How will they fare on "Judgment Day"? Reputational Herding in Professional Communities

Alexandra Oprea Adviser: Michelle Connolly Adviser: Hussein Yildirim

Abstract

This paper presents a theoretical model of decision-making by professionals within the same field. The decision-makers consider reputation effects that result both from their relative actions and from the absolute outcomes. By combining the idea of relative actions from the managerial reputation herding models and that of outcomes from the information cascades models, one can make comparisons across a wide range of professions. The importance of outcomes, as a proxy for the relative importance of the two considerations, is a function of time. The model hopes to predict a herding equilibrium is more likely when 'judgment day' – the time when the outcome is reveled – is far into the future.

I. Introduction

It is the beginning of senior year and students are busy making decisions about the future. One option is to write a senior thesis. For some talented researchers, it's an easy choice – they *want* to conduct research. Mike, however, is just an average student looking to score his Washington dream job. With job applications and interviews happening next month and thesis completion a distant affair, should he try to signal his talent to the market by starting such a project? Had interviews happened after his project was graded, everyone would know exactly what Mike's ability was. How does the timing of interviews change the incentive structure?

Such reputation concerns often loom large in the minds of individuals, particularly in professional contexts where one is looking to establish oneself as an expert. Assuming all the previous seniors have graduated and are twenty years later happily employed in their respective fields; let us compare Jenny's and Mark's situations. Jenny is a doctor recommending a treatment to a patient. She selects a more experimental procedure than the majority of her peers, and they snub their noses at her. But low and behold – soon thereafter medical tests confirm her recommendations were right. Her reputation goes up significantly and wealthy patients line up to get an appointment. Mark, however, is a macroeconomist working in the National Central Bank (NCB). He is recommending a monetary policy stance which differs from that of his peers. Much like Jenny's colleagues, his peers show disbelief in his ability.¹ If the policy were to be implemented, years would be spent trying to evaluate the effects of his suggestion. During this time economic conditions may change, financial crises may hit, and the world can change many

¹ We can here also talk about another graduate, Sarah, who is an artist. In the art profession, the presumption of bias towards the status quo can be dropped and replaced with a symmetric assumption of bias against the status quo. In other words, in these examples, innovative actions are wrong until proven right, while in the art world they are right until proven wrong.

times over. Having difficulty in judging whether Mark was right, the NCB would lose confidence and change its stance to the more conventional one. Mark's reputation goes down as he becomes an outlier and his career prospects look dim. Both professionals take a bold stance against the conventional methods. What accounts for their different success? Was Mark's action wrong, even though he firmly believed his suggestion was appropriate?

In the example outlined above, professionals concerned with their reputation make a decision, taking into account two factors. First, they consider the decisions of their peers and whether their own prescriptions match or deviate. Second, they are concerned with the validity of their decision which will only be discovered at a future time. In Jenny's case, both aspects end up playing a significant role. She gets negative reputation by going against her peers, but shortly thereafter receives significant reputation benefits by being proven right. In Mark's case, however, the outcome does not get revealed and the market has only his decision and that of his peers to base its evaluation on. As a result, he only has the negative reputation from going against the crowd. This paper aims at the broadest level to understand the difference between these two situations. To do so, however, it begins with a simple model similar to Mark's situation which can then further be extended on a profession-by-profession basis.

I start off by outlining a very specific case of two central bank governors within different countries making decisions about monetary policy. We can think of them as two experts within the same profession: central banking. The two decision-makers receive signals about what an appropriate policy would be. They act sequentially with the second player observing the action of the first. To make the situation clear, imagine the first player is the European Central Bank's president (viewing him as a unitary actor does not affect the model), whose reputation is already established in the labor market to be high: he is considered 'smart'. The second player is the

Romanian National Central Banks' president, who can be either 'smart' or 'dumb' with the same prior probability. The question this model serves to answer is: What will the second decisionmaker rationally choose to do? Will he ignore his own signal and 'herd' with the ECB president or will be follow his own signal regardless of what the ECB decides to do? What if the second agent has more information about his own type?

II. Literature review

Keynes (1936) was the first economist to remark that "[w]orldly wisdom teaches that it is

better for reputation to fail conventionally than to succeed unconventionally", introducing the possibility of reputational concerns playing a role in decision-making.² A number of theoretical models have been developed to deal with the phenomenon of herding, described by Scharfstein and Stein (1990) to represent the case when "managers mimic the investment decisions of other managers, ignoring substantive private information."³ Most theoretical models have not been concerned with defining the general phenomenon of herding they are explaining, preventing an optimum generalization of their results beyond the specific circumstances of their model. I expand the Scharfstein and Stein (1990) definition to cover *situations in which professionals adopt the actions taken by others in their field, regardless of whether they privately believe those actions to represent the best course of action with regards to outcome.*⁴ I make the last specification in order to differentiate between concerns about outcome and concerns about reputation. A professional's decision to 'herd' may not produce the best outcome, but it can provide the most reputation benefits.

² Keynes, J. M. "The General Theory of Employment, Interest and Money" (1936: 158)

³ Sharfstein and Stein (1990) p. 465

⁴ There is a particular situation in which the professional does not hold any personal beliefs about the appropriate course of action. In other words, he has no private information. In this scenario, we will by default consider him to be herding, without any loss of generality to the model. For more on information gathering costs, there will be (hopefully) an extension of the model in Part IV.

In their interdisciplinary integrative review of herding, Parker and Prechter (2005) classify the theoretical approaches to explaining herding from different disciplines on eight dimensions. The model presented in this paper retains the assumption of rationality in order to explain behavior, accepting the fact that other circumstances than those presented here may require a loosening of this assumption. In doing so, I incorporate the importance of choosing the right outcome from models using information cascades and the importance of convincing an outside observer (call it the 'labor market') of your ability from the social psychology models (also referred to as principle-agent models).

In information theory models, the seminal examples are Bikhchandani, Hirshleifer et al. (1992), heretofore referred to as BHW, and Banerjee (1992). In both models information externalities represent the key explanatory variable since agents receive information from observing the actions of others. The conformity of behavior or herding is described using models of informational cascades. BHW define an informational cascade as occurring when "...it is optimal for an individual, having observed the actions of those ahead of him, to follow the behavior of the preceding individual without regard to his own information."⁵ This definition of information cascades is the same as the definition for herding in the reputation model. The difference only comes through in the reason for the cascade. Rather than herding to increase reputation, agents in these models herd because they trust the information revealed by the actions of previous movers. The only concern of decision-makers in these two information cascade models is to select the right outcome (the one with the intrinsic maximum payoff), and their payoff is independent of the choices made by others. ⁶ The other agents matter only by adding

⁵ Bikhchandani, Hirshleifer and Welch (1992) p. 992

⁶ "There is a set of options represented by a line segment, and within this set there is one correct option. The aim of the game is to find the correct option." Banerjee (1992:799)

information about which decision is correct. I will now briefly review the two models since they represent the foundations of thinking about herding.

In the model of Banerjee (1992), there are N agents making a choice between multiple options or assets. There is a unique asset with a positive return and all others give a payoff of zero. These N agents maximize identical risk-neutral utility functions defined on the space of asset returns and use both their private information and the information conveyed by the actions of others before them in order to make their choice. The 'herding externality' appears because the choices made by agents are not "...sufficient statics for the information they have."⁷ In simpler terms, the actions don't show the underlying information on which the decision was made. One feature that is underemphasized in Banerjee and later critics like BHW is the timing of discovering the payoff. For the purposes of my paper, timing plays a crucial role in explaining the equilibrium. In Banerjee (1992), it is clearly stated that all the choices of individuals are only tested after everyone has made their decision. At that point, if any of the choices "turn out to work", the respective individuals receive their rewards. The moment when your choice 'turns out to work' is very important, since information cascades would not occur if people could immediately discover whether the previous choices were right or wrong. To my knowledge, there has been no research on exactly what the impact of the timing of this revelation is on herding behavior.

In BHW, there is also only one right choice, but it is between only two alternatives: adopting or rejecting some behavior. All individuals have cost C and gain V (which is either 0 or 1) from adopting the behavior and the goal is to figure out what is the correct value of V. The mechanism for information cascades forming is the same as in Banerjee (1992), but the fragility

⁷ Banerjee (1992: 809)

of cascades is emphasized in this model, since public information releases can completely overturn cascades. These models work well in predicting behavior when the choices we make are not taken to say anything about our ability. This would only happen in one's personal life. Making the wrong choice may means you lose some money or forgo a meal at the better restaurant, but it doesn't make the market reevaluate your ability to make decisions (it might, however, make your spouse do so). This reputational concern, however, is very real within one's professional life.

The reputation or principle-agent models explain herding as the result of professionals, narrowly defined as managers, wanting to avoid being revealed to be of low-ability. Scharfstein and Stein (1990) is the foundational paper for this literature and contains many of the elements I will be using in the model described in this paper. Managers can be of two types: 'smart' or 'dumb', and each receive a signal about a potential investment. Managers themselves are unaware of their own types, a plausible beginning-of-career assumption. The informative (true) signals received by smart managers are correlated with reality, while the uninformative signals (noise) received by dumb managers are random. From the beginning, being against the herd is a sign that your signal is random and increases the probability of the market judging you as having low ability⁸.

The Scharfstein and Stein (1990), Banerjee (1992) and BHW(1992) can all be described as 'follow the leader' models of herd behavior.⁹ However, the more agents know about their type and the certainty of their private information, the more incentives appear to put in motion the

⁸ <u>Scharfstein and Stein (1990</u>) have inspired a number of other models of reputational herding [<u>Trueman (1994</u>), <u>Zwiebel (1995</u>)]. <u>Ottaviani and Sørensen (2000</u>) generalize beyond the differential conditional correlation, arguing that reputationally concerned agents always herd, unless they know enough a priori about their own information quality. In the last case, the conditional correlation informing the 'sharing the blame' effect is necessary ⁹ Trueman (1994: 4)

opposite behavior: anti-herding. A number of recent models have attempted to generalize the situations under which anti-herding may occur [Avery and Chevalier (1999), Levy (2004), Effinger and Polborn (2001)]. These models present a much more nuanced approach and will be helpful in looking at extensions for the present model. Effinger and Polborn (2001) effectively summarize the underlying concepts in strategic reputation models as follows. Managers do not care about successful investment per say, but want to appear to be smart. Making the right decision is an indicator that a manager is smart, but since smart managers receive the same signal, unanimity is also an indicator of smartness. The model in this paper directly introduces time concerns in order to understand the relative importance of these two indicators of ability. Specifically, the further into the future the accuracy of the decision is decided, the more important unanimity becomes and vice-versa. This insight will allow for applications of the herding models across different types of decisions and different professional environments.

III. Theoretical Framework

i. The Model

The set-up of the information structure in this model closely follows Scharfstein and Stein (1990). It applies most literally to the specific example briefly presented in the introduction of two central bankers making decisions about monetary policy. The entire model is summarized in Figure 1 below.

Figure 1: Which Game are We Playing?

Players	Signal	Predictive	Realized State	Outcome
(dumb, smart)	(left, right)	Power	of the World	
The Expert	Signal is	Predicts correctly	Left	Choosing Left
(smart ex-ante)	informative	with probability		results in payoff
		-		
			Right	Choosing Left
				results in payoff
The Newbie	Signal is	Predicts correctly	Left	Choosing Left
(dumb ex-post)	uninformative	with probability –		results in payoff
			Right	Choosing Left
				results in payoff
The Newbie	Signal is	Predicts correctly	Left	Choosing Left
(smart ex-post)	informative	with probability		results in payoff
		-		
			Right	Choosing Left
				results in payoff

A. Timing and Information Structure

There are two central bankers, called A and B respectively. These central bankers make decisions about monetary policy sequentially, with A moving first. The two central bankers can be of two types: 'smart' or 'dumb'. Being 'smart' can be interpreted as having the necessary skills and education required in monetary policy making. Being 'dumb' can be interpreted as being less skilled or less educated in monetary policy making. The prior probability of being 'smart' is θ , while the prior probability of being 'dumb' is *1*- θ . In this particular model, we assume that central banker A, modeled as a recognized expert, has already developed a reputation for being 'smart'. Uncertainty, however, remains for B, modeled as a newcomer to the field with potential to either be a great expert or a blundering failure.

At t=1, A decides whether or not to use a particular instrument of monetary policy. This can be interpreted as a conventional instrument of monetary policy such as changing interest rates or conducting open-market operations. This is a binary choice between using and not using this instrument and will depend on the underlying conditions prevailing in the world. Under some conditions, using the instrument will be a good idea and provide tangible benefits. Under different conditions, using the same instrument would be ineffective or, even worse, detrimental.

For simplicity, I choose use a higher level of abstraction and represent these two choices as either going "Left" or "Right". There are then two possible outcomes at date *T*: either the 'Left' state of the economy, in which case the action "Left" was optimum and the action yields a payoff , or the 'Right' state of the economy, in which case the action "Left" is suboptimal and the action yields a payoff It should be pointed out that these are payoffs for the decision maker, not for the economy as a whole. The prior probabilities of these two states are and respectively. The outcome is publicly observable even if neither central banker chooses to use the policy instrument. To understand the end-game situation better, I provide the simple decision-tree in Figure 2 below.





In making his decision, A has access to a private signal, which can take on one of two values: (a "Left" signal) or (a "Right" signal). The signal is only informative to 'smart' decision-makers. Formally, we can describe this for A to mean that

In the case of player B, we have the added uncertainty over whether he is a 'smart' or 'dumb' decision maker, captured by the variable θ . In the case in which player B is dumb, his signal is uninformative, meaning that he has a random chance of correctly predicting the outcome, captured below.

However, if player B is smart, then his signal (let's call it q) will also be informative.

The relationship between p and q is not specified at this point. In this model, I will be considering the situation in which . This means that all experts who are smart reach the same level of accuracy in their signals and thus their estimation of the future state of the economy. This scenario is reasonable, since we can assume that - the probability that an expert is wrong – can be the same idiosyncratic value for all smart individuals. In an extension to this paper, one could allow for the possibility that - implying that different experts can have different abilities to interpret the conditions and thus different accuracies. For the purposes of this model, however, we continue to assume that - implying that different accuracies.

For each player, we know they can use Bayes' rule to calculate the posterior probability of each state of the economy given the private signal which they have received. As a final result, we can observe for instance the posterior probability of the "Left" state for player A conditional on the signal that he has received to be

To simplify the calculations, we can assume that the prior probability of each state of the world (to be equal to –. This is a situation in which there are no reasons to believe that one particular state happens more often that the other – which corresponds to the prior

expectations of the players. Biased models can be constructed that assume one state to be more common and the entire analysis that follows can be reproduced by the interested reader keeping α as a parameter. For -, we get the following results for the probability of , conditional on the signal received

-

We can also calculate the posterior probability of a "Right" state after receiving the private signal. Knowing that , the posterior probability of the "Right" state after receiving a "Left" signal for player A is . Similarly, the posterior probability of a "Right" state after receiving a "Right" signal is simply

•

After A makes a decision at time t=1, at time t=2 it is B's turn to decide. He will consider the information given by his own signal, but will also be interested in what A has done. Let's start with just his signal, without taking into account player A's action. Since B does not know whether he is 'smart' or 'dumb', the posterior probability of the "Left" outcome occurring from before becomes

Even before taking the action of player A into account, player B's probabilities of getting the appropriate prediction of the state of the world (in other words, of being correct) is lower than A's due to the added uncertainty over his status as 'smart' or 'dumb'. This uncertainty exists both for player B himself (who does not know whether he is 'smart' or 'dumb') and for the outside observer (the labor market).

B. Utility Functions

Central bankers in the model have an overriding concern of gaining reputation – and as much of it as possible. Reputation can be translated into job security, as well as potential for career advancement. Alan Greenspan's reputation helped him keep his job as Chairman of the Fed for almost 20 years.¹⁰ One form of reputation, applicable in particular to a new comer in a

¹⁰ Not to mention the proceeds he made from selling his autobiography. As an aside, it's a delightful read. One would never have imagined him to be a talented jazz musician.

professional community, is measured as the labor markets' updated probability that the central banker is 'smart'. This revised probability is and is a function of three things:

- A's choice of a course of action:
- B's choice of a course of action:
- The labor market's perception of the likelihood of herding:

Therefore, in this model, the central banker will make decisions in order to maximize the value of . The labor market initially revises its probability that the up-and-coming central banker is smart by comparing the relative actions of the two players. It also considers the possibility that the consensus between A and B may be the result of herding. These 'reputational effects' (gains or losses) will continue to accrue to the player, discounted at each period by a factor of . This implies that central bankers prefer to receive 'reputational effects' in the present rather than in the future.

However, at moment T, the prevailing state of the world is also revealed and the labor market can then revise its reputational rewards to both players by also taking into account whether the decision is correct. This moment is referred to as 'judgment day', when the players are 'judged' by the labor market, taking all available information into account, and are rewarded or punished depending on whether they made the correct call.

Putting together these two measurements of reputation, one can obtain a general-form utility function. This total utility function will have two components, namely intermediary reputational effects and judgment day evaluation reputation. This will closely resemble B's utility function: Intermediate Reputational Effects

Judgement Day

Because of the difference in experience between A and B, the importance of these two components of the utility function will vary. For A, the reputable expert in central banking, the intermediate reputation gains do not matter since his reputation has already reached a steady state and it will not be affected by not taking the same action as his peer. As a result, A's utility function is just concerned with the final results on judgement day:

Having outlined the timing of the game and the utility functions that the players seek to maximize, it is time to consider what the equilibrium will be.

C. Equilibrium

Model only with Judgement Day

The first model we consider focuses only on the final utility component. The model without intermediate reputation effects is essentially a 'judgement day' model in which there are no intermediate reputation gains to be taken into consideration since the market's evaluation of a player's ability compared to that of the expert does not factor in. This can be the case if the state of the world is revealed right away.

We start off at T=1 and presume that player A has just received the "Left" signal. The question we need to answer is: Under what circumstances will A choose to take action "Left"? In other words, we want to know whether A will always prefer to when he receives this

signal. We assume that choosing "Left" in the state of the world gives the player a payoff of , while choosing "Left" in the state of the world gives the player a loss of ε^{11} . To understand the decision, we calculate the expected utility of A as

12

We can now compare the expected utility from and conditional on observing . We can set equal to , which implies that

Since we know that p, the choice of following one's own signal only depends on the value of p, as follows:

, if -, if -, if -

¹¹ We know that assume that

¹² JDU stands for Judgement Day Utility.

Given these possibilities, the probability of choosing "Left" after receiving signal "Left" increases with p. If the prior probability of each state of the world is equal to – and private signals are uninformative at –, player A is indifferent between the two options. However, within the theoretical constraints of this model, this will not be the case. The model is set up so that p is always greater than – (we expect experts to be better than coin tosses). Therefore we can construct proposition 1 below. The proof is the same for the alternative situation in which player A receives signal .

Proposition 1: Without reputational concerns, player A will always choose to follow his own signal.

Player B is in a similar situation. The difference is that he acts at t=2, after having observed player A's action. Given Proposition 1, he knows that player A will *always* follow his signal, so he can deduce the signal from the choice and use it in the evaluation of the prior probabilities of each state.

We are still assuming that choosing "Left" in the state of the world gives the player a payoff of , while choosing "Right" in the state of the world gives the player a payoff of ε . I take here the particular situation in which Player A has taken action and Player B's private signal is , since this is the most interesting situation for comparison with the model with intermediate reputational concerns. The set up in which Player A has taken action and Player B's private B's private signal is is identical to this one¹³. The expected utilities for each choice are

¹³ For the situations in which Player A has taken action and Player B's private signal is or in which Player A has taken action and Player B's private signal is , there is no herding to be analysed, since the player will always follow the direction of the two signals.



The comparison for player B is less straightforward than for Player A, mostly because he has to take into account both his own signal and the action of Player A acting before him. In order to make the calculations easier, we can give values to the end-state payoffs and ε . All that is important for our purposes is that $>\varepsilon$, so we can assign the values =1 and ε =0. This

implies that the player gets a positive payoff of 1 for being correct and gets nothing if he is wrong. Now, we can simplify the judgement day utility functions to

We can now compare the expected utility from and conditional on observing and player A's choice . We can set equal to , which implies that

Since we know from the model set-up that -, whether player B chooses and follows his own signal or and herds depends on the value of θ . There are three possibilities, described below

We know that and as a probability it can only take values between 0 and 1. In the corner case in which and player B is considered to be an expert from the beginning, the two opposing signals 'cancel' each other out and the result is that the player is indifferent between the two choices. In the usual situation in which , however, player B will always prefer to herd rather than follow his own signal. This is intuitively plausible, since the uncertainty regarding player B's performance leads him to follow the informative signal of the more reliable player. This results in proposition 2 below.

Proposition 2: Player B will always prefer to herd as long as there are no intermediate reputation effects (provided –).

In order to see this result graphically, we can plot the difference between the two utilities for different values of θ , keeping *p* constrained between – and 1. I provide below the graphs of this function for $\theta = 0.5$ and $\theta = 0.75$.

The Effect of p on The Advantage of Herding for $\theta{=}0.5$

р

As we can see from the charts, the higher the value of θ , the more the relationship between the advantage of herding¹⁴ and the value of *p* becomes exponential. In other words, if *p*=0.6 and θ =0.5, then the advantage of herding is equal to 0.102. If *p*=0.6 and θ =0.75, then the advantage of herding is equal to 0.05. While the exact values aren't of interest here, the change in magnitude shows that herding is less advantageous for each value of *p* as the likelihood of being smart ex-ante increases. At the extreme, when ______, player B is indifferent between herding and not herding.

Model with only Intermediate Reputation Concerns

The model with only intermediate reputation concerns can be thought of as a model in which the player chooses to completely discount the future and only cares about his

¹⁴ The 'advantage of herding' is defined as the difference between and

immediate gains in terms of reputation. These intermediate functions reflect the labor market's probability that player B is smart and are calculated based on the relative actions of the two players, not on their private signals, as well as on the labor market's belief in the likelihood that herding occurs. The situation analyzed here is the same conflicting context in which player A has chosen action L and player B has a signal of In terms of actions and signals, Proposition 1 makes it easy for the market to deduce player A's signal from his action, since he acts first and always chooses to follow his signal.

For player B, however, choosing action L is less informative. The market has to take into account that either he received signal and followed his own signal (which the market can take as an indication of being 'smart') or he received signal and decided to herd in order to increase his reputation (which the market would not consider a behaviour worth rewarding even if the player is "smart"). In addition, the market uses the signal of the first player to evaluate the likelihood of different states of the world occurring, taking into account that the expert himself may have been wrong due to that idiosyncratic margin of error that would affect any smart player. As a result, we have the following intermediate reputation function for action L

We can think of the numerator of this function as all the possible configurations of moves player B can make given that he is smart and the denominator represents all the possible configurations a smart central banker will do and all the configurations a dumb central banker will do. Figure 3 below summarizes the circumstances in which a 'smart' player can find himself given his selection of "Left".

Figure 3: V	Nhat are a	Smart New	bie's O	ptions?
-------------	------------	-----------	---------	---------

Expected State of the	Signal	Probability	Herding	Weight
World				
			Not Herding	
			Herding	
			Not Herding	
			Herding	

For a quick check of whether this intermediate reputation function for herding accurately reflects the logic of the model, we can evaluate what happens when takes on the values of 0 and 1. When is equal to 1, reflecting the labor markets perception that player B always herds, we have the following values for

Sincehas to be between 0 and 1, we need to check for which values thefunction fits within the bounds. The numerator of this function is always less than or equal tozero sinceThe denominator's sign depends on whether. Since

- , then , then , which means that . Since

as well, the product and thus the denominator is less than or equal to zero. This implies that the entire fraction is positive.

The circumstances under which intermediate reputation is equal to zero are those in which the informativeness of the signal for smart players is either 0 (outside of the constraints of the model) or 1 (perfectly informative signal). In addition, when θ is zero, implying that the player is known to be dumb, the reputation function is null¹⁵.

Our interest, however, is to see how the ex-post evaluation of the probability that player B is smart compares to the ex-ante probability that was simply θ . In order to analyze this, we calculate which is equal to

In this situation of anticipated herding, there will be no reputation benefits and, in fact, will be lower than θ – the a priori probability that player B is smart. The exceptions are the situations in which player B is already known to be dumb with certainty (θ is 0), for which the entire reputation function is equal to 0 as seen from the previous equation, or those in which player B is already known to be smart with certainty (θ is 1) for which there is no more

¹⁵ In fact, the null reputation function reflects the fact that the intermediate reputation doesn't change from its ex ante value (namely 0). The player was known to be 'smart' with probability 0 and remains 'smart' with probability 0.

reputation to be gained. The negative result makes intuitive sense. The market in this situation uncovers the plan of player B to pretend that he has the same signal as the expert and therefore the same good 'professional intuition' and revises its probability that B is smart downwards.

When is equal to 0, reflecting the labor market's perception that player B will never herd, we get the following values for

The denominator is positive since we know that. The numerator canbe simplified further to. We know that, and we know that for-,. This means that. We can now compare this functionto θ in order to see what effect on the prior probability of B being 'smart' this labor marketperception has

Since we know that , the numerator is a positive number, and the entire fraction is greater than or equal to zero. This implies that there is an added benefit to reputation if the market does not expect you to herd. Intuitively, this makes sense, since you are succeeding in convincing the market that your choice was based on receiving the appropriate signal and therefore you can increase the probability of being 'smart'.

For action R, we have a different intermediate reputation function. The market's evaluation of the probability that player B is herding does not factor into this calculation since herding is defined as taking the same action as player A while disregarding one's own signal.

Since player B is choosing R, his choice cannot qualify as herding. However, player B's signal is not known and the choice of action "right" can imply two possibilities. Either his private signal was in fact and he decided to follow the signal or his private signal was and he decided to ignore both his signal and the expert's. This function also takes into account the probability of each state's occurring based on the action of the expert beforehand.

The function is

As we can observe, this function is equal to . When the market always expects player B to herd, his intermediate term payoff for herding is equal to his intermediate payoff for going against the expert and both imply a decrease in the player's reputation.

We can now set the intermediate reputation functions and equal to each other and calculate the solution

Solving this equation provides the following solutions for the equality

As anticipated from the previous discussion, player B is indifferent between herding and not herding when he is either known already to be smart or dumb, or the market expects him to always herd or the signal of the expert is just as uninformative as the signal of a dumb player. Now we can provide some comparative statics to understand the effects of changes in p, θ and

The Effect of Changes in p

In order to analyze the effect of a change in p, I set the other two parameters θ and equal to – (the midpoint value) and then graph the difference between reputation from herding and reputation from not herding for – . I prefer the visualization to calculating the first order derivative since it offers a clearer understanding to the reader. A more rigorous proof of these effects can be provided on demand using Mathematica 8.0.0.

р

As anticipated, the advantage of herding increases as p becomes greater. As the market expects the expert's signal to be more informative, it is better for player B to follow this signal. The lower error rate of the expert implies that choosing action R is a stronger indicator that player B is dumb.

The Effect of Changes in

-

In order to analyze the effect of a change in θ , I set the other two parameters p and equal their midpoint values – –) and then graph the difference between reputation from herding and reputation from not herding for .



θ

The inverse U-shape of the graph suggests that the advantage of herding is maximum when player B's ability is most uncertain (in other words, when θ is very close to –). When the labor market's prior belief is skewed towards believing player B is more likely to be dumb (– , the advantage of herding will be smaller since the market can evaluate the situation with more certainty. The situation when labor market's prior belief is skewed towards believing player B is more likely to be smart (– is the same. In other words, herding tends to give the highest pay off in the intermediate run when you are truly a new player with little information regarding his/her skills.

The Effect of Changes in

In order to analyze the effect of a change in P, I set the other two parameters p and equal their midpoint values - -) and then graph the difference between reputation from herding and reputation from not herding for



The relationship between the labor market's perception of herding and the advantage of herding is straightforward. The more likely the market believes herding to be, the less gain can a player get from herding. At the opposite end, where is closest to zero, the player can trick the market into believing he is smart and increase his reputation in the intermediate run.

The model with only intermediate reputation concerns has shown that player B will find it more advantageous to herd, as long as the market does not expect herding with certainty. This corresponds to the predictions of the model with only judgement day concerns and leads us to expect that herding will occur in the general model as well.

Equilibrium with Intermediate Reputation and Judgement Day

This analysis brings together the previous two parts of the model in order to discuss the decision to herd within a professional community. The first situation I will look at focuses on a three period situation, in which the outcome would be revealed in period t=3. For player B, this

implies that there will be one period of intermediate reputation effects after making his decision and one period of discounted judgement day utility.

Bringing the two functions together results in the following total utility for choosing "Left" (the herding choice)



The same function can be produced for the choice of "Right" (the non herding choice)



Setting these functions equal to each other, we can see under which values for parameters of the model (if any) the player will be indifferent between herding and not herding. The expectation is that the player will always prefer to herd, given the results of the analysis with only judgement day and with only intermediate reputation. To simplify the calculations, we set and . This implies that and can be calculated out to This entire equation simplifies down to

The solutions for indifference between herding and not herding are

As anticipated from the previous discussion, the only time player B is indifferent between herding and not herding is when he and the market both already know that he is smart or when the expert himself has a signal that is uninformative (unrealistic expectation). The third solution for is negative, therefore does not fit the constraints of the model. Now we can look at the comparative statics to understand the effects of the parameters

The Effect of Changes in p

For the comparative statics, I set the other two parameters θ , δ and equal to – (the midpoint value) and then graph the difference between reputation from herding and reputation from not herding for – .

As anticipated, the advantage of herding increases as p becomes greater. As the market expects the expert's signal to be more informative, it is better for player B to follow this signal. The lower error rate of the expert implies that choosing action R is a stronger indicator that player B is dumb.

The Effect of Changes in θ

For the comparative statics, I set the other three parameters p, δ and equal their midpoint values - - -) and then graph the difference between reputation from herding and reputation from not herding for .

The Effect of on The Advantage of Herding



The relationship between the advantage of herding and θ is concave and decreasing in θ . The more likely it is that player B is smart, the less utility he receives from herding. This result makes intuitive sense since herding is less appealing the more you trust your own judgement. Being more likely to be 'smart' increases one's trust in one's own signal. In addition, the utility function adds the effects of on the judgement day and intermediate reputation functions. While the judgement day function decreases in a liniar fashion with , the intermediate function was an inverse U-shape. Together, these two effects provide the shape of the overall difference between the utility of herding and that of not herding.

The Effect of Changes in δ

In order to analyze the effect of a change in δ , I set the other three other parameters p, θ and equal their midpoint values - - -) and then graph the difference between reputation from herding and reputation from not herding for .

The Effect of on The Advantage of Herding



The result is an increasing linear function. The closer the discount factor is to 1 (in other words, the less player B discounts the future, the more advantageous herding becomes. Given that in both the judgement day only model and the intermediate reputation functions only model, herding was always more profitable than not herding, this result is showing that herding is even more profitable when you give the future more weight. The more eagerly the player awaits judgement day, the more he is willing to engage in herding.

The Effect of Changes in

In order to analyze the effect of a change in , I set the other three other parameters p, θ and equal their midpoint values - - -) and then graph the difference between reputation from herding and reputation from not herding for .



As increases (in other words, as the market finds it more likely that player B is herding), the advantage of herding decreases. However, the function never reaches a value of zero because of the judgement day utility component. For these values of the other three parameters, at $\$, the difference between herding and not herding is equal to 0.14 which is also the value obtained in the judgement day only model for - -.

From this analysis using the full model, we can conclude that player B will always prefer to herd as long as – We have also been able to observe the way changes in the different parameters affect the comparative advantage of herding.

IV. Conclusion

In this paper, I have attempted to construct a model through which one can analyse the phenomenon of herding in professional communities. Due to the simple set up of the model, I predict that herding will occur both in situations in which the players are only concerned with the future and in situations in which players are only concerned with their immediate reputation. Despite its simplicity, taking into account the significant advantage provided by herding can serve as a cautionary tale for those receiving advice from members of a professional community. Policy-makers in particular, as consumers of professional advice, are in a position to remedy this situation in order to obtain maximum social utility, in this particular example when consulting central bankers.

The goal is to be able to vary the parameters of the model in order to recreate the conditions faced by professionals in different environments from that of central banking used in this example. There are a number of changes that can be made to the model in order to obtain situations in which not herding is more profitable. A simple solution is to force new comers to take decisions before reputable experts. However, in many circumstances, this would not be feasible. Another way of allowing for more variation of outcomes is to allow for the possibility that new comers may have better predictive power from their signals than established experts. In many fields where education has kept the pace with technological innovation, it is not unlikely that fresh graduates may have a higher potential than their peers. Furthermore, allowing players to know more about the quality of their intuition beforehand can offer them an incentive to signal to the market their ability by going against the herd.

Regardless of the extension of choice, the model provides an intuitive foundation on which the analysis of communities of experts can be based. Keeping in mind that herding carries a social loss by ignoring potentially informative signals, one should always be on the lookout for new ways to incentivise risk taking by going against the herd.

Works Cited

- Avery, C. N. and J. A. Chevalier (1999). "Herding over the career." <u>Economics Letters</u> 63(3): 327-333.
- Banerjee, A. V. (1992). "A Simple Model of Herd Behavior." <u>The Quarterly Journal of Economics</u> 107(3): 797-817.
- Bikhchandani, S., D. Hirshleifer, et al. (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." <u>The Journal of Political Economy</u> 100(5): 992-1026.
- Effinger, M. R. and M. K. Polborn (2001). "Herding and anti-herding: A model of reputational differentiation." <u>European Economic Review</u> 45(3): 385-403.
- Keynes, J. M. (1936). <u>The general theory of employment, interest and money</u>, London, Macmillan and Co., ltd., 1947.
- Levy, G. (2004). "Anti-herding and strategic consultation." <u>European Economic Review</u> 48(3): 503-525.
- Ottaviani, M. and P. Sørensen (2000). "Herd Behavior and Investment: Comment." <u>The</u> <u>American Economic Review</u> 90(3): 695-704.
- Scharfstein, D. S. and J. C. Stein (1990). "Herd Behavior and Investment." <u>The American</u> <u>Economic Review</u> 80(3): 465-479.
- Trueman, B. (1994). "Analyst forecasts and herding behavior." <u>Review of Financial</u> <u>Studies</u> 7(1): 97-124.

• Zwiebel, J. (1995). "Corporate Conservatism and Relative Compensation." <u>The Journal</u> of Political Economy 103(1): 1-25.