Full name

Flavio Toxvaerd

Faculty or Department

Faculty of Economics

Your E-Mail Address

fmot2@cam.ac.uk

Research Title and Abstract as provided upon application to CERF

Rumours and Information Transmission in Real and Financial Markets

Key Research Findings to date

In this project, I analyse the propagation of rumours and information in populations in which individuals make economic decisions such as investments, decisions about where to live and work and which portfolios to hold. In doing so, I make use of rumour models developed within the field of mathematical epidemiology. Using the analogy of an "infection of the mind", these models characterise how information, spreading from one individual to another, propagates through a population. In contrast to the existing theory, I couple these (purely mechanical) models with models of economic and financial decision making.

Comments and Additional Information

Final CERF Fellow's Report for Flavio Toxvaerd

In this project, I have worked towards a rigorously micro-founded model of information transmission in real and financial markets, by building on two separate disciplines, namely so-called rumour theory as developed within mathematical epidemiology and economics/finance.

Rumour theory has the strength of being mathematically elegant and amenable to formal analysis and as such is an attractive framework for analysis. The one major drawback of this approach, is that it builds on the assumption that individuals are automata that mechanically spread (or cease to spread) rumours or information. This approach, while workable from a practical and mathematical perspective, is undesirable from an economics perspective, because it makes strong assumptions directly on behaviour (in contrast to making assumptions on preferences and information). The challenge then, is to make use of the nice aspects of rumour theory while making use of the methodology of economics to derive and characterise (rather than assume) individual decisions on information propagation and how these change endogenously over time.

It turns out that even very small and seemingly trivial departures from the assumption of non-optimising behaviour on the part of individuals, makes formal analysis extremely difficult. The reasons are twofold. First, sophisticated optimising individuals must try to forecast what effects their present decisions may have on future information propagation and thus ultimately on payoffs. This is a difficult task, even in the simplest environments. This difficulty can be sidestepped by assuming that individuals are myopic and thus only maximise present payoffs. But even with myopic individuals, there is an additional obstacle that seriously complicates the analysis, namely that individuals must form beliefs about possible histories of the rumour. To be specific, an individual who receives information must try to determine the veracity of the information and in doing so, must assess the most likely source of the information, the most likely route through which the information travelled before reaching the individual and the preferences and incentives of all past propagators of the information to truthfully pass on the information (or indeed to garble the information or to otherwise misrepresent it).

As can be readily appreciated, this is an extremely difficult task to perform, both for the individual and for the modeller. The few contributions in the literature that have in fact accomplished this, build on somewhat restrictive assumptions that are not easily generalisable.

An alternative approach to analysing information propagation in structured populations, is using techniques from the study of cellular automata. In particular, such techniques can be used to simulate the dissemination of information and behaviour in populations, by embedding individuals in two-dimensional grids and specifying rules of behaviour of each individual as a function of the behaviour of neighbouring individuals. The main problem with this approach, while technically not trivial, is a philosophical one. In general, the profession has been lukewarm to work using such techniques, because it is difficult to interpret the results that one obtains. In particular, many researchers have raised the objection that such analyses have little predictive power and that it is unclear what conclusions about real markets can be drawn from an analysis of interacting, but largely non-optimising, individuals (i.e. that are very boundedly rational, ill-informed and myopic). In some sense, the cellular automaton approach suffers from some of the same drawbacks as mathematical rumour theory.

My work on this project has helped me clarify a number of very delicate issues in the modelling of rational rumours and to understand the important limitations that rational decision making imposes on the analysis of person-to-person information transmission in structured populations. As I have built a clearer picture of the difficulties involved, I have also started pursuing new modelling approaches that I hope will be both tractable and yield interesting results.