Learning to Herd

Paul Shea*

Bates College

March 25, 2020

Abstract

We introduce a new approach to modeling herd behavior. Households choose whether to pay a cost in order to form a more reliable forecast. Under expectations herding, households obtain disutility from deviating from the population's average forecast, while under expectations anti-herding, they obtain utility from doing so. The model exhibits distinct states including a state with high consumption volatility where all agents choose to form the more reliable forecast, and a state with low consumption volatility where all agents choose the less reliable forecast. As the desire to herd initially becomes stronger, the model spends more time in the high volatility state If the desire to expectations herd is above a threshold level, however, then the model abruptly starts to spend all its time in the low volatility state.

JEL Classification: E17, E21, E27.

Keywords: Adaptive Learning, Herd Behavior, Multiple Equilibria, Forecasting.

*pshea@bates.edu

1 Introduction

An extensive literature in both economics and finance examines herd behavior where heterogeneous agents distort their actions to match or deviate from ("anti-herding") those of the population. Behavioral economics has cited herd behavior as a critical driver of speculative bubbles, such as the housing bubble that preceded the Great Recession.[1] Herd behavior has been introduced into theoretical settings mostly through information cascades.[2] Here, agents infer information from the actions of others in a sequential game so that the population may converge to a common belief despite individual agents' their own private information. Another approach is to assume that agents have a reputational incentive to ignore their private information and mimic the actions of agents who have acted before them.[3]

This paper presents and examines a new method of modeling herd behavior which we call *expectations herding.* Households choose between forming a simple expectation with a large mean squared forecast error, or enduring a disutility in order to form a better forecast with a smaller mean squared forecast error. Households differ in the size of this disutility, allowing for heterogeneous expectations. Expectations herding assumes that households also obtain disutiliy from having their expectation differ from the average expectation across all households. Likewise, expectations anti-herding assumes that agents obtain utility from standing out from the crowd.

We introduce this expectations herding into a simple, two-period overlapping generations model. We find that if agents know the correct distribution of forecast errors, then expectations herding often induces multiple equilibria. One such equilibrium may be an interior solution where both types of expectations coexist and all agents obtain some disutility from departing from the average expectation. One or two corner solutions may also exist where all agents coordinate so that they all form one type of expectation. We thus assume that agents use adaptive learning to learn about forecast errors. In addition to being more plausible than assuming that agents know the distribution of forecast errors, learning acts as a selection criteria among multiple equilibria.

Under learning, when anti-herding is strong enough, the model is always at an interior equilibrium where some agents form each type of expectation. Once ant-herding is made sufficiently weak, however, the model starts to spend some time at a corner solution where all agents form the more reliable forecast and the model's behavior is identical to rational expectations. This state exhibits maximum consumption volatility and consumption is perfectly correlated with output. We show that as anti-herding continues to weaken, and then herding starts to strengthen, the model

[1] See Akerlof and Shiller (2009), and Shiller (2000).
[2] See Bikhchandani *et. al.* (1992) and Fajgelbaum *et. al.* (2017) for two prominent examples.
[3] See Scharfstein and Stein (2000).

spends more time in this state. Surprisingly, we never observe both corner solutions in the same simulation and for low enough values of herding, we never observe the state where all agents form the less reliable forecast.

As the desire to herd becomes stronger, it passes a threshold level where the model suddenly displays very different behavior. Above this value, we almost never observe the high volatility state. Instead, the model converges to the state where all agents form less reliable forecasts and the model rarely leave this state. Here, the model's consumption volatility is minimized.

The paper is organized as follows. After a brief review of the empirical evidence on expectations herding, Section 2 presents the overlapping generations model and expectations formation process. Section 3 reports simulation results. Section 4 concludes.

1.1 Empirical Evidence on Herding and Anti-Herding

Testing for herding behavior among forecasters is common in the finance literature. Trueman (1994), and Werner *et al.* (1999) find that analysts tend to bias their earnings' forecasts towards those previously released by other analysts. Galariotis *et. al.* (2015), and Clements *et. al.* (2017) find evidence of herd behavior in stock price forecasting in both the United States and United Kingdom. Pierdzioch and Rülke (2012), however, find the opposite result, where forecasters anti-herd in an effort to stand out from the crowd.

Several papers suggest that anti-herding may be more common among macroeconomic forecasters. Lamont (2002) finds evidence that as macroeconomic forecasters become older and more established, they anti-herd by issuing more radical forecasts. Rülke *et. al.* (2016) also find widespread evidence of anti-herding among business cycle forecasters. Pierdzioch *et. al.* (2010) also find widespread anti-herding among professional oil price forecasters. An exception is Bewley and Fiebig (2002), who document herding behavior among interest rate forecasters.

2 Model

We use a simple two-period overlapping generations model with exogenous output to illustrate the expectations herding mechanism. Each period there exists both a continuum of young agents on the unit interval, and a continuum of older agents. Agents receive an exogenous endowment of income in each period, Y_t. They choose how much to consume in period one. They may costlessly save their output (we do not include a formal credit market) and freely borrow against future income. Young household i faces the following problem:

$$Max_{C^i_{1,t},f^*} \quad ln(C^i_{1t}) + \beta E_t[ln(C^i_{2,t+1)}] - jf^*i - \tau(E^i_t[Y_{t+1}] - E_t[Y_{t+1}])^2 \tag{2.1}$$

s.t:

$$C^i_{1,t} + C^i_{2,t+1} = Y_t + Y_{t+1} \tag{2.2}$$

$$Y_t = Y^{\rho}_{t-1} e_t \tag{2.3}$$

where $u_t = ln(e_t)$ is a white-noise, mean-zero, productivity shock. Young households' choice of consumption yields an unremarkable consumption Euler Equation:

$$\frac{1}{C_{1,t}} = E_t \left[\frac{\beta}{Y_t + Y_{t+1} - C_{1,t}} \right] \tag{2.4}$$

Log-linearizing (2.4) and (2.3) yields:

$$c^i_{1,t} = E^i_t \left[\frac{y_t + y_{t+1}}{1+\beta} \right] \tag{2.5}$$

$$y_t = \rho y_{y-1} + u_t \tag{2.6}$$

where lower-case indicates percentage-deviations from the steady state. If agents form rational expectations, then $E^i_t[y_{t+1}] = \rho y_t$ for all i and:

$$c_{1,t} = \frac{1+\rho}{1+\beta} y_t \tag{2.7}$$

Young households must also choose which forecast to form, $f^* = 0, 1$. Choosing $f^* = 1$ allows households to form the correctly specified AR(1) forecast for y_{t+1}. Doing so, however, requires that the household suffer a heterogeneous disutility equal to ji where $j > 0$ is an exogenous parameter. This disutility reflects the effort needed to determine and calculate the superior expectation. If the young household instead chooses $f^* = 0.$, then they instead form a naive expectation where $y^e_{t+1} = \bar{y}$=0, output's steady state value.

The parameter τ is novel to this paper and allows for expectational herding. The term $E_t[y_{t+1}] = \int_0^1 E^i_t[y_{t+1}]di$ is the population's average forecast. If $\tau > 0$, households obtain disutility from deviating from the average forecast and they have an incentive to herd. This captures situations where the reputational cost of a forecaster standing out from the crowd and being wrong, is greater than the benefit of standing out and being correct. If $\tau < 0$, however, agents gain utility from standing out from the crowd.[4]

We begin by assuming that when choosing f^*, agents know the distribution of both types of forecast errors, as well as the distribution of how they will differ from the average expectation. It is natural to question how agents could know these distributions without actually calculating the AR(1) forecast? We soon address this issue by assuming that agents must learn about the reliability and conformity of each type of forecast using only past data.

The forecast error for the AR(1) forecast of y_{t+1} is simply u_t, yielding a mean-squared forecast error of σ_u^2. The forecast error for the naive forecast is thus $\rho y_{y-1} + u_t$, yielding a mean-squared forecast error of $\rho^2 y_t^2 + \sigma_u^2$. The expected utility gain from being better able to smooth consumption by choosing the AR(1) forecast is well-approximated by:

$$\kappa(y_t) = ln(\bar{C_{1,t}}) - \frac{1}{2} ln\left(\bar{C_{1,t}} + \frac{\rho^2 y_t^2}{1+\beta}\right) - \frac{1}{2} ln\left(\bar{C_{1,t}} - \frac{\rho^2 y_t^2}{1+\beta}\right) \tag{2.8}$$

A young household then chooses to form the AR(1) expectation if and only if:

$$ij + \tau((i-1)\rho y_t)^2 \leq \kappa(y_t) + \tau(i\rho y_t)^2 \tag{2.9}$$

When simulating the model in Section 3 (without learning), we solve for the equilibrium share of households using the AR(1) forecast using:

$$\hat{i} = \left\{ \begin{array}{ccc} \frac{\kappa(y_t) - \tau\rho^2 y_t^2}{j - 2\tau\rho^2 y_t^2} & if & \frac{\kappa(y_t) - \tau\rho^2 y_t^2}{j - 2\tau\rho^2 y_t^2} \in [0,1] \\ 1 & if & \frac{\kappa(y_t) - \tau\rho^2 y_t^2}{j - 2\tau\rho^2 y_t^2} > 1 \\ 0 & if & \frac{\kappa(y_t) - \tau\rho^2 y_t^2}{j - 2\tau\rho^2 y_t^2} < 0 \end{array} \right\} \tag{2.10}$$

An interior equilibrium does not always exist. Likewise, if $\tau > 0$, then the model is prone to multiple equilibria with corner solutions where $\imath = 0$ or $i = 1$. Intuitively, if households care about having their expectations conform to the average expectation, then it might be possible for all agents to either choose the naive expectation or the AR(1) expectation. In fact, as $\tau \to \infty$, three equilibria always coexist where $i = 0, \frac{1}{2}, 1$. The equilibrium where $i = \frac{1}{2}$ here represents a coordination failure where agents would be better off coordinating on either corner solution. When we simulate learning in Section 3, we show that the interior equilibrium is generally unstable and the model instead usually settles near one of the corner solutions.

The $i = 0$ corner equilibrium exists if:

[4]One example is Zillow's "Crystal Ball" Award given to the forecaster who is the most accurate at predicting housing prices. If a forecaster wants to win this award, they may have an incentive to distinguish their forecast from the competition, even if doing so increases their mean squared forecast error.

$$\tau(\rho y_t)^2 > \kappa(y_t) \tag{2.11}$$

The $i = 1$ corner equilibrium exists if:

$$j - \tau(\rho_t^y)^2 < \kappa(y_t) \tag{2.12}$$

2.1 Learning

We now introduce learning into the model. Learning serves two purposes. First, it acts as a selection criteria when multiple equilibria exist. Second, it is a plausible behavioral assumption for how households might choose among the two types of forecasts without actually knowing the AR(1) forecast.

We assume that all households are able to observe the history of how each forecast has deviated from the true value of y_t, and the mean forecast, $E_{[}y_{t+1}]$, even if they might not understand how the AR(1) forecast is formed. They then update their estimate of the difference in expected forecast errors (a_t) using:

$$a_t = (1-\gamma)a_{t-1} + \gamma\left[y_{t-1}^2 - (\rho y_{t-2} - y_{t-1})^2\right] \tag{2.13}$$

where ρy_{t-2} is the AR(1) forecast from period $t-2$ for output in period $t-1$. The naive forecast always equals zero by construction. The parameter γ is the gain and represents how fast agents learn. By making the gain a constant, the learning dynamics are persistent. The most common rational for constant-gain learning is that agents worry about structural changes to the economy as thus want to weigh recent observations more heavily.[5] By chance, the naive forecast may temporarily outperform the AR(1) forecast and, less often, agents might believe that the naive forecast is better. This will push agents towards the naive expectation, but does not ensure that all households choose it if $\tau \neq 0$.

Households also learn how each forecast deviates from the population average forecast. The learning parameter b_t is the estimated squared deviation for the AR(1) forecast while c_t is the estimated squared deviation of the naive forecast. These are updated using:

$$b_t = (1-\gamma)b_{t-1} + \gamma((\hat{i}-1)\rho y_t)^2 \tag{2.14}$$

[5] Sargent (1999), and Evans and Honkapohja (2001) provide overviews.

$$c_t = (1-\gamma)c_{t-1} + \gamma(\hat{i}\rho y_t)^2 \tag{2.15}$$

Young households then choose the AR(1) forecast if and only if:

$$ij + \tau b_t \leq a_t + \tau c_t \tag{2.16}$$

The equilibrium share of households choosing the AR(1) forecast is then:

$$\hat{i} = \left\{ \begin{array}{ccc} \frac{a_t+\tau(c_t-b_t)}{j} & if & \frac{a_t+\tau(c_t-b_t)}{j} \in [0,1] \\ 1 & if & \frac{a_t+\tau(c_t-b_t)}{j} > 1 \\ 0 & if & \frac{a_t+\tau(c_t-b_t)}{j} < 0 \end{array} \right\} \tag{2.17}$$

Aggregate young consumption then equals:

$$c_{1,t} = \frac{1+\hat{i}\rho}{1+\beta} y_t \tag{2.18}$$

Importantly, because households are simply using data to estimate how each forecast will deviate from the population average forecast, equilibrium under learning is unique.

3 Simulation Results

We simultaneously simulate five versions of the model. The first is rational expectations where all households have complete information, equivalent to imposing $\hat{i} = 1$. We consider two cases without learning where (2.9) and (2.10) determine expectations: one without expectations herding ($\tau = 0$), and one with expectations herding ($\tau \neq 0$). Finally, we consider two cases with learning where (2.16) and (2.17) determine expectations: again one without expectations herding and one with it. Our focus is one how expectations herding and learning affect the volatility of consumption, its correlation with output, and the dispersion of expectations.

Our baseline calibration sets $\rho = 0.95$ and $\beta = 0.95$, common values in the literature. We set σ_u, the standard deviation of productivity shocks, to 0.05. The steady state level of output is set to one, alternative values have virtually no effects on our results. We set γ, the gain for the learning versions of the model, to 0.02. This is similar to agents using rolling window window of about 50 periods. The parameters τ and j are novel to our paper. We fix them at 1 and 0.05 respectively and then show the effects of alternate values.

3.1 Alternate Herding/Anti-Herding Preferences

We first consider the effects of changing τ, the desire to herd, while holding j constant at 0.03. Figure 1 shows the effects on the average value of $\hat{\iota}$. Note that the no herding cases and the rational expectations case are independent of τ and the results are thus constant. Focusing on the the model with learning, the results for the model with expectations anti-herding (τ) seem surprising. In each period, ant-herding pushes ι closer to one-half. Yet under anti-herding, the average value of ι across periods is further from one-half.

Figure 1: Effects of Varying τ on $Mean(\hat{\iota})$

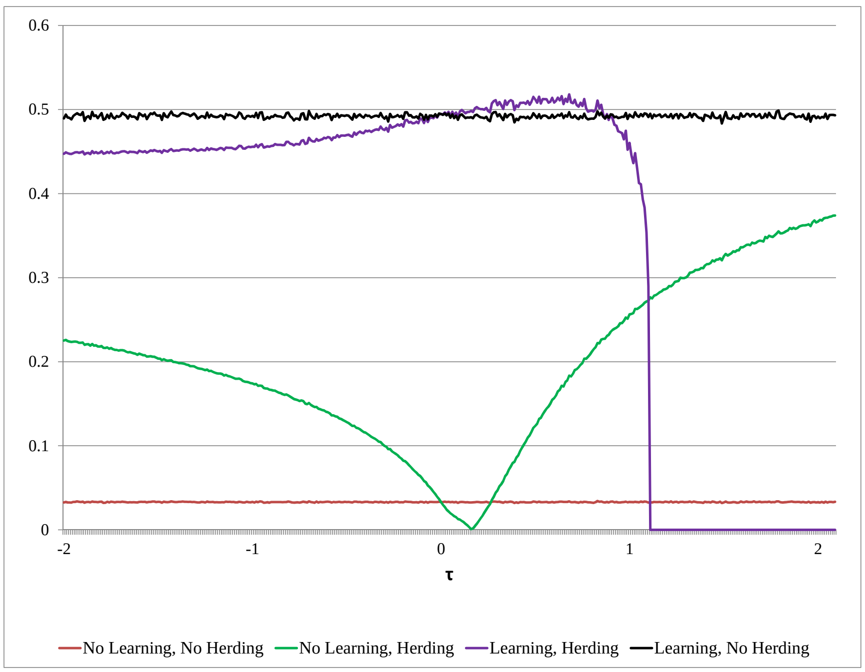

This result occurs because, without expectations herding, the model is prone to periods of high volatility. When random productivity shocks drive the estimated benefit of the AR(1) forecast very high, $\hat{\iota} = 1$, and $c_{1,t}$ is as volatile under rational expectations. Ant-herding, however, reduces the frequency of this high volatility state. This causes $\hat{\iota} = 1$ to average a lower value under anti-herding even though it is closer to one-half in every period. As shown in Figure 2, because it reduces the frequency of this high volatility state, strong anti-herding reduces consumption volatility.

Figure 2: Effects of Varying τ on $StDev(c_{1,t})$

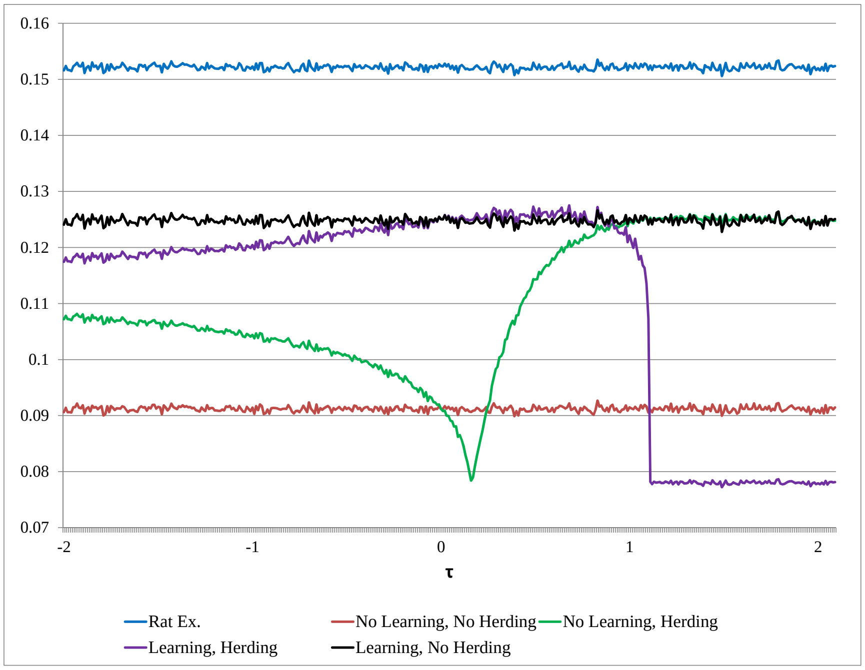

As τ rises above zero and expectations herding starts to occur, the high volatility state where all households choose the AR(1)forecast occurs more often. The average value of $\hat{\imath}$ and $c_{1,t}$'s volatility are thus higher. For this set of simulations, once τ rises above 1.10, an abrupt change occurs. Now expectations herding is so strong that full herding occurs and the model spends all of its time at $\hat{\imath} = 0$ where consumption volatility is minimized at a level equal to $\frac{1}{1+\beta}$ of the rational expectations level.

The model getting stuck at $\hat{\imath} = 0$ instead of $\hat{\imath} = 1$ made us worry that this could be a result of the starting parameters we used to initialize the learning algorithm. We thus repeated the simulations with starting parameters that endowed households with an implausibly strong preference for the AR(1) forecast. But for $\tau \geq 1.11$, all agents nevertheless switched from the AR(1) to the naive forecast within 200-400 periods.

Figure 3 reports the effect of varying τ on the correlation between young households' consumption and output. Prior to the point where $\hat{\imath}$ becomes trapped near near zero, both learning versions reduce this correlation below 1. Once $\hat{\imath}$ becomes stuck near zero, however, $c_{1,t} \approx y_t$ and the correlation is again very lose to 1.

Figure 3: Effects of Varying τ on $Cor(y_t, c_{1,t})$

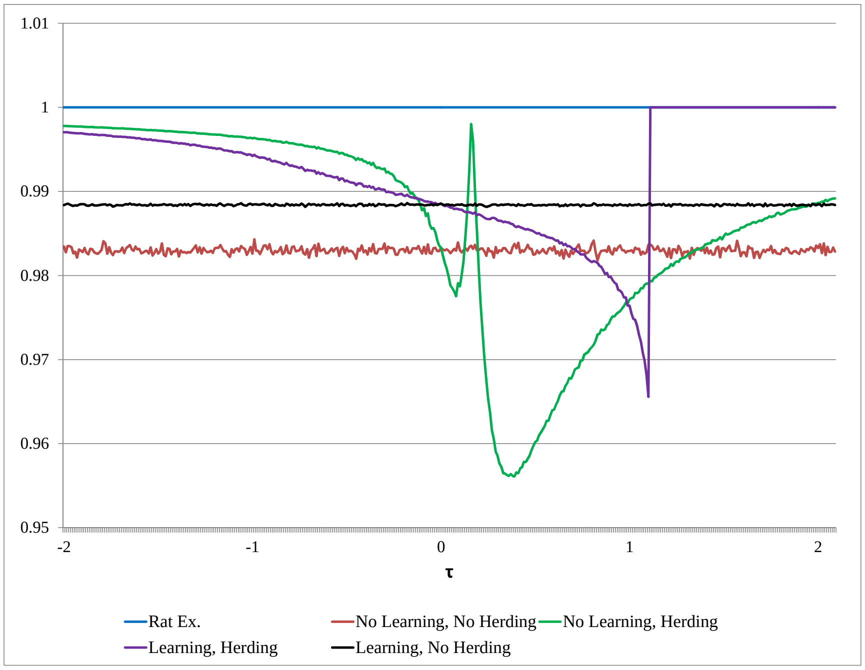

3.2 Volatility Cycles

For values of τ greater than 1.10, our use of adaptive learning causes the model to exhibit two distinct states. The *high volatility state* occurs when productivity shocks cause the estimated difference in mean-squared errors to be large enough so that all agents choose to form AR(1) expectations. Here, the model is identical to rational expectations. $c_{1,t}$ is perfectly correlated with output, and its volatility is maximized.

The *moderate volatility state* occurs when the difference in estimated mean squared forecast errors is smaller and some agents choose to form naive expectations. Here, $c_{1,t}$'s volatility is lower. The properties of these states are highly similar regardless of the value of τ. Table 2 shows the volatilities for each variable in each state for a simulation where $j = 0.03$ and $\tau = 0.5$.

Table 1: Standard Deviations Under Each State

	y_t	$c_{1,t}$(Rat. Ex)	$c_{1,t}$(Learning)
Moderate Volatility State	0.133	0.133	0.113
High Volatility State	0.197	0.197	0.197

Output and consumption under rational expectations are also more stable when the learning model is in the moderate volatility state. This is because periods of output stability tend to lower the benefit of choosing the AR(1) forecast and thus cause the low volatility state under learning.

Note, however, that learning causes a further 15% drop in consumption volatility (from 0.133 to 0.113). Learning thus acts to further dampen random periods of lower volatility.

Herding and anti-herding affects the frequency of high volatility state. When anti-herding is strong enough, $\tau < -0.9$, we never observe the high volatility state. Increasing τ then causes the high volatility state to occur more often until it occurs most frequently, 18% of the time, around $\tau = 1.1$. This is intuitive, if the model yields $\iota = 0.9$ without herding, the desire to conform to the population's average expectation might induce the remaining 10% of agents to pay for the AR(1) expectation if herding is strong enough.

Figure 4: Frequency Spent in Each State

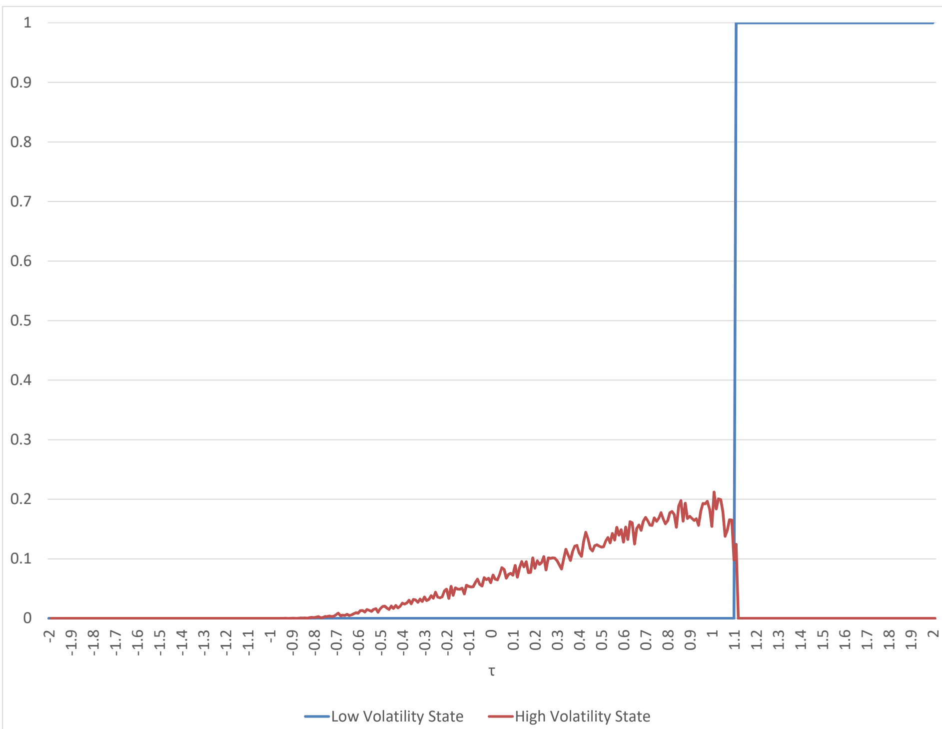

Surprisingly, an abrupt transition occurs once τ rises above 1.10. Once herding is strong enough, we very rarely observe the high volatility state. Instead the model always quickly converges to a third state, the low volatility state, where all agents form the naive expectation.[6]

Table 2: Standard Deviations in the Low Volatility State

	y_t	$c_{1,t}$(Rat. Ex)	$c_{1,t}$(Learning)
Low Volatility State	0.152	0.152	0.082

[6]The transition is exceptionally abrupt. At $\tau = 1.1052$, a simulation of 10,000,000 periods never yields the low volatility state. By $\tau = 1.1053$, however, the model spends 96% of the time in the low volatility state.

Convergence to the low volatility state always occurs even when we impose initial conditions on the learning algorithm where agents strongly prefer the AR(1) forecast. This low volatility steady state is never observed when $\tau < 1.11$. Thus while herding initially increases volatility, there exists a threshold level of herding that causes a dramatic and abrupt decline in volatility.

The threshold value of τ where the model transitions to the low volatility state occurs at the point where even agents with $\iota \approx 0$ will start to choose the naive expectation if all other agents do so as well. The desire to herd is strong enough that even periods of high volatility will not induce them to switch and the model remains in this state for the rest of the simulation (5,000,000 periods).

3.3 Alternate Costs of Forming the AR(1) Expectation

We next the effects of varying j, the disutility from forming the AR(1) forecast. For each value of j, we simulate all five versions of the model for $1,000,000$ periods. The rational expectations version of the model is equivalent to imposing $\iota = 1$ so that (2.7) determines consumption, and it is unaffected by changing j. Rational expectations thus exhibits a perfect correlation between $c_{1,t}$ and y_t, as well as the highest volatility of consumption of all versions. Figure 5 reports the average share of households paying to form the AR(1) forecast.

Figure 5: Effects of Varying j on $Mean(\hat{\iota})$

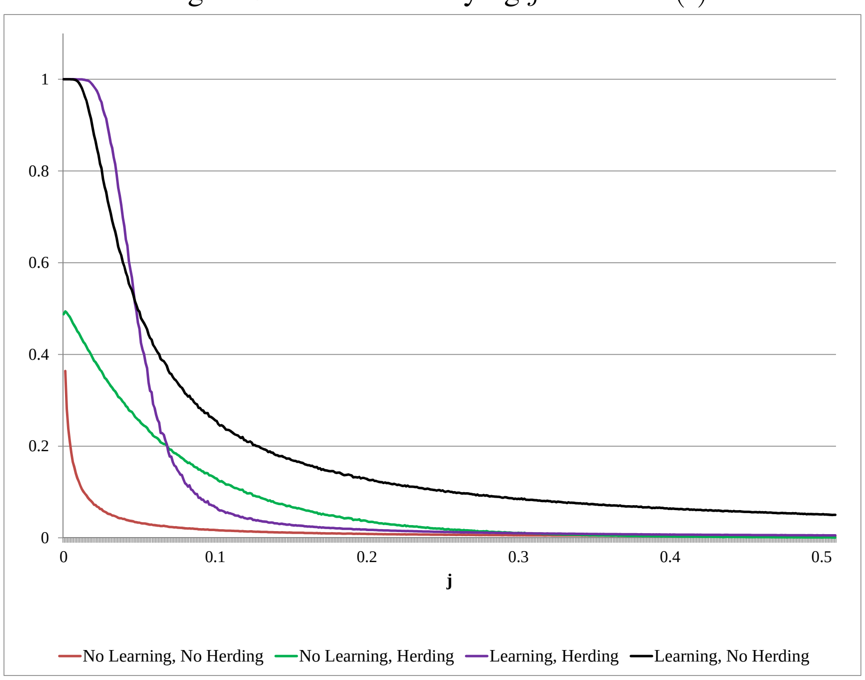

Unsurprisingly, as j becomes large, more households in all four boundedly rational versions of the model choose the naive forecast, and as $\jmath \rightarrow \infty$ $\hat{\iota} \rightarrow 0$ and all four versions become identical. It may seem surprising that, without learning, including herding ($\tau = 1$) leads to a higher average value of ι and one that is closer to one-half, suggesting more dispersed expectations. This is driven, however, by our selection of the interior equilibrium from (2.10) instead of corner solutions. Figure 5 shows that herding moves ι away from one-half. For low values of j, AR(1) expectations are cheap, and $\iota > \frac{1}{2}$. Herding then causes some households to shift to the AR(1) expectation, increasing ι toward 1. For higher values of j, however, the AR(1) expectation is expensive so $\iota < \frac{1}{2}$. Herding now moves some households towards the naive expectation, pushing ι down toward zero.

Figure 6 reports the standard deviation of $c_{1,t}$. These results follow from the average value of ι, higher values cause consumption to respond more to productivity shocks, increasing consumption's volatility. Focusing only on the learning versions, herding thus causes there to be increased consumption volatility for low values of j, but reduced volatility for higher values of j.

Figure 6: Effects of Varying j on $StDev(c_{1,t})$

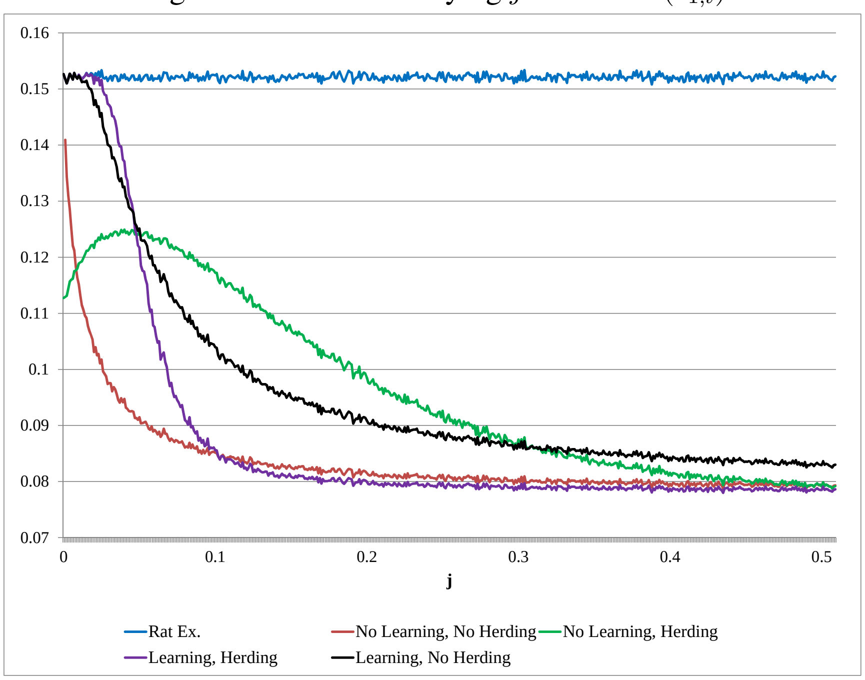

Rational expectations not only yields the highest levels of volatility, but also yields a perfect correlation between $c_{1,t}$ and y_t. Figure 7 shows that all four boundedly rational versions reduce this correlation.

Figure 7: Effects of Varying j on $Cor(y_t, c_{1,t})$

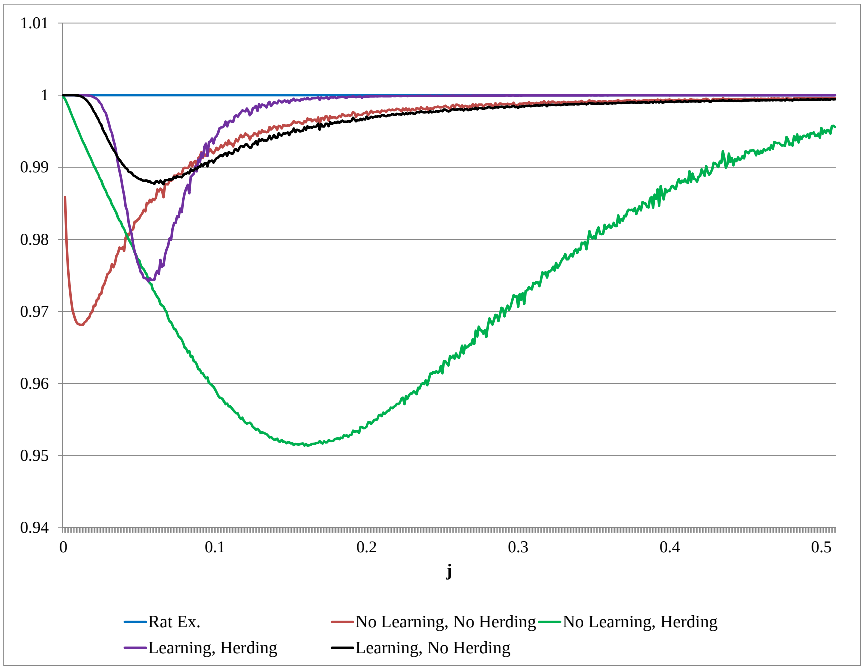

4 Conclusion

An extensive empirical literature demonstrates both the presence of herding and ant-herding in different settings. This paper has developed a new approach for modeling herd behavior by assuming that boundedly rational households must choose whether or not to pay a cost in order to allow them to form the best possible forecast. If the desire to herd is strong enough, then agents will coordinate on the low information state where the model underreacts to shocks.

This paper has used a simple endowment economy to illustrate the potential impacts of herding or anti-herding. Future research may extend this framework into other applications. One example is a real business cycle model where agents choose their labor supply as well as consumption. As in the OLG model, herding initially causes the high volatility state to occur more often. Above a threshold level of herding, however the model again abruptly transitions to always being in a low volatility state. Now, however, lower consumption volatility corresponds to higher investment volatility which also leads to greater investment volatility. Furthermore, the low-consumption volatility state (with mostly naive forecasts) also exhibits a much lower correlation between consumption and output.

References

Akerlof, G., and R. Shiller. 2009. *Animal Spirits: How Human Psychology Drives the Economy and Why It Matters for Global Capitalism.* Princeton, NJ: Princeton University Press.

Bewley, R., and D. Fiebig. 2002. "On the herding instinct of interest rate forecasters." *Empirical Economics,* Vol. 27(3): 403-425.

Bikhchandani, S., Hirshleifer, D., and Welch, I. 1992, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy,* Vol. 100(5): 992-1026.

Clements, A., Hurn, S., and S. Shi. 2017. "An empirical investigation of herding in the U.S. stock market." *Economic Modelling,* Vol. 67: 184-192.

De Bondt, W., and W., Forbes. 1999. "Herding in analyst earnings forecasts: evidence from the United Kingdom." *European Financial Management,* Vol. 5(2): 143-163.

Evans, G, and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.

Fajgelbaum, P., Schaal, E., and M. Taschereau-Dumouchel. 2017. "Uncertainty Traps." *The Quarterly Journal of Economics,* Vol. 132(4): 1641–1692.

Galariotis, E., Rong, W., and S. Spyrouc. 2015. "Herding on fundamental information: A comparative study." *Journal of Banking & Finance,* Vol. 50: 589-598.

Lamont, O. 2002. "Macroeconomic forecasts and microeconomic forecasters." *Journal of Economic Behavior and Organization,* Vol. 48(3): 265-280.

Pierdzioch, C., and J. Rülke., 2012. "Forecasting stock prices: Do forecasters herd?" *Economics Letters,* Vol. 116(3): 326-329.

Pierdzioch, C., Rülke., J., and G. Stadtmann. 2010. "New evidence of anti-herding of oil-price forecasters" *Energy Economics,* Vol. 32(6): 1456-1459.

Rulke, J., Silgoner, M., and J. Worz. 2016. "Herding behavior of business cycle forecasters." *International Journal of Forecasting,* Vol. 32(1): 23-33.

Sargent, T. 1999. *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.

Scharfstein, D., and J. Stein. 1990. “Herd Behavior and Investment.” *American Economic Review,* Vol. 80(3): 465-479.

Shiller, R. 2000. *Irrational exuberance.* Princeton, N.J. Princeton University Press.

Trueman, B. 1994. “Analyst Forecasts and Herding Behavior.” *The Review of Financial Studies,* Vol. 7(1): 97-214.