Essays on Herd Behavior — Theory and Criticisms

Annika Westphäling*

“Four eyes see more than two”−that information gets more precise being aggregated from more people is proverbial. Problems occur only when the aggregation of information does not run smoothly, for example if information is not communicated directly but revealed through people’s actions. This is what the branch of literature studying imitative and herd-like behavior is concerned with. The concepts of herd behavior and informational cascades were first formalized in 1992 by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) (hereafter BHW). Herd behavior in financial markets and the resulting price bubbles had been observed since long before, without any available theoretical explanation; so when the paper by BHW was published, the model was embraced to serve as an explanation of these phenomena. However, while the results prove to be highly applicable to examples from cultural change to the spread of technology, the applicability to the functioning of financial markets is more delicate, as we will show in this essay.

The basic model involves agents acting one after another, maximizing their expected gain, the payoff depending on their decision and the underlying state. Their information consists of a private signal about the state and the actions of their predecessors. We call herd behavior the situation where a history of actions

* Annika Westphäling received her degree in Economics (B. Sc.) from the University of Bonn in 2010. The present article refers to her bachelor thesis submitted in August 2010.

leads an agent to take the same action as her predecessor, regardless of her private signal. Agents are ex-ante equal, in the sense that the precision of the signal received is the same for everybody; so when an agent starts herding, learning stops altogether, as her actions do not provide any additional information, leaving her successor in the same position where the evidence of the history outweighs any private signal. The conformal behavior of all subsequent agents is called an informational cascade.

A Classic Model of Herd Behavior

In the paper by BHW, individuals are asked whether to adopt or reject a certain behavior. Both the gain of adopting the behavior and the signal the agents receive can take either value 1 or 0, and for every agent the probability of receiving a signal that equals the real value is $p > 1/2$. The cost of adoption is constant at $1/2$ for all individuals, so an agent will adopt whenever she estimates the value being 1 more likely than the value being 0. In case of indifference, a coin flip will decide on the action. Consider the case where the first agent receives a signal 1 and adopts, thus revealing her signal. If the second agent receives a signal of 1, she will adopt; receiving a signal of 0, she is indifferent, and she adopts or rejects with equal probability.

The history (adopt; adopt) does not allow to perfectly deduce the history of signals, but indicates a high value of the behavior; indeed, it can be shown that the third agent adopts regardless of her signal, thus starting an informational cascade. If the third agent observes a history of actions (adopt; reject), she will know that the history of signals must have been $(1; 0)$, which is equally probable given a high or low value, so she finds herself in the very same position as the first agent. It follows that after an even number of agents we find ourselves in a cascade if the history features two more actions of one type than the other;

moreover, the probability of this occurring is one in the limit. The more precise the private signal, the more likely it is for a cascade on the correct action to start early on, as a precise signal favors subsequent agents receiving the same signal and thus taking the same optimal action, starting what we call a "fully revealing" cascade. However, we may not let this term mislead us: agents can only state with a certain probability whether the action they take is optimal, but learning stops as soon as the first agent engages in herd behavior.

Considering the high loss of information induced by both the coin flip and the short sequence of actions sufficing for the start of an informational cascade, it is not surprising that even with a precise signal $p >> 1/2$, the probability for a non-fully revealing cascade is very high.

Compared to the calamitous picture that is painted above, stock markets seem downright docile. In fact, both the strength and the weakness of this model lies in its simplicity: we will see that as we adapt the model to better account for the complexities of financial markets, the possibility for an inefficient outcome will be overruled.

First Critique: Modifying the action space

The first criticism to the direct application of BHW's model to describe how financial markets operate was brought in by Lee (1993) and it concerns the richness of the action space available to our decision makers. Compared to the coarseness of the adoption or rejection of a behavior, in financial markets, it is more sensible to assume a continuous action space, as the possible quantity traded can be approximated to a continuum.

Lee allows for a finite number of states, each of which is characterized by the probability of receiving a signal 1 it induces. The action space is a subset of the interval $[0, 1]$ and it includes these probabilities. The agents have to assess the odds of receiving a signal 1 considering their signal and the actions of their predecessors; formally, they minimize the expected squared difference between their action and the underlying probability, which is their loss function.

This leads us also to alterations concerning the definition of an informational cascade: here, an informational cascade occurs whenever the history of actions converges, whether in finite time or in the limit. In the first case, the conformal action is optimal only with a certain probability, so if a cascade is fully revealing it is so by chance, whereas if actions converge only in the limit, we will see that the limes is bound to be optimal given the underlying state.

Lee states that, with any discrete action space, the emergence of a cascade as per BHW is almost certain when the number of agents goes to infinity. This follows from the observation that the longer the history, the less an agent's action will deviate from her predecessor's. For the smallest gap in the action space, there exists a history length after which an individual is no longer ready to "jump over" this gap to accommodate her signal. So after any history of that length, we find ourselves in a cascade.

Furthermore, the positive probability of a cascade is equivalent to the positive probability of an inefficient outcome. In fact, any history is induced by a sequence of private signals. Every signal has a positive probability of occurring in any state. It follows that if there exists a finite history so that a given action minimizes the expected loss independently of the signal, this history has a positive probability of occurring in any state, hence also in a state where the action is not optimal.

Moreover, note that the non-existence of such a history is sufficient for an optimal outcome. If the action somebody takes is never the same under a high as under a low signal, an agent's action reveals her private information. It follows that agents can deduce the history of private signals from the history of actions, and the strong law of large numbers states that by increasing the history length we can almost certainly deduce the true state of the world. With a continuous action space and a non-degenerate prior, an agent will always adapt her action according to her signal, albeit to an arbitrarily small degree as the length of the history grows. As we see above, this fact eliminates the possibility of herding and guarantees an efficient outcome, thus setting one major critique to the applicability of "standard" herding models to financial markets.

Second Critique: Introducing a Pricing Mechanism

Another critique to the direct applicability of BHW-like models to financial markets is founded on the assumption that in a functioning financial market, all public information is reflected in the price. In Lee, the action chosen, if informative, reveals the signal received and it equals the public belief. In the trading model by Avery and Zemsky (1998) (hereafter AZ), the estimation of the predecessor, which equals the public information, is reflected in the price, and an agent's buy or sell order reveals whether her assessment is higher or lower than her predecessor's.

Consider a financial market where agents sequentially trade one asset that can take value 1 or 0. Differently from BHW, where the cost of adoption is constant, here the price at which to buy or sell an asset depends on the history of trading. In each period, one agent from a continuum of traders is randomly selected, where λ is the fraction of informed traders, and $(1 - \lambda)$ the fraction of noise traders.

The agents may either buy or sell one unit of the stock or refrain from trading, and we assume that noise traders choose each action with equal probability; on the other hand, informed traders are risk neutral and maximize their expected profit updating their beliefs about the value of the asset after observing the history and a signal having precision p.

We need to modify the definition of herding to account for flexible prices. For

AZ, an agent engages in herd buying, if two conditions are met: first, the expected value of the asset given the initial prior has to be higher than the prior updated with her signal, so that if the agent had been the first to trade, she would have sold. Secondly, the market maker estimates the value of the asset higher than in the beginning, so the trading history must have been positive, what implies a preponderance of buying orders.

So for herding as per AZ, a signal which would have had a negative effect on the expected value of an asset in the beginning has to lift the expected value of the asset even above an ask price, which is at least as high as the estimated value of the asset given the history.

At any point in time, having observed the same trading history unfolding and having exactly the same information about it, both the traders and the market maker have the same valuation of the asset. However, once called to trade, informed traders receive a signal which moves their valuation either above or below the market maker's. The latter fixes bid and ask prices conditionally on receiving a sell and a buy order respectively. As she makes zero profits in equilibrium, and due to the presence of noise traders in the market, the market maker will always fix an ask price lower than the valuation of a trader with a high signal and a bid price higher than the valuation of a trader with a low signal. Therefore when prices are not fixed but change with the trading history and where there is only uncertainty about the value of the asset traders always follow their signal and herding is not possible, and even though the existence of noise traders blurs the information, in the long run the history of trades will reveal the true value of the asset.

To allow for herd behavior in financial markets, however, AZ offer an extension of their model, accounting for the fact that traders may be informed of events that have an impact on an asset's value, whereas the market maker is not. Traders in turn do not know for sure whether the reported event, hereafter referred to as information event, decreases or increases the value of the asset, but have to rely on the history of trades and on their private signals that work by the same principle as already established in BHW. The true value of the asset can now be either 0, $1/2$ or 1, and informed traders receive a signal $1/2$ if and only if the true value is $1/2$; otherwise, the signal is correct with probability $p > 1/2$. The market maker can never exclude $1/2$ as a possible state but at this point traders can and they interpret the trading history differently from the market maker. If the probability of an information event is small, the valuation of the market maker stays close to $1/2$ while traders' valuation can diverge to a point in which one of them could buy with a signal 0 or sell with a signal 1. Price rigidity is recreated in proportion to how close the prior probability of state $1/2$ is to 1, allowing for an arbitrarily long sequence of herd behavior that follows. During this time the traders' valuation does not move, and the market maker's valuation re-aligns with the traders' until normal trade recovers.

Note, that in this respect, herd behavior is even informationally efficient. Creating conform behavior of informed agents helps the market maker realize that the value of the asset has deviated from $1/2$. In the long run, however, the price smoothly converges to the true value unless a new shock arises. It follows that uncertainty about an information event fails to link herd behavior to catastrophic market events and extreme price distortions.

References

Avery, C., and P. Zemsky (1998): "Multidimensional Uncertainty and Herd Behavior in Financial Markets," *The American Economic Review*, 88(4), 724–748.

Banerjee, A. (1992): "A simple Model of Herd Behavior," *The Quarterly Journal of Economics*, 107(3), 797–817.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100(5), 992–1027.

Lee, I. H. (1993): "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61(2), 395–411.