



Learning is a social activity

LSE Research Online URL for this paper: <http://eprints.lse.ac.uk/103155/>

Version: Accepted Version

Article:

Goodhart, C. A. E. (2020) Learning is a social activity. Review of Behavioural Finance. ISSN 1940-5979

Reuse

Items deposited in LSE Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the LSE Research Online record for the item.

Learning is a Social Activity

By C.A.E. Goodhart¹
Financial Markets Group, London School of Economics

Time is the real constraint

In my earlier textbook, *Money, Information and Uncertainty* (1975), I began with the following paragraph,

“Time is the ultimate constraint on mortals. Lack of time contributes also to limited information. Shortage of time and information make it harder for agents to achieve Pareto equilibrium matchings of wants with endowments. This had led, as described in Section 1, to the evolution of several social artifacts to relax such constraints, namely the establishment of both organised markets and money.”

In much of the standard economics literature, the constraint that is used is either income or wealth; but this is a secondary and adjustable constraint. If I devote more time to earning work, I can obtain more income. If I had infinite time in my life, I could have infinite income. Similarly, if I had more earning time, I could add to my wealth, to an indefinite degree. At any point of time, available income and wealth are given, but, over time, income and wealth are a function of how the individual allocates their time. Moreover, of course, there is intertemporal substitution. If I spend more time on trying to enhance my human capital at an early age, thereby giving up my use of time in other ways, e.g. for recreation, relaxation, earning, etc., then I may be able to enjoy more income and wealth at a later age. While everyone has to decide how to use their available income best at a point in time, the main optimisation decisions that anyone has to make on a continuous basis is how they want to use their time best. Humans have no ability to affect the inexorable passage of time, (ignoring space travel possibilities and other such sci-fi fantasies). At best, humans may have a choice between a healthier lifestyle and enjoying the immediate pleasures of food, drink and drugs.

¹ My thanks are due to George Evans for guiding me through parts of the literature, but all the views and the errors are mine.

In my earlier book (ibid), I went on from this initial discussion of time as the essential constraint to discuss how markets and money can act to improve our allocation of our available time. But one feature that I did not discuss then was how treating time as the ultimate constraint should affect understanding and analysis of learning and expectations. This is the purpose of this paper. More generally, it is my conjecture that a shift in macroeconomics to the treatment of time as the effective, ultimate constraint, would have significant and beneficial effects on many aspects of macroeconomic analysis. But I have been waiting in vain for any such revolution. It simply is not happening.

Learning takes time

Again as a generality, learning often requires little or no expenditure in income, nor the use of wealth. It is therefore treated in much of macroeconomics as a free good. As a result, agents are supposed to go on absorbing information and learning, until their expectations approximate to the best available underlying model of economic system. This is treated as 'rational'. Admittedly not all learning processes are free, notably now higher education, and books are expensive. But the web is largely free to use; and primary and secondary education is provided free by the state to all students.

But the real cost of learning is that it takes time. Every hour that I spend learning some new skill is an hour that I cannot use for some alternative purpose. Many, perhaps most, of which alternative uses of time are more immediately attractive and satisfying. Note that the young are prone to prefer holidays to being made to study at school. Because parents and the state believe that the young are incapable of appreciating the underlying longer-term benefits of schooling, the young are not allowed to decide for themselves on their allocation of time. But as they get older, everyone has greater freedom to allocate their scarce and limited time to a range of activities.

With learning requiring time, and time being scarce, we commonly consciously decide *not* to learn about most aspects of modern life, and we tend to specialise our learning into those areas where we feel, rightly or wrongly, that we have some comparative advantage. Take languages for example. Since English has become the common universal language, a large proportion of English speakers will

not bother to learn a foreign language. In contrast, most of those in professional classes in non-English speaking countries will spend time to become proficient in English.

One of the important questions in education is the degree to, and age at, which children should be allowed and encouraged to specialise. What are the macroeconomics of that decision?

A normal human will become individually expert in a relatively small range of subjects. Few will know how electronics, or the internal combustion engine, or television actually work. We just take it for granted. By the same token, only a tiny proportion of the population has studied subjects like climate change or the workings of a modern macroeconomy. For most people it is simply not worth the time and effort to do so. And this is a rational and optimal use of time.

When a subject arises to which we have not devoted our personal specialisation of learning, how do we form an opinion, or an expectation about it? For example, how do we expect climate change to influence the world over future decades, or Brexit to affect the UK economy over the medium term? What we do *not* do is to go back to the original scientific articles on climate change, nor do we reread the latest macroeconomic textbooks. In most cases we would not understand them; and in any case the cost in time would be excessive. So, what we do when we form opinions, and also expectations, about such issues is to react to, often simplified, publicly available, narratives about such issues.

Only a small handful of academic economic studies note that learning takes time, and that such an allocation of scarce time represents a cost. These include such experts as Evans and Ramey (1992) and Hommes (2013). But even in these examples the choice that they present lies between making no attempt to learn, and using an accurate model of the economy at some positive cost (Evans and Ramey, p. 210; Hommes, p. 15); though how the agents get hold of such an accurate model is not described. Neither recognize that, when uncertain of the validity of our own forecasts, the standard, and obvious, response is to ask, or consult, someone else who, we have reason to believe, is better informed than ourselves. When suffering an illness, do we form an expectation of the malady and its cure by reading the Lancet or a medical dictionary, or by consulting a doctor?

Learning is a social exercise

Basically we learn from others in the majority of cases, rather than from our own independent study. We simply do not have time to do the latter. When we form our opinions and expectations, this will largely be done by interactions with others, either directly by talking with people whose views we take seriously, or indirectly by reading narratives in books or other media. Should there be no particular narrative that we find gripping, the default position in almost any case is to assume that the future will be like the recent past.

But when events change violently, and the more important such events may be to ourselves, the more that we will search around alternative narratives about what may happen in future, and at the limit, but very rarely, undertake some independent learning and research for ourselves. This suggests, inter alia, that strongly autoregressive expectations will be most common when the system is moving in a normal smooth pattern, whereas expectations will come to focus much more on narrative models for future developments when events have become violent. For example, when nothing much is happening in the Middle East, most people will tend to assume that the price of petrol next month or next year will remain constant at its present level. But should a war break out in the Middle East, we would all look much more carefully at what experts were saying about the implications of such shocks. Thus the idea that expectations are generated in some constant fashion is itself an error. The generating structure of expectations is itself endogenous to the distribution of shocks and the importance of the outcome of such shocks on our own individual wellbeing.

Most studies of learning, and resultant expectations, treat the individual as being an isolated atom, not interacting with anyone else. This is, of course, nonsense. As already noted, most of our opinions and views derive from social interactions with others. Since the world is complex and the future unknowable, there are many alternative ways of trying to explain how the system works; in many cases these cannot be verified, e.g. what happens after death? In such normal cases there are a range of narratives about such issues, and the individual will choose between the narratives that they find most attractive based on many factors such as subjective personality, the persuasiveness of the narrative, what the authorities are claiming, etc., etc. It is perfectly possible that, from time to time, we may focus upon an invalid narrative, even within the natural sciences (phlogiston). The

likelihood of focusing on an invalid narrative is, of course, even greater in the social sciences and humanities, let alone religion.

With competing narratives on offer, there are bound to be bubbles and busts in asset markets. If a narrative indicating that asset X will rise in price appears to be successful, it will gain adherents, particularly when people see others making money, irrespective of the fundamental truth of that narrative. But when the bubble finally bursts and the narrative becomes discredited, then a rapid bust is likely to happen. But none of this is irrational. It simply follows from the fact that learning takes time and is, hence, costly. As a result, we try to choose from a variety of narratives about what the implications of any action might be. One of the tricks involved in this exercise is to make one's narrative persuasive, brief, straightforward and simple. For example, in the Brexit referendum, the phrase 'Take back control' was hugely influential, although more detailed analysis would have indicated that the degree of extra control obtained by leaving was limited.

The idea that we learn individually, by a Bayesian process of updating our own forecasts on the basis of the divergence of our prior forecast from the ex post actual, is deeply engrained in macro-economists. In practice, we explore whose other forecasts, and narratives, appear to have explained recent phenomena better. Hommes (2013) gets nearest to this, where he allows his subjects to choose between four different heuristics, adaptive expectations (ADA), a weak trend rule (WTR), a strong trend rule (STR) and a learning anchor and adjustment rule (LAA), (p. 33 and Chapters 7 and 8). He shows that, in laboratory experiments, agents will switch towards the currently more successful heuristic. Such social learning approaches sit within a broader category of agent based models, a field with a substantial literature, see e.g. C.H. Hommes, 2006.

But switching from one simplified heuristic to another, e.g. on the basis of past performance, is uncommon. Instead, we are, perhaps, more likely to switch allegiance from one expert to another, e.g. on inflation between Monetarists and Keynesians. Such a switch of heuristics is not unknown. Jim Bullard, another expert in this field, (e.g. Bullard, Evans and Honkapohja, 2010), has argued for switching from the heuristic that, absent shocks, the macro-economy will revert to its trend equilibrium, (as in most DSGE models), to the heuristic that, absent clear evidence to the contrary, the economy will remain in its current state.

Nevertheless, if the need is for a formal model of learning, the best currently available is the ‘Social Learning’ approach of Arifovic, Bullard and Kostyshyna (2012); also see Granato, Guse and Wong (2008). ABK state (p. 39) that,

“In this article we study an alternative approach to learning, one that can be viewed as more realistic in terms of actual learning in complicated market economies. In it, agents are initially heterogeneous with respect to the models they use to forecast the future. Forecast rules are updated via genetic operators, meant to simulate the process of learning from neighbours and others in the economy. Results are not analytic but are based on computational experiments. We will call this alternative approach *social learning*.”

The results of such exercises do not suggest that reasonable learning processes necessarily converge to a rational expectations equilibrium, and in many cases such algorithms can allow exaggerated volatility to recur, i.e. booms and busts.

Conclusion

To conclude, it is hard to disagree with Bray and Kreps (1987), who wrote (p. 623) that,

“Perhaps the most accurate statement that we can make is that our analysis also applies to *learning about* equilibria when (i) the learning is rational (Bayesian) on the individual level and (ii) agents’ models are consistent (and consistent with reality). This need not and should not be read as an assertion that only models that satisfy such stringent requirements are interesting. In fact, we are of the opinion that such models are of less interest than those that have in place some level of individual irrationality or some level of inconsistency or both. It seems incredible indeed that agents are so rational and consistent. The models analysed here provide a benchmark, but we strongly believe that this is a sterile benchmark and that there is much more to be learned by studying models that are somewhat less sterile. A failing of economic theory in general has been that it has proved remarkably resistant to movements away from models with full rationality and consistency; here, perhaps, is an area where there has been a bit of progress, and the promise of a good deal more.”

Over 30 years on, we are still waiting for a combination of good sense and empirical research to break the grip of the individual rational expectations model.

Much of the behavioural literature has focussed on various features of human behaviour, such as loss aversion, myopia, short-term impatience, ignoring extreme possible outcomes, overweighting, recent experience, etc., etc. All of these are, in my view, largely correct. Certainly the work of the behavioural economists, such as Kahneman and Tversky. (e.g. 1972, 1973, 1979, 1982, 1992, 2011), and Richard Thaler (1994, 2009, 2016), and others who have followed in their path, are praiseworthy; and indeed show important aspects of human behaviour and the use of heuristics.

But, in my view, most of the literature still largely misses the crucial point. This is that the process of learning by oneself is extremely costly in terms of time allocation. As a result, almost all learning is a social exercise, whereby we learn directly, or indirectly, from the expressed views of others. Put it another way, informational contagion is baked into our social and economic systems.

References

Arifovic, J., Bullard, J., and O. Kostyshyna, (2012), 'Social Learning and Monetary Policy Rules', *The Economic Journal*, 123 (March), pp 38-76.

Bray, M., and D.M. Kreps, (1987), 'Rational Learning and Rational Expectations', Chapter 19 in Feiwel, G.R., *Arrow and the ascent of modern economic theory*, Palgrave Macmillan.

Bullard, J., Evans, G.W., and S. Honkapohja, (2010), 'A Model of Near-Rational Exuberance', *Macroeconomic Dynamics*, 14, pp 166-188.

Evans, G.W., and G. Ramey, (1992), 'Expectation Calculation and Macroeconomic Dynamics', *The American Economic Review*, Vol. 82, No. 1, March, pp 207-224.

Goodhart, C.A.E., (1975), *Money, Information and Uncertainty*, Macmillan Education Ltd.

Granato, J., Guse, E.A., and M.C. Sunny Wong, (2008), 'Learning from the Expectations of Others', *Macroeconomic Dynamics*, 12, pp 345-377.

Hommes, C.H., (2006), *Computational Economics: Heterogeneous Agent Modeling*, Amsterdam: North Holland.

Hommes, C., (2013), *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*, Cambridge University Press.

Kahneman, D., and A. Tversky, (1972), 'Subjective Probability: A Judgment of Representativeness', *Cognitive Psychology* 3, pp 430-54.

Kahneman, D., and A. Tversky, (1973), 'On the Psychology of Prediction', *Psychological Review*, 80, no. 4, pp 237-51.

Kahneman, D., and A. Tversky, (1979), 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica*, 47, no. 2, pp 263-91.

Kahneman, D., and A. Tversky, (1982), 'The Simulation Heuristic', in *Judgment under Uncertainty: Heuristics and Biases*, edited by D. Kahneman, P., Slovic and A. Tversky, Cambridge University Press, pp 3-22.

Kahneman, D., (2011), *Thinking, Fast and Slow*, New York: Farrar, Straus and Giroux.

Thaler, R.H., (1994), *The Winner's Curse: Paradoxes and Anomalies of Economic Life*, Princeton University Press.

Thaler, R.H., and C.R. Sunstein, (2009), *Nudge: Improving Decisions About Health, Wealth and Happiness*, Penguin.

Thaler, R.H., (2016), *Misbehaving: The Making of Behavioural Economics*, Penguin.

Tversky, A., and D. Kahneman, (1992), 'Advances in Prospect Theory', *Journal of Risk and Uncertainty*, 5, pp 297-323.