that the given value of ϕ_2 is not an individual maximum, and under local utility adaptation there would be a desire to crawl uphill by moving the parameter ϕ_2 out to larger values. It is interesting that the values for adjusted n are actually slightly smaller. This suggests that the dynamic rule change is actually doing better for the agent using the wrong utilities. There is nothing inconsistent with this and the logit fractions, since we don't know that the logit really is an optimal forecast of the future gains from different strategies. It is only a backward looking response to past data.

Moving to the case where all agents change their strategies, and price impact is allowed, a completely different situation appears. The gradient is again not flat, but this time it is downward sloping. Agents, as a whole, would be worse off if they all simultaneously increase their forecast parameter ϕ_2. It is probably the case that increased trend following behavior in the population has endogenously reduced the trendiness of the data, but some further tests on this are necessary. This divergence between individual and group outcomes is a classic case in economics where a lack of individual coordination can lead to reductions in utility for the population as a whole. See Schelling (1978) for many classic examples of this.

Another interesting part of this result, reflects on the stability properties of the model. The previous section showed that the probability of model instability is increasing as ϕ_2 increases. The gradients here are all pointing in the direction of instability. If individual agents were increasing ϕ_2 in response to positive gradients, then they would be taking this model closer to an explosive situation.

In figure 4 the experiment is repeated for ϕ_1. This situation is quite different. In this case, all the objective functions are moving in the same direction. They all suggest that agents would be interested in reducing ϕ_1 individually. However, in this case, this would also be a social improvement if everone followed this.

These figures suggest a much richer dynamic than is given in the original two agent model. To begin to explore this a model covering a uniform set of ϕ's inside the original model support is created and simulated. In this case the strategy space is opened up to 15 values of ϕ from 0.947 to 1.017. This involves generalizing the pricing and adaptation equations as follows,

$$x_t = \frac{1}{R^*}(\sum_{j=h=1}^{H} n_{h,t}E_{h,t}x_{t+1}) + \epsilon_t, \tag{28}$$

$$n_{h,t} = \frac{e^{\beta U_{h,t-1}}}{\sum_{h=1}^{H} e^{\beta U_{h,t-1}}}. \tag{29}$$

Figure 5 displays the time average of $n_{h,t}$ for all the H forecasts, indexed by h. The value is far from uniform, and shows large amounts of forecasting mass on both the lowest and highest levels of ϕ. This is consistent with the earlier gradient graphs. If agents are indeed interested in shifting, then the levels of $n_{h,t}$ should reflect this. They will not go completely to the extremes because of the random nature of the discrete choice machinery. It always imposes some noise on the system. However, the relative values of n should reflect the learning gradients and they do. Also, the gradients were steeper for lower values of ϕ (left side of the plot), and this appears as a stronger probability mass for n on that side then for the larger ϕ

on the right side. This plot also suggests that the agents are constrained in what they do since the fractions are pushed against both end points for ϕ.[9]

4 Large parameter changes

In the previous section only local changes to the parameters were considered. Pointing out local gradient slopes is useful for understanding the underlying internals of the model, but may not have a big impact on the actual dynamics of the model observable price/dividend ratios. This section explores bigger changes in the parameters. Specifically, it will examine increases in the trend following parameter ϕ_2 to much larger values. This introduces the problem that was brought up in the earlier section, models with larger values of ϕ_2 tend to be unstable. This problem has been addressed in earlier two type models such as Gaunersdorfer & Hommes (2007) by adding a stabilization equation that eventually adjusts the populations when the price becomes too far from the fundamental. Trend followers lose faith in their models and push toward mean reversion. This is accomplished here by adjusting the fraction of trend followers using,

$$\tilde{n}_{2,t} = n_{2,t} e^{-\frac{x_{t-1}^2}{\alpha}}, \tag{30}$$

where $n_{2,t}$ corresponds to the original logit fraction of trend followers. The parameter α is set to an arbitrarily large value of 10000 which doesn't the impact dynamics of the model while x_t remains relatively small, but still shuts down explosive trajectories.

Figure 6 repeats the utility comparisons of the last section, but now allows for a larger range of increases in ϕ_2 which are enabled by the model stabilizer. It is clear that the utility gains from increasing ϕ_2 are not local, but continue out for a large range of ϕ_2. There is essentially no change from the earlier estimates. It is also intereresting that now the magnitude of the increases are getting larger with an increase in nearly 10 percent for $\phi_2 = 1.05$.

This figure again suggests that the presence of a larger ϕ_2 rule may be chosen, and could possibly impact the market dynamics. To test this a 3 agent model is considered, using the original parameters, but adding a third forecasting rule with ϕ_3 set to some larger values. The model stabilization in equation 30 is done only for $n_{3,t}$, not for the other forecasts. This is sufficient to keep the trajectories stable.

Figure 7 displays the time average of $n_{h,t}$ for the three strategies for two larger values of ϕ_3, 1.05 and

[9] It should be noted that $n_{h,t}$ is not a static object. It moves through time in interesting ways with the data. This unconditional mean snapshot of $n_{h,t}$ can only be taken as a crude summary of what the overall model looks like.

1.15. These are all results of a complete simulation of the three type model with pricing taking into account all the agents in the market. The figure shows that in terms of populations, the additional strategies matter. They are able to attract a large fraction of n with amounts slightly over 20 percent in both cases. In the second case, ϕ_2 is driven to a very low value, near 0.05. This indicates that in the new market, not only does the new strategy survive, but it may thrive.

These new strategies survive in the model, but are they going to actually impact the market dynamics? This is tested by looking at the time series for x_t, and some of its key moments to get an initial picture for how the pricing dynamics is impacted. Figure 8 displays a 100 year snapshot (400 quarters) for x_t for the three type models with $\phi = 1.05$, and $\phi = 1.15$ respectively. They show a distinct difference in that increasing the trend following parameter increases market volatility, and makes the big cycles more pronounced and regular than they are in the first simulation. On the other hand, the graph for 1.05 suggests a time series that visually, is not too far from some of the earlier simulations at the estimated parameters.

Table 3 describes some basic moments for the data across different parameter sets and models for a $10,000,000$ quarter run. The four models are labeled (B) corresponding to the benchmark parameter set (B) from table 1. The next row, labeled, $H = 15$ corresponds to the model with 15 different types, but still in the range of parameter set (B). The final two rows use the last two three type models with $\phi = 1.05$ and $\phi = 1.15$ respectively. Moments are reported for the simulated values of x_t which correspond to price/dividend ratios. Also, given the near random walk behavior of x_t, it is also useful to examine some moments generated from the first difference $z_t = x_t - x_{t-1}$. The table reports the first order autocorrelation for z_t and z_t^2. The final column is a measure of mean reversion which is closely related to Dickey/Fuller tests. It correlates the level of the system, x_t with the next period change z_{t+1}. In a true random walk process, this value would be zero. In a stationary process it will be negative.

The first two rows of the table demonstrate that adding the additional agents in between the original ρ_h values has little or no impact on the estimated properties of the system. This is an interesting defense of the two type model framework. It supports the power of the simplification in that it gives analytic tractability without giving up much in terms of the dynamic process. This result changes as the values of ϕ_3 are added outside of the model's support. The last two rows of the table show that this changes the time series in several important ways. First, it increases the volatility as shown by the increase in standard deviation. It also moves it closer to appearing random walk like. The first order autocorrelation moves toward one, and the mean reverting measure (last column) moves toward zero. This would appear to be driven by the strong trending behavior of the ϕ_3 traders. They push toward instability reducing the underlying long range stability in the process. However, they also have one unusual change which is not consistent with this. The table reports the autocorrelations of the first difference in ρ_z. For the first two models this is near zero as it would be for a true random walk. However, in the last two cases this value becomes positive. For $\phi_3 = 1.15$ it is is over 0.3. These results are both different from the underlying model (B), and from a true random walk, and they may indicate that the nonlinearities may be getting stronger as this trend parameter increases. This is consistent with the qualitative features in the previous figure.

The last two rows present the corresponding moments from the U.S. data presented in figure 1. These are both the raw price/dividend ratio, and the ratio adjusted for share repurchases. For most moments the actual data shows a good alignment with parameter set (B) which makes sense since that model was estimated with close related data. The adjusted data are a little different with lower volatility, lower autocorrelation, and lower kurtosis. All of these are probably related to how it attenuates the run up in prices during the dot com bubble. One curious feature in the data is that there is strong evidence for persistence in volatility of the first differences. This feature is not shared by any of the models. Finally, the adjusted data shows a much stronger case for mean reversion which is probably also caused by the adjustment of the data around the turn of the century. Except for the persistence of volatility the unadjusted data shows good agreement with model (B), but increasing ϕ_3 moves the model away from the data.

5 Conclusions

Estimable few type models present a new and interesting direction for agent-based models and their empirical validation. This paper demonstrates several different features of one of these models that has been empirically estimated with U.S. equity market data. In looking into the details of the model several new and interesting features are uncovered. Some of these are supportive of this framework, while others can be viewed as critical or at least cautious.

First, the two type model appears robust to the addition of additional forecasting rules that are of the same form, and exist inside the support of the original two type model. In other words in a model successfully fit with forecast parameters in the range of $[a, b]$ adding more agents inside this range does not impact the results. This may be partly a general feature related to the "large type limit" of Brock, Hommes & Wagener (2005), and it also may be related to the result that agent mass appears to concentrate at the extremes, yielding a nearly two type model. This should all be viewed as good news for the empirical framework.

The second result is less supportive. This paper's key experiment was to examine the micro consistency of the forecast rules in the two type model. It was found that they are not consistent with local optimization, or simple adaptive behavior, as long as this behavior was allowed to make small changes to the forecasting parameters. Agents operating individually would seek to move up their objective gradients by changing the key forecast parameters from the estimated values. This is important since the macro coordination imposed by the model and its estimation is not consistent with the underlying agent-based model and local behavior that it implies. Also, the model gives a very clear picture of a case where micro and macro objectives give opposing recommendations for behavior at the individual level.

Finally, this local adaptation and, in particular, the desire to increase the trending forecast parameter to larger values may be a problem for underlying model dynamics. First, it looks likely that it would move the model out of the stable region to a parameter set that is unstable. Furthermore, once model stabilizing components are added, the addition of a third stronger trending trader has a major impact on the underlying time series of the model.

These problems may not be insurmountable for these models. They might be assumed away by imposing bounds on the forecasting behavior ex-ante. For example, simply saying that agents would keep their trend following forecast parameters below some plausible upper bound would be enough to save the original model. However, it is important to then note that parameter estimation involves both estimating the upper bound along with agent behavior. Also, determining good reasons for such an upper bound are necessary.

References

Boswijk, H. P., Hommes, C. H. & Manzan, S. (2007), 'Behavioral heterogeneity in stock prices', *Journal of Economic Dynamics and Control* **31**(6), 1938–1970.

Branch, W. A. & Evans, G. W. (2011), 'Monetary policy and heterogeneous expectations', *Economic Theory* **47**, 365–393.

Brock, W. A. & Hommes, C. H. (1997), 'A rational route to randomness', *Econometrica* **65**, 1059–1097.

Brock, W. A. & Hommes, C. H. (1998), 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model', *Journal of Economic Dynamics and Control* **22**(8-9), 1235–1274.

Brock, W. A., Hommes, C. H. & Wagener, F. O. O. (2005), 'Evolutionary dynamics with many trader types', *Journal of Mathematical Economics* **41**, 7–42.

Chiarella, C., Dieci, R. & He, X.-Z. (2009), Heterogeneity, market mechanisms, and asset price dynamics, *in* T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', Elsevier, USA, pp. 277–344.

Chiarella, C., He, X.-Z. & Zwinkels, R. C. J. (2014), 'Heterogeneous expectations in asset pricing:Empirical evidence from the s&p 500', *Journal of Economic Behavior and Organization* **105**, 1–16.

Gaunersdorfer, A. & Hommes, C. (2007), A nonlinear structural model for volatility clustering, *in* A. Kirman & G. Teyssiere, eds, 'Micro Economic Models for Long Memory in Economics', Springer-Verlag, pp. 265–288.

Hommes, C. H. (2006), Heterogeneous agent models in economics and finance, *in* L. Tesfatsion & K. L. Judd, eds, 'Handbook of Computational Economics', Elsevier, pp. 1109–1186.

Hommes, C. H. & in 't Veld, D. (2014), Booms, busts and behavioral heterogeneity in stock prices, Technical report, CeNDEF, University of Amsterdam.

Hommes, C. H. & Wagener, F. (2009), Complex evolutionary systems in behavioral finance, *in* T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', North-Holland, pp. 217–276.

Kirman, A. P. (1991), Epidemics of opinion and speculative bubbles in financial markets, *in* M. Taylor, ed., 'Money and Financial Markets', Macmillan, London, pp. 354–368.

LeBaron, B. (2006), Agent-based computational finance, *in* L. Tesfatsion & K. L. Judd, eds, 'Handbook of Computational Economics', Elsevier, pp. 1187–1233.

LeBaron, B. (2013), Heterogeneous agents and long horizon features of asset prices, Technical report, Brandeis International Business School.

Lof, M. (2012), 'Heterogeneity in sock prices: A star model with multivariate transition function', *Journal of Economic Dynamcis and Control* **36**(12), 1845–1854.

Lux, T. (1997), 'Time variation of second moments from a noise trader/infection model', *Journal of Economic Dynamics and Control* **22**, 1–38.

Lux, T. (2009), Stochastic behavioral asset pricing models and the stylized facts, *in* T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', North-Holland, pp. 161–215.

Schelling, T. (1978), *Micromotives and Macrobehavior*, Norton, New York, NY.

ter Ellen, S. & Zwinkels, R. C. (2010), 'Oil price dynamics: A behavioral finance approach with heterogeneous agents', *Energy Economics* **32**(6), 1427–1434.

Table 1: Stability

Model	ϕ_1	ϕ_2	β	ω	σ_ϵ^2	Fraction unstable
(A)	0.936	1.026	1.000	0.824	14.12	1.00
(B)	0.947	1.017	2.443	0.800	14.09	0.00
(C)	0.940	1.026	10.000	0.852	13.87	1.00
B+	0.947	1.020	2.443	0.800	14.09	0.23
B++	0.947	1.030	2.443	0.800	14.09	1.00

Model stability estimates. Fraction stable reports the fraction of explosive runs for the given sets of parameters.

Table 2: Sample moment accuracy

Moment	$E(x_t)$	σ_x	ρ_x	$E(n_t)$	$E(U_{1,t})$
Mean	-0.005	18.39	0.978	0.677	0.218
Std.	0.046	0.035	0.001	0.001	0.004

Mean and standard deviations across 250 simulations of sample sizes of $10,000,000$.

Table 3: Feature comparisons

Model	σ_x	ρ_x	Kurtosis(x)	ρ_z	$\rho_{z_t^2}$	$corr(x_t, z_{t+1})$
Parameter set (B)	16.3759	0.9730	3.6586	0.0205	0.0020	-0.1163
$H = 15$	16.6222	0.9740	3.6379	0.0129	0.0012	-0.1140
$\phi_3 = 0.05$	23.8550	0.9869	3.8264	0.0778	0.0055	-0.0811
$\phi_3 = 0.15$	54.4639	0.9963	2.5516	0.3442	0.0774	-0.0428
US adjusted	6.6689	0.9495	2.9218	0.0309	0.2157	-0.1580
US actual	14.1804	0.9778	6.1768	-0.0306	0.2725	-0.0889

Moments estimated over single 10,000,000 period run. x_t is the price/dividend ratio, and $z_t = x_t - x_{t-1}$. ρ is the first order autocorrelation.

Figure 1: U.S. price/dividend ratios

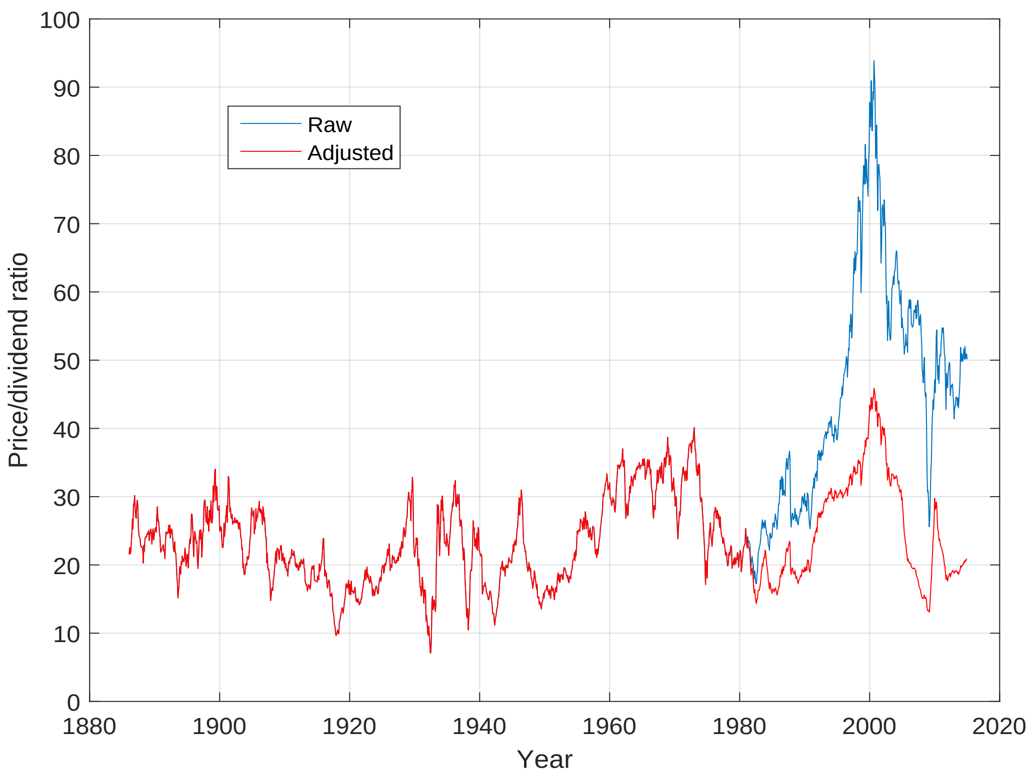

Figure 2: Simuated annualized price/dividend ratios

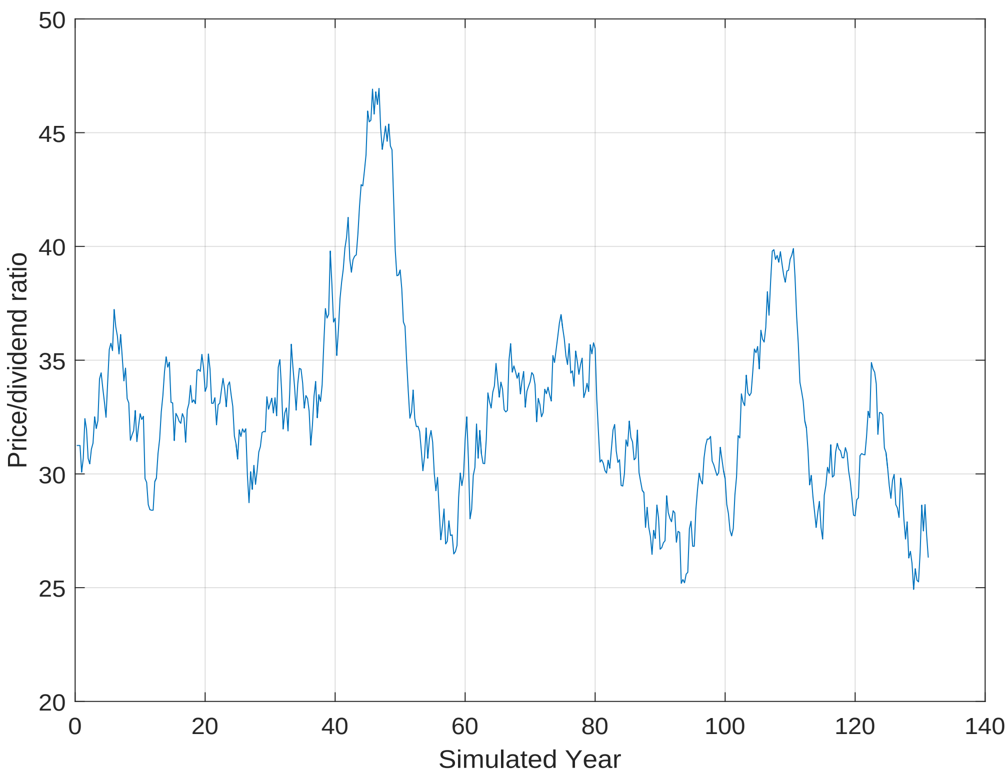

Figure 3: Fitness changes from local forecast changes ϕ_2

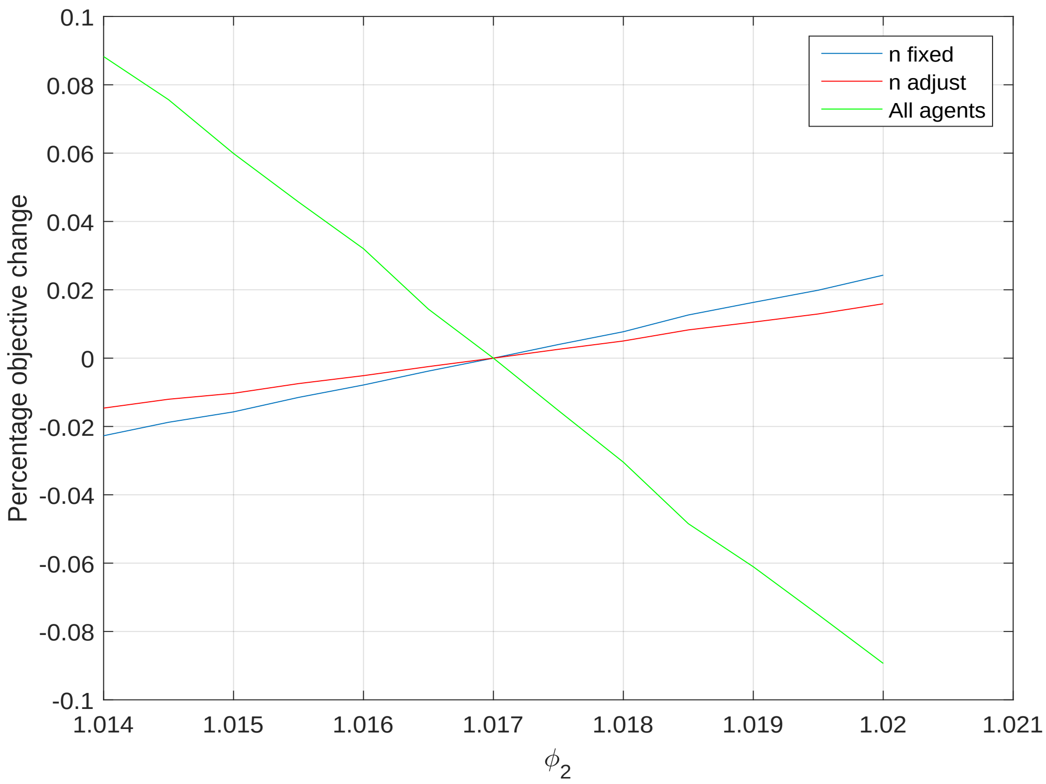

Figure 4: Fitness changes from local forecast changes ϕ_1

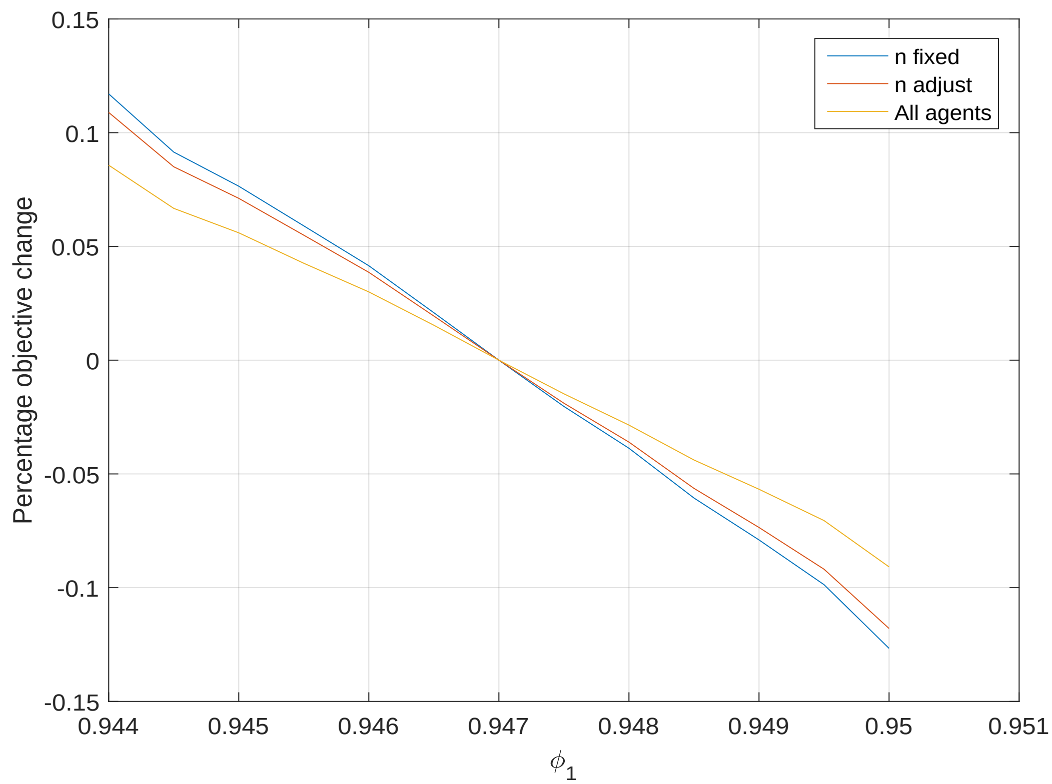

Figure 5: Agent type distribution

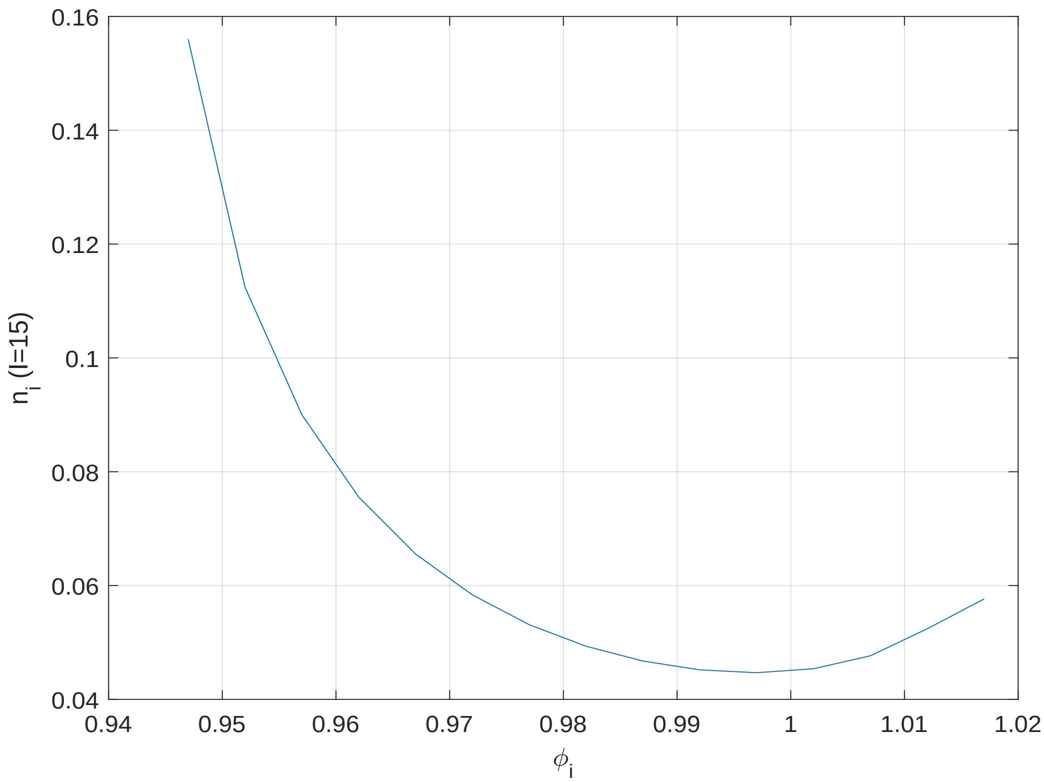

Figure 6: Fitness for large forecast changes

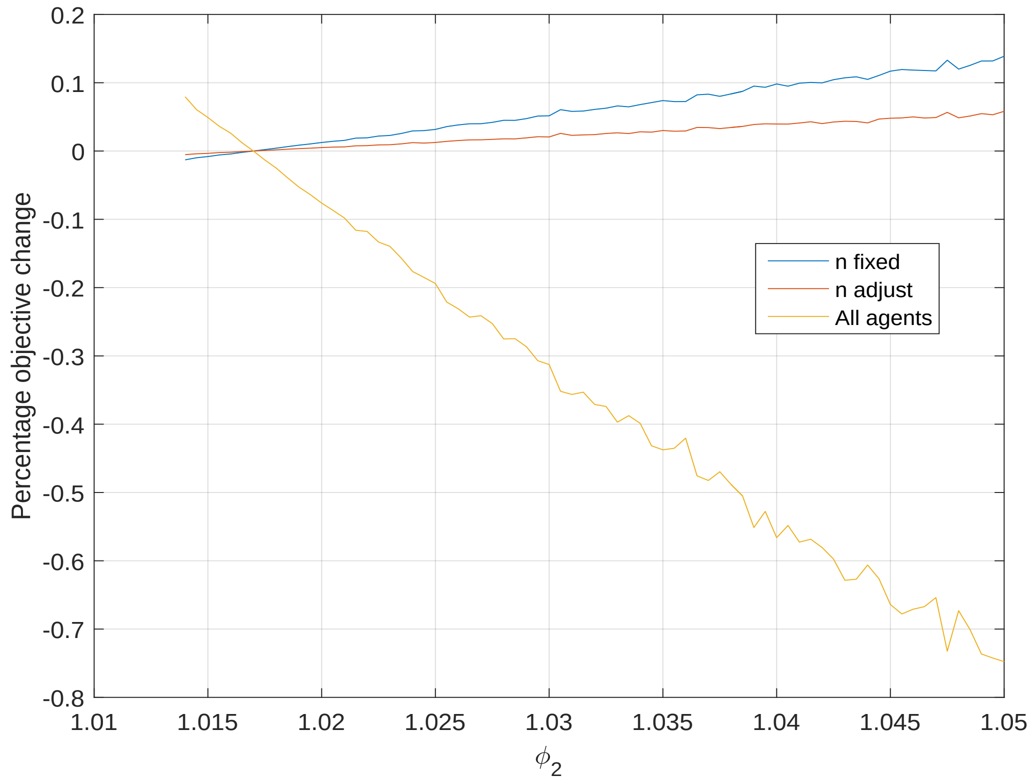

Figure 7: Type fractions

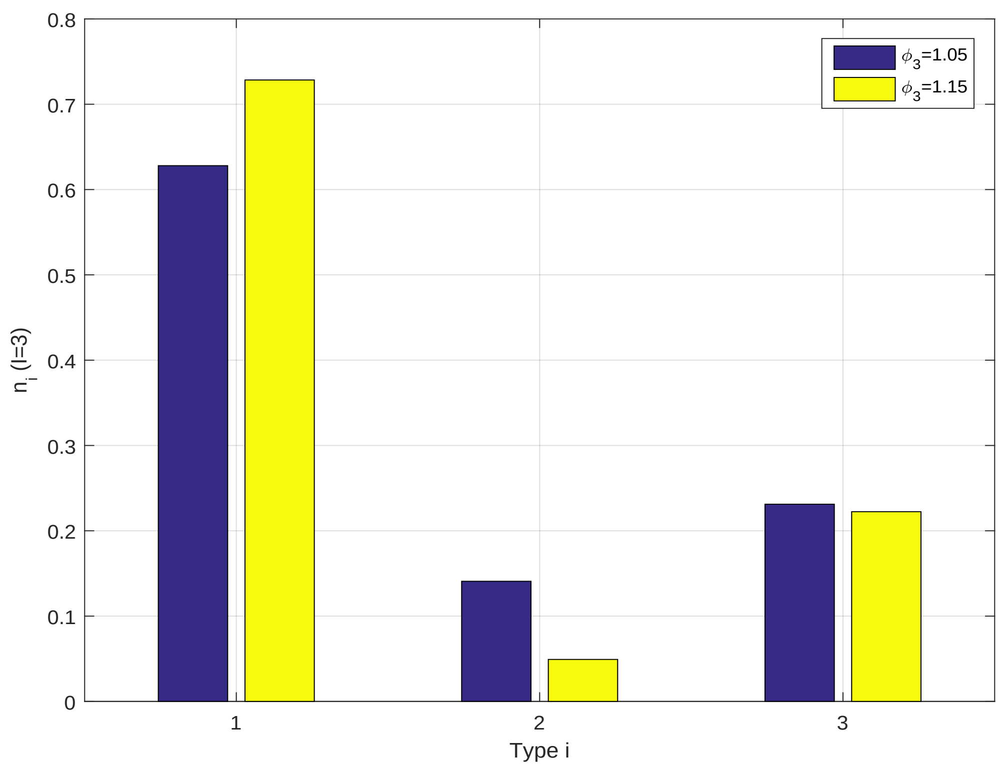

Figure 8: Comparison time series