Information Cascades and Rational Conformity

Intermediate article

Lisa R Anderson, College of William and Mary, Williamsburg, Virginia, USA
Charles A Holt, University of Virginia, Charlottesville, Virginia, USA

CONTENTS

Introduction
Conformity incentives
Rational information cascades
Applications to markets and other social institutions
Conclusion

An information cascade is a pattern of matching decisions. A cascade can occur when people observe and follow 'the crowd', which can be rational if the information revealed in others' earlier decisions outweighs one's own private information.

INTRODUCTION

When individuals obtain private information and make publicly observed decisions in a sequence, the first decisions tend to act as 'signals'. If early decisions show a clear pattern, the information inferred from them may outweigh any one person's private information. This inference can cause people to 'follow the crowd', even when the group consensus conflicts with their own private information. This type of sequential conformity is termed an 'informational cascade' in a seminal paper by Bikhchandani *et al*. (1992).

For example, consider an employer who interviews a job applicant and forms a good impression. The employer, however, does not offer the job after discovering that the worker has been turned down previously by other employers. Even though this decision is in conflict with the employer's own information, it may be rational if the employer concludes that the other interviews went badly, and that the aggregate information implied by past interviews more than offsets the employer's own positive evaluation. Even a qualified applicant may make a bad impression on any given day, so may have difficulty finding a job after several unsuccessful attempts.

Bannerjee (1992) uses the term 'herd behavior' to describe similar patterns of conformity that arise in models where individuals must decide which type of financial asset to purchase. Indeed, much of the interest in cascade behavior arises from attempts to explain temporary patterns in investment behavior.

CONFORMITY INCENTIVES

There may, of course, be non-informational factors that produce conformity in social interactions. Sometimes people prefer to behave like the others in a group. Such behavior has even been recommended (Post, 1927, chap. XXXIII): 'to do *exactly as your neighbors* do is the only sensible rule'.

There may be social stigmas and punishments associated with nonconformity. For example, an economic forecaster may prefer the chance of being wrong with everybody else to the risk of providing a deviant forecast that turns out to be the only incorrect guess. In the words of John Maynard Keynes (1936, p. 158): 'worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally'.

Some research suggests that people prefer to maintain the status quo (Samuelson and Zeckhauser, 1988). For example, subjects in experiments were provided with a scenario in which they inherit a portfolio of cash and securities and are asked whether to leave the portfolio intact or to change it by investing in other securities. There was a strong tendency for individuals to retain portfolio *A* in preference to *B* when it was listed as the current portfolio, and to retain portfolio *B* in preference to *A* when it was listed as the current portfolio. A similar tendency was observed in response to other matched pairs of questions that alternated descriptions of the *status quo* choice.

Such a 'status quo bias' may explain herding behavior in sequential decision situations. But these decision patterns do not allow us to distinguish between behavior based on a preference for conformity and behavior that is motivated by information inferred from prior decisions. Thus, someone inheriting a portfolio from a rich uncle might conclude that the uncle's wealth was due to wise portfolio choices that should be imitated. Laboratory experiments can be used to set up and control information flows in order to distinguish among alternative explanations of herding behavior.

RATIONAL INFORMATION CASCADES

We will use a numerical example to illustrate the concept of a rational information cascade (Anderson and Holt, 1997). In this example, there are two equally likely events, A and B, which might represent whether or not a particular patent is marketable. Decision-makers obtain private signals, *a* and *b*, which are correlated with the events. In particular, $P(a|\mathrm{A}) = P(b|\mathrm{B}) = \frac{2}{3}$, so the error rate is $\frac{1}{3}$ for each signal. (For example, the signal might be the result of a consultation with experts.) This model assumes that each person's private signal is correlated with the event but is independent of the other signals. After observing their signals, individuals are approached one by one in a sequence and are asked to make a prediction about which event has occurred. People find out the prior predictions, if any, made by others, but they cannot observe others' private signals. Thus the prediction made by the first person is based only on that person's signal, and hence will reveal that signal, since the signal is more likely to be correct than not.

Suppose the first person sees a *b* signal and publicly predicts event B. If the second person in the sequence sees a *b* signal, it is rational for this person to also predict B. If the second person sees an *a* signal, the observed and inferred signals essentially cancel each other, and each state is equally likely. We observe from laboratory experiments that individuals almost always use their own information in such cases, and therefore, the second decision will reveal the person's private signal, whether or not it conforms to the first prediction. When the first two individuals in the sequence observe the same signal, their decisions will also match. In this case, the information inferred from the matching decisions is greater than the information inferred from any one private signal. In particular, if the first two people choose B, then the third person should also choose B, even if that person's private signal is *a*. Information cascades form in this manner, and the effect is that all subsequent decision-makers will follow a pattern established by the first ones in the sequence.

This example was used by Anderson and Holt (1997) in a laboratory experiment in which subjects were paid a cash reward for each correct prediction. The events were referred to as 'urn A' and 'urn B'. Each urn was a cup with three colored balls, which we will refer to as *a* or *b* signals. There were two *a* balls and one *b* ball in urn A, and there were two *b* balls and one *a* ball in urn B.

A random device was used to select the urn, with each event being equally likely, and therefore, each of the six balls is equally likely to be drawn. Suppose that the draw is *b*. Since two of the three *b* balls are in urn B, the posterior probability of urn B given a draw of *b* is $\frac{2}{3}$. (This is an example of the application of Bayes' rule; or more precisely, the ball-counting heuristic used here corresponds to the conditional-probability calculations that are referred to as Bayes' rule. See Holt and Anderson (1996) for a discussion of how the simple counting heuristic can be used to make Bayesian calculations in more complicated settings, and an intuitive explanation of the mathematical expression of Bayes' rule.) Of course, it is intuitively obvious that the probability of urn B is greater than $\frac{1}{2}$ when the signal is *b*. Holt and Anderson (1996) show that the probability of urn B is still greater than $\frac{1}{2}$ when first two predictions are B and the third person's signal is *a*. Thus, a cascade can begin with two matching decisions, and all subsequent decision-makers should follow the pattern established in this manner.

Information cascades may not form immediately if there is not a bias in early predictions. Suppose, for example, that the first two predictions are A and B, so the third person would consider each urn to be equally likely, prior to seeing a private signal. If the third and fourth decision-makers both predict B, then this bias in favor of B would cause the fifth person to predict B, regardless of that person's private signal.

These calculations are based on a model of perfect rationality. The purpose of an experiment is to determine how people actually behave in such situations. Anderson and Holt (1997) used the 'ball and urn' scenario described above in order to remove any preference for conformity that is not based on informational considerations. Subjects were placed in small cubicles and were shown a single ball drawn from the relevant urn, but they could not see which urn was being used. Subjects were selected in a random order to make their predictions, which were announced by a neutral assistant who did not know the signals or which urn was being used. (Allowing subjects to announce their own predictions could have given them the chance to convey additional uncontrolled information, by tone of voice, for example.) After all predictions had been announced, a non-decision-making subject serving as a 'monitor' announced which urn had actually been used. Those with correct predictions were paid $2, and others earned nothing for that trial. There were 15 trials and six decision-making subjects in each session. Altogether, there were 12 sessions in the experiment, each of which was conducted on a different day.

The sequences of draws made cascades possible in more than half of the trials, and cascades actually formed in about 70% of the trials in which they were possible. A particularly interesting trial is shown in Table 1. In this case, all six individuals earned $2, since urn B was actually used for the draws.

Generally, prior information is informative, and cascade behavior tends to increase the accuracy of predictions, and hence to increase earnings. It is possible, however, for initial predictions to be incorrect, which may create an incorrect cascade. This occurred in another trial, shown in Table 2. In this case the first person's prediction revealed their signal, but the second person made a serious error in predicting B after seeing an *a* signal and a prior A prediction. The third person predicted B, which was also an error, since previous predictions were balanced and their own signal suggested urn A. All of the remaining individuals predicted B, which turned out to be incorrect, since urn A was used.

The most common type of error occurred when a person saw a signal that was inconsistent with the implication of prior predictions. In this case, the optimal Bayesian prediction is to follow the established pattern, but subjects deviated in about a quarter of the cases in which their own private information was inconsistent with this pattern.

In summary, the general tendency was for subjects to use the information implied by previous decisions correctly, which produced rational information cascades. There were, however, deviations that could either break a cascade or result in an incorrect cascade. Incorrect cascades also occasionally resulted from 'unlucky' incorrect decisions observed by early decision-makers.

This pattern of results was replicated by Hung and Plott (2001), who added some interesting treatment variations. In one of their experiments, the incentive structure was altered so that subjects received a positive pay-off only if the majority of the group made the correct prediction. (This is somewhat like a jury whose decision is determined by a majority vote.) The effect was to reduce conformity for early decisions because individuals have an incentive to signal their information so that others can make better decisions. A second treatment rewarded conformity directly: subjects received a positive pay-off only if their decision matched that of the majority.

Table 1. A trial from Anderson and Holt's (1997) 'ball and urn' experiment. Three of the last four decision-makers follow the pattern established by the first two despite contradictory evidence

Subject	58	57	59	55	56	60
Observes	*b*	*b*	*a*	*b*	*a*	*a*
Predicts	B	B	B	B	B	B

Table 2. A trial from Anderson and Holt's (1997) 'ball and urn' experiment. The second and third subjects both made errors, which led to an incorrect cascade (urn A was used)

Subject	57	58	59	55	60	56
Observes	*a*	*a*	*a*	*a*	*a*	*a*
Predicts	A	B	B	B	B	B

APPLICATIONS TO MARKETS AND OTHER SOCIAL INSTITUTIONS

Strong movements in stock prices are sometimes attributed to herd-like behavior. Keynes (1936) noted the similarity between investment decisions and a guessing game in which participants have to predict which contestant in a beauty contest will get the most votes. In this game, each person has to think about whom the others think is attractive, and also about whom the others think others will find attractive, and so on. Similarly, when investing in stocks, one would like to guess which enterprises will become popular with other investors, since a strong demand will raise the prices of those stocks. A herd-like response may move asset prices out of line with market fundamentals, and thus set the stage for an equally strong 'stampede' in the other direction. Such behavior could be due to seemingly irrational 'animal spirits', to use Keynes's colorful term, or it could be due to a rational tendency to follow others' decisions when they are based on independent sources of information. Christie and Huang (1995), for example, argue that it can be rational to follow others' decisions during surges or declines in stock prices; i.e., to rely on inferences derived from information that is aggregated by market prices.

Other applications are suggested by the majority-voting treatment of Hung and Plott (2001). In a trial, for example, jurors may form independent judgments about the guilt or innocence of a defendant, but such judgments are often changed in the process of voting and deliberation. In this manner, herd behavior can create the consensus needed to avoid a 'hung jury'. Some experiments that simulate sequential jury voting have been conducted, and strong patterns of cascade-like conformity are observed in many cases.

Decisions may occur in sequence in these applications (e.g., as stock purchases appear on a ticker tape), but the order of decisions is not exogenously specified as it was in the experiments discussed above. The order of voting is not exogenous in jury voting unless the foreman chooses to take votes by going around the table. Similarly, stock purchases or responses to initial public offerings are not subject to order requirements. Plott *et al.* (1997) report some pari-mutuel betting experiments with an endogenously determined order of play. The incentive structure is like that of a horse race in which the purse is divided among those who bet on the winning horse, in proportion to their bets. The experiment was presented as a choice between six assets, with only one of them offering positive earnings, depending on the realized state of nature. Investors had private, noisy information about the assets, and they could observe others' bets as they were made. A considerable amount of information aggregation was observed in these experiments, with the asset prices accurately indicating the correct state in most cases. In some cases, however, heavy purchases of an asset triggered more purchases, even though the asset being purchased turned out not to be the one that paid off. This corresponds to an incorrect cascade. The application of information cascade theory to asset markets in richer and more realistic settings is a prime area for future research.

CONCLUSION

Theoretical models of 'herding' pertain to situations where individuals observe private signals that are correlated with some unknown event. Predictions about the event are made in sequence, with later decision-makers being able to base their predictions on their own signal and on information inferred from prior decisions. The first few decisions tend to reveal the private signals, which may establish a pattern of matching predictions that others follow, even if the conforming predictions are different from the prediction that would be best given only the person's own private signal. This type of 'information cascade' can produce conformity that is rational, because the information content in prior decisions may outweigh that in one's own private signal. There is some laboratory evidence to support such theories of rational cascades, although individuals do make mistakes, and behavior is sometimes influenced by biases and heuristics that may lead to non-Bayesian decisions.

References

Anderson LR and Holt CA (1997) Information cascades in the laboratory. *American Economic Review* **87**: 847–862.

Bannerjee AV (1992) A simple model of herd behavior. *Quarterly Journal of Economics* **107**: 797–817.

Bikhchandani S, Hirshleifer D and Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* **100**: 992–1026.

Christie WG and Huang RD (1995) Following the Pied Piper: do individual returns herd around the market? *Financial Analysts Journal* **51**: 31–37.

Holt CA and Anderson LR (1996) Classroom games: understanding bayes' rule. *Journal of Economic Perspectives* **10**: 179–187.

Hung AA and Plott CR (2001) Information cascades: replication and an extension to majority rule and conformity-rewarding institutions. *American Economic Review* **91**: 1508–1520.

Keynes JM (1936) *The General Theory of Employment, Interest, and Money*. London, UK: Macmillan.

Plott CR, Wit J and Yang WC (1997) Paramutuel betting markets as information aggregation devices: experimental results. Working Paper, California Institute of Technology.

Post E (1927) *Etiquette in Society, in Business, in Politics, and at Home*. New York, NY: Funk and Wagnalls.

Samuelson W and Zeckhauser R (1988) Status quo bias in decision making. *Journal of Risk and Uncertainty* **1**: 7–59.

Further Reading

Anderson LR (2001) Payoff effects in information cascade experiments. *Economic Inquiry* **39**: 609–615.

Anderson LR and Holt CA (1996) Classroom games: information cascades. *Journal of Economic Perspectives* **10**: 187–193.

Asch SE (1952) *Social Psychology*. New York, NY: Prentice-Hall.

Asch SE (1956) Studies of independence and conformity: a minority of one against a unanimous majority. *Psychological Monographs* **70**(9).

Camerer C (1995) Individual decision making. In: Kagel JH and Roth AE (eds) *The Handbook of Experimental Economics*, pp. 587–703. Princeton, NJ: Princeton University Press.

Davis DD and Holt CA (1993) *Experimental Economics*. Princeton, NJ: Princeton University Press.

Devenow A and Welch I (1996) Rational herding in financial economics. *European Economic Review* **40**: 603–615.

Kahneman D and Tversky A (1973) On the psychology of prediction. *Psychological Review* **80**: 237–251.

Pound J and Shiller RJ (1989) Survey evidence on diffusion of interest and information among investors. *Journal of Economic Behavior and Organization* **12**: 47–66.

Thaler RH (ed.) (1993) *Advances in Behavioural Finance*. New York, NY: Russell Sage Foundation.

Welch I (1992) Sequential sales, learning, and cascades. *Journal of Finance* **47**: 695–732.

Information Processing

Intermediate article

James T Townsend, Indiana University, Bloomington, Indiana, USA
Kan Torii, Indiana University, Bloomington, Indiana, USA

CONTENTS

Definition and history
Serial and parallel processing, flow diagrams, and continuous-flow systems
Independence, capacity, and stopping rules
Mathematical models of human information processing and cognition
Codes in information processing
Systems-factorial methodologies for determining architecture
Conclusion

The information-processing approach in cognitive science assumes mental architectures with various components that interact in order to manipulate incoming information.

DEFINITION AND HISTORY

Any field of study, in this case cognitive science, always has as a backdrop a kind of super-theory or metatheory, which may itself be rather vague but which provides a philosophy or way of thinking about problems and helps to motivate the theoretical and experimental investigations that are carried out. The information-processing approach is a good candidate for the central metatheory of cognitive science. Cognitive science began to emerge, slowly at first, in the 1950s, and developed rapidly in the 1960s and 1970s, becoming the dominant paradigm in psychology and attracting researchers in linguistics, computer science and philosophy.

A popular introductory text on cognitive science (Ashcraft, 2002) defines the information-processing approach as 'the coordinated operation of active mental processes within a multicomponent memory system'. The essential element of this definition is the specification of 'active mental processes' indicating dynamic ongoing activity. The term 'processes' emphasizes this dynamic aspect, and its early and continued use led to the term 'process model', again suggesting a model or theory that tells what happens cognitively across time. A 'model' is a theory constructed to explain and make predictions about a relatively restricted set of phenomena, while a 'theory' may aim to cover a much wider range. However, the terms 'model' and 'theory' are often used interchangeably. Much of psychology has traditionally been oriented more towards static descriptions (e.g., most personality trait models are static in nature), so the idea of being composed of active processes was a major innovation in the study of human thought and behavior. The words 'coordinated operation' also suggest that the model builder is paying close attention to how the operations of the various components involved in a mental task might start, continue, and stop relative to each other. A notable omission in the above definition of information processing is the arrangement, or architecture, of the processes. Architecture has been the subject of many studies.

In the 1950s, a number of well-known experimental (and early cognitive) psychologists employed information theory (e.g., Shannon and Weaver, 1949) as a theoretical and methodological tool. Claude Shannon, an engineer, invented information theory to measure how much information is