

Social learning and financial crises

Marco Cipriani and Antonio Guarino, NYU

Introduction

The 1990s witnessed a series of major international financial crises, for example in Mexico in 1995, Southeast Asia in 1997-8, Russia in 1998 and Brazil in 1998-9. These episodes have revived interest among economists in the study of financial system fragility. Theoretical research has analysed various problems, such as bank runs, currency attacks and international contagion. Although many approaches have been taken, two main perspectives have emerged. One part of the literature has emphasised the relation between financial crises and weak fundamentals of the economy. Another has stressed that crises may just be due to random events and self-fulfilling prophecies, with variables unrelated to the real economy acting as “sunspots”. Early macroeconomic models of currency crises such as Krugman (1979) (for a survey, see Flood and Marion (1999)) are an example of the first perspective. Microeconomic models of bank runs, such as Diamond and Dybvig (1985), are an example of the second. Rather than alternative explanations, these two views are now considered complementary: financial crises do not occur only in the presence of weak fundamentals, but weak fundamentals can trigger bank run psychology and this in turn can have disproportionately bad effects on the real economy.

To see how difficult it is to reconcile some crisis episodes with a fundamental analysis, let us consider Figure 1, taken from Kaminsky (1999). The chart refers to the case of Malaysia. The solid line is the probability of a crisis estimated using Malaysia’s macroeconomic fundamentals. The figure shows that fundamental variables may not be sufficient to forecast the occurrence of a crisis. For instance, the index failed to forecast the crises of the second half of the 1970s and of the second half of the 1990s. In contrast, it forecast a crisis in the mid-1980s, which failed to materialise. This finding is common in much of the empirical work on financial crises. Fundamentals do help to predict when a crisis will occur; nonetheless, crises may occur when the fundamentals seem sound or may not occur when fundamentals are weak.

A possible explanation of why sound fundamentals may not be reflected in asset prices is that information about these fundamentals is spread among economic agents (ie investors) and prices may fail to aggregate it. In particular, this would happen if investors, instead of acting according to their own private information, simply decided to follow the actions of previous traders, ie if they herded. Herd behaviour may, therefore, be a reason why we can observe a misalignment between prices and asset values.

Herd behaviour and information revelation

Several recent models of social learning have shown that herding is not necessarily an irrational phenomenon.¹ The argument was originally made in two seminal papers by Banerjee (1992) and Bikhchandani et al (1992).² These papers show that, if people act in sequence and observe the action of their predecessors, the information contained in the history of actions may overwhelm private information. When this happens, agents will disregard their own private information and follow the actions of their predecessors, thus joining a herd.

¹ We limit our analysis to information-based herding in financial markets and do not discuss herd behaviour due to reputation or compensation schemes. A comprehensive survey of herd behaviour in financial markets is offered by Bikhchandani and Sharma (2001).

² After these papers, much effort has been dedicated to the topic. See, among others, Chamley and Gale (1994), Chari and Kehoe (2000) and Smith and Sørensen (2000). For a critical assessment, see Gale (1996).

The mechanism can be illustrated with a simple example. Let us consider an economy in which agents can trade a good (ie a financial asset) that can take two values, \$0 or \$100, with equal probability. In this economy, agents do not trade among themselves, but with a market-maker, who sets the price at which traders can buy or sell the asset (by going short). Let us assume that the market-maker sets a price equal to \$50, the expected value of the asset. This price is kept fixed, ie the market-maker does not change it after observing a buy or a sell. Each agent, before making his trading decision, receives some private information (a binary signal) on the value of the asset. This signal has a 70% chance of being correct. Suppose that the value of the asset is \$100, but the first two agents arriving in the market receive the wrong (ie low) signal and therefore sell the good. What will the third agent do? Even if his signal is high, he realises that the two previous agents (who sold) had low signals. The negative information contained in the first two sell orders overwhelms his private information. Therefore, he will also sell and will not reveal his (positive) information on the asset value. All the following agents will be in the same position as the third, since they realise that the third agent's action did not depend on his private information. Therefore they will all join the sell herd. Although the value of the asset was \$100, everyone will sell and the true state of the world will never be revealed (as it would be, by the law of large numbers, if all agents traded according to their own private information). The actions of the first two agents have a disproportionate and pathological effect on the history of trades.

One of the characteristics of the previous example is that the price does not adjust to the order flow. Indeed, we have assumed that even after a series of buy orders the price is fixed at the level of \$50. This is a perfectly reasonable assumption in many economic contexts. For instance, Bikhchandani et al (1992) refer to the choice of adoption of a new technology whose cost is fixed.

In financial markets, however, prices are certainly not fixed. Avery and Zemsky (1998) have shown that, in this case, the argument of Banerjee (1992) and Bikhchandani et al (1992) no longer holds.³ The presence of a flexible price induces people to follow their own private information since the price adjusts in order to factor in the information contained in the past history of trades.⁴

Let us repeat the example in the previous paragraph. Let us assume, however, that the price is not fixed at \$50, but is set equal to the expected value of the asset given the past history of trades. After the first two traders sell, the market-maker will lower the price from \$50 to \$15.50⁵ to take into account that the first two sells came from agents with a low signal. The third agent knows that the two previous traders had a negative signal. If his signal is high, his expected value of the asset will be \$30. Given that he faces a price of \$15.50, he will buy, ie he will follow his own private information. By the same argument all traders will always follow their private information. Since the signal that they receive is correct 70% of the time, over time the price will converge to the fundamental value of the asset, thus aggregating the private information dispersed among traders. Therefore, when prices are set efficiently, agents will follow their own private information and the price will aggregate the information spread among traders. Consequently, we should not observe misalignments of the price with respect to the fundamentals.

³ Avery and Zemsky base their analysis on the Glosten and Milgrom (1985) model of a specialist market.

⁴ Avery and Zemsky argue that herd behaviour arises in their model when there are multiple dimensions of uncertainty, ie when agents are uncertain not only on the asset realisation, but also on whether an informational event has occurred. Their definition of herd behaviour, however, is different from the one that is standard in the literature and refers more to "swings" of the traders' beliefs. They say that there is herd behaviour when an agent who is originally more pessimistic (optimistic) than the market on the asset value becomes more optimistic (pessimistic) after seeing a sequence of buy (sell) orders. Whereas multiple dimensions of uncertainty make this type of "swing" possible, they do not make informational herding (which the authors call an "informational cascade") possible (see Proposition 2 on page 728 of Avery and Zemsky's paper).

⁵ The value of \$15.50 is obtained by using Bayes's formula.

Explaining rational herds in financial markets

The point made in the previous example is a powerful one. Flexible prices seem to rule out situations in which rational traders choose to disregard their own private information. Given this result, can we still relate the observed price misalignments to rational herding?

In the example, for prices to be able to destroy herds, the traders and the market-maker must value the asset in the same way. In this case, traders find it convenient to use their informational advantage (their private signal) and never herd. The expected utility gained from buying or selling a financial asset, however, may be different across different classes of traders, or between the traders and the market-maker. In other words, a wedge can exist between the way the traders and the market-maker value an asset after observing the same history of trades. When this is the case, traders may decide to disregard their own private information and herd. In the remainder of the section we will briefly describe two papers of ours, in which we analyse an economy where the expectations of traders and market-makers diverge and, as a result, herds arise.

In Cipriani and Guarino (2001b) we show that a divergence between the trader's and the market-maker's valuations can arise when there is uncertainty on the degree of informativeness in the economy (for example, on the proportion of traders who act for informational reasons). Because of these different valuations, even a trader with a negative (positive) signal may decide to buy (sell) because he believes that the asset is undervalued (overvalued). Therefore, there may be situations in which all traders buy or sell independently of the information they have and the price is unable to aggregate the information dispersed among traders. Consequently, the price remains far away from its fundamental value for a period of time longer than if agents always followed their own private information. Eventually, however, the uncertainty on the degree of information in the economy will be resolved (ie it will be learned by the traders and the market-maker) and people will resume trading according to their private information. Therefore, the mechanism outlined in this paper can account for misalignments of the price with respect to the fundamentals, but these misalignments are not long-lasting. There are only a limited number of periods in which people disregard their own private information.

In Cipriani and Guarino (2001a) we consider another source of asymmetry between the traders' and the market-maker's valuations. Different valuations can arise because traders themselves are heterogeneous, ie they may have different degrees of risk aversion or different propensities to save or consume. Different valuations can also be the result of different hedging needs that make some traders more willing to buy an asset, and others more willing to sell it. Differences in the preferences of economic agents are a fundamental feature of markets, which is usually overlooked in the financial market microstructure literature. In many market microstructure papers, agents are assumed to trade only for information reasons (ie because they have a signal about the value of an asset). What characterises markets, however, is that agents are heterogeneous and there are gains from trade. Trade is not driven simply by asymmetric information.

When preferences are homogenous across agents, the price that the market-maker sets is equal to the expected utility that all agents enjoy from the asset. In contrast, when preferences differ across traders and between the trader and the market-maker, the price cannot be equal to the expected utility of each trader. At a given price, some agents will find the asset overvalued, and some will find it undervalued. This wedge between the market-maker's and traders' valuations implies that, for instance, even traders with good information on the value of the asset may decide to sell because the price is simply too high according to the utility that they expect from the asset.

In other words, in a market where traders' preferences are not identical, agents trade not only because they have different information, but also because the asset gives them different utilities. It may happen that this second reason becomes more important than the informational one and traders simply decide to disregard their own private information. In this case, a trade does not reveal anything about the traders' private information, which is therefore not aggregated by the market price. The price remains far away from the fundamental value forever.

In Figure 2, we consider an asset that can take two values, 1 and 2, with equal probability. We show a simulated price path for this asset. Although the realised value of the fundamental is 2, the price converges to a value close to 1. The prevalence of sell orders at the beginning of trading induces all following informed traders to neglect their private information. Given that the market-maker realises that trades are independent of private information, he does not revise the price, which remains stuck at

a low value. This price behaviour may explain why, as we discussed in the introduction, models based on the fundamentals may fail to predict the occurrence of a financial crisis.⁶

Financial contagion

Financial crises often spill over from one country to the other, even when these countries are not closely linked. Consequently, asset prices show a correlation in excess of that between the fundamentals. This phenomenon, labelled as financial contagion, is of great relevance, as an economy with sound fundamentals might be affected by a financial crisis just because another economy has been hit. In the 1990s, episodes of contagion were the “tequila” effects of the Mexican currency crisis of 1994-5, the “yellow fever” during the Asian crisis of 1997-8, the asset market crises following the Russian default in 1998 and the Brazilian devaluation in 1999.

We believe that herding can explain why we observe co-movements in asset prices that cannot be accounted for by the fundamentals. In Cipriani and Guarino (2001a) we show that sell orders in one market can affect the price path of another and make its price settle to a lower value. Of course, informational spillovers are to be expected in asset markets, as long as there is some degree of correlation between the fundamentals. We show, however, that these informational spillovers can have pathological consequences. Sell orders in country A not only depress the price of financial assets in country B, where fundamentals are good, but can also cause herd behaviour to arise in this country. Given that herding prevents the revelation of private information, asset prices in country B can remain below the fundamentals even in the long run.

Some evidence

During the 1990s, parallel to the development of the theoretical analysis of herding, many scholars made a significant effort to capture the presence of herd behaviour in empirical data. Starting with the seminal work of Lakonishok et al (1992a), several studies have tried to understand whether different categories of fund managers cluster their decisions (for a survey, see Bikhchandani and Sharma (2001)). These empirical investigations, however, do not estimate any theoretical model of herding, but test whether fund managers cluster their decisions significantly more than one would expect if they acted independently. The reason why there have been no attempts to test a model of informational herding is quite clear. There are no data on the information available to individual traders, and, therefore, it is difficult to gauge whether they disregard their private information when they trade.

An alternative route to test herding models is to gather experimental data. Experimental analysis allows us to test herd behaviour directly because we can control the information available to economic agents. For this purpose we have constructed a laboratory financial market to test whether people tend to imitate their predecessors (Cipriani and Guarino (2002)). In our study, experimental subjects traded in sequence an asset that could take values of \$0 or \$100. In a situation where all agents traded only for informational reasons and the price adjusted in a Bayesian fashion to the order flow, most experimental subjects did follow their own private information. This seems to show that prices destroyed the incentive of agents to herd. As a result, the price converged to its fundamental value.

When, however, there was a wedge between the expectations of the traders and of the market-maker (for instance, because of non-informational reasons to trade or trade costs), the behaviour of the experimental subjects changed substantially. After the first agents had traded, the following ones stopped trading according to their own private information. Consequently, the price did not always converge to the correct fundamental value. Figure 3 shows the histogram of the last prices (ie the

⁶ Note, however, that the price will converge close to 1 more frequently when the fundamental is 1 than when the fundamental is 2. Therefore, when fundamentals are bad, a crisis is more likely to happen. The analysis does not show that the fundamentals are not useful in predicting a financial crisis; rather that some financial crises cannot be predicted looking at the fundamentals.

price after all agents had traded) for the runs of the experiment in which the realised value was 100. The histogram shows that in 10% of cases the price settled on a value far below the fundamental one.

Therefore both the chart in Figure 2 and the histogram in Figure 3 show a behaviour qualitatively similar to the one experienced by Malaysia in the 1970s and 1990s (see Figure 1). That is, in both the theoretical simulation and the experiment we have a financial crisis (ie a large number of periods in which the price is low) that cannot be justified by the fundamentals, but is only due to the inability of the price mechanism to aggregate private information.

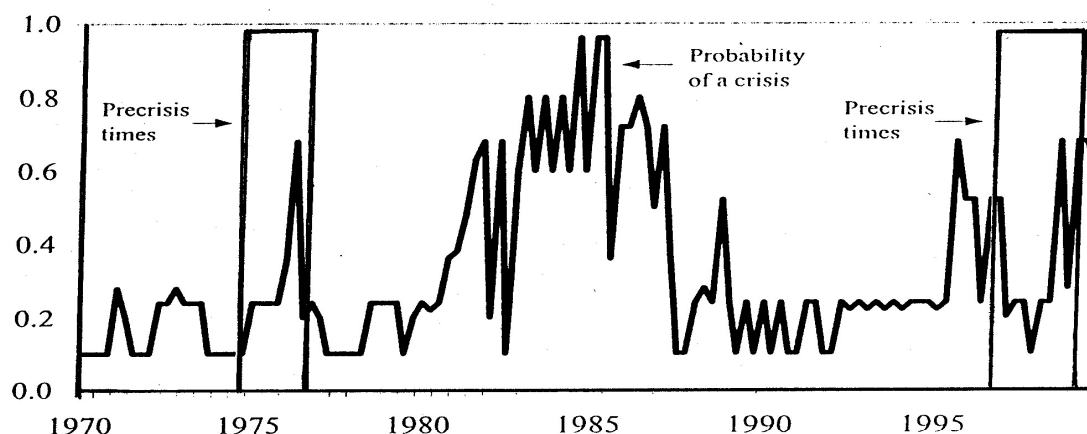
It is the mechanism itself by which prices are formed in financial markets that can explain why we sometimes observe a financial crisis when the fundamentals are good. Even when prices are flexible, rational traders may find it convenient to disregard their own private information. When this happens, the market price may fail to aggregate the information dispersed among traders and long-lasting misalignments may occur.

Conclusions

Our theoretical analysis shows that, in a financial market, the mechanism of price formation may lead traders to disregard their own private information and herd. When this happens, the price does not aggregate traders' information and does not reflect the fundamental value of the asset. Consequently, a financial crisis may occur even when the fundamentals of the economy are sound. Experimental data show that this phenomenon is observed in a laboratory financial market, where experimental subjects choose to disregard their own private information and herd.

Figure 1


Probability of a currency crisis in Malaysia according to fundamentals.



Note: The solid line represents the estimated probability of a financial crisis in Malaysia computed using fundamental variables. The rectangles represent 24-month windows before the occurrence of crises.

Source: Kaminsky (1999).


Figure 2
A simulated path for the asset price



Note: The figure shows the simulated price of a security with a realised value of 2. The price starts from the unconditional expected value of 1.5 and then, after a predominance of sell orders at the beginning of trading, it decreases to a value close to 1. When traders start herding, the price does not change and fails to converge to the fundamental value.

Source: Cipriani and Guarino (2001a).

Figure 3
The histogram of prices in the experiment, $V = 100$



Note: The figure shows the histogram of last prices in the experimental study for all runs in which the asset value was 100. The last price is the price recorded after all experimental subjects had traded. In 10% of cases, the price converged to a level far from the fundamental value.

Source: Cipriani and Guarino (2002).

References

- Avery, C and P Zemsky (1998): "Multidimensional uncertainty and herd behavior in financial markets", *American Economic Review* 88, 724-48.
- Banerjee, A (1992): "A simple model of herd behavior", *Quarterly Journal of Economics* 107, 787-818.
- 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, 992-1027.
- Bikhchandani, S and S Sharma (2001): "Herd behavior in financial markets", *IMF Staff Paper*, vol 47, no 3.
- Chamley, C and D Gale (1994): "Information revelation and strategic delay in a model of investment", *Econometrica* 62, 1065-86.
- Chari, V and P Kehoe (2000): "Financial crises as herds", mimeo, University of Minnesota.
- Cipriani, M and A Guarino (2001a): "Herd behavior and contagion in financial markets", mimeo, NYU.
- (2001b): "Herd behavior and informational uncertainty: a note", mimeo, NYU.
- (2002): "Herding and price convergence in a laboratory financial market", mimeo, NYU.
- Diamond, D and P Dybvig (1983): "Bank runs, deposit insurance, and liquidity", *Journal of Political Economy* 91, 401-19.
- Flood, R and N Marion (1999): "Perspectives on recent currency crises literature", *International Journal of Finance and Economics* 4, 1-26.
- Gale, D (1996): "What have we learned from social learning", *European Economic Review* 40, 617-28.
- Glosten, L and P Milgrom (1985): "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders", *Journal of Financial Economics* 14, 71-100.
- Kaminsky, G L (1999): "Currency and banking crises: the early warnings of distress", mimeo.
- Krugman, P (1979): "A model of balance of payments crises", *Journal of Money Credit and Banking* 11, 311-25.
- Lakonishok, J, A Shleifer and R W Vishny (1992a): "The impact of institutional trading on stock prices", *Journal of Financial Economics* 32(1), 23-43.
- Milgrom, P and N Stokey (1982): "Information, trade and common knowledge", *Journal of Economic Theory* 26, 17-27.
- Smith, L and P Sørensen (2000): "Pathological outcomes of observational learning", *Econometrica* 68, 371-98.
- Welch, I (1992): "Sequential sales, learning, and cascades", *Journal of Finance* 47, 695-732.

