



BANK OF ENGLAND

Working Paper No. 503

Peering into the mist: social learning over an opaque observation network

John Barrdear

August 2014

Working papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee.



BANK OF ENGLAND

Working Paper No. 503

Peering into the mist: social learning over an opaque observation network

John Barrdear⁽¹⁾

Abstract

I present a model of social learning over an exogenous, directed network that may be readily nested within broader macroeconomic models with dispersed information and combines the attributes that agents (a) act repeatedly and simultaneously; (b) are Bayes-rational; and (c) have strategic interaction in their decision rules. To overcome the challenges imposed by these requirements, I suppose that the network is *opaque*: agents do not know the full structure of the network, but do know the link distribution. I derive a specific law of motion for the hierarchy of aggregate expectations, which includes a role for *network shocks* (weighted sums of agents' idiosyncratic shocks). The network causes agents' beliefs to exhibit increased persistence, so that average expectations overshoot the truth following an aggregate shock. When the network is sufficiently (and plausibly) irregular, transitory idiosyncratic shocks cause persistent aggregate effects, even when agents are identically sized and do not trade.

Key words: Dispersed information, network learning, heterogeneous agents, aggregate volatility.

JEL classification: C72, D82, D83, D84.

(1) Bank of England. Email: john.barrdear@bankofengland.co.uk

The views expressed in this paper are those of the author and not necessarily those of the Bank of England. This paper was developed at the London School of Economics before the author's employment by the Bank of England. I am indebted to Kevin Sheedy, Wouter DenHaan, Kristoffer Nimark, Andrea Prat, Francesco Nava, Tony Yates, one anonymous referee and various seminar and workshop participants at the London School of Economics and the Bank of England for their invaluable feedback. This work was partially funded by the ESRC. This paper was finalised on 23 June 2014.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

Summary

Communication pervades human existence, and economic behaviour is no exception to this rule. In addition to the myriad of cultural interactions, people directly share economic information such as job opportunities and prices, and indirectly reveal information to each other as they trade goods and services. The study of how information is shared over a network of interactions is therefore an important field of economic research.

The topic of social learning – examining if, how, and how quickly people's beliefs might converge – when people communicate via a network has been examined extensively in the microeconomic literature. There has been little to no application to questions of macroeconomics, however, despite the common acceptance that imperfect access to information is critical to explaining the movement of aggregate variables. For example, firms' price-setting decisions may be influenced by observing each other's individual prices.

The reason that network learning has not been previously explored within macroeconomics is that three other features commonly deemed essential to the discussion of macroeconomics – that agents act repeatedly; that agents act strategically, with their payoffs a function of other agents' actions; and that although imperfectly informed, agents' expectations are (close to) rational – make comprehensive analysis of network learning intractable in anything other than trivially small networks.

This paper presents a solution to this problem by proposing a simplifying assumption: that the network is 'opaque' in that economic players ('agents') such as households or firms do not know exactly who is connected to whom. Instead, it is supposed that agents know the probability distribution from which everybody draws the identity of their observees. That is, it is known that agent 1 is observed with a specific probability, agent 2 is observed with another probability, etc. The model also includes a key feature of actual networks by supposing that while most agents are unlikely to be observed, some groups of agents are disproportionately highly observed, even as the number of people in the network becomes very large.

Agents are attempting to learn about an unobserved or hidden 'state' variable (e.g., the level of demand) by observing each other's actions. In the paper, the way that the possible expectations of this state (the 'hierarchy' of expectations) adjust over time is derived. With an opaque network, the hierarchy includes the average expectation regarding the hidden state, the average expectation of the average expectation, etc., but also includes an infinite sequence of weighted-average expectations and higher-order combinations between them.

Following a shock to the hidden state, average expectations respond more quickly than they do when agents do not observe each other in a network, but also temporarily overshoot the truth in a kind of herding behaviour that relies on the agents' observations of each other and their strategic motives (strategic meaning that they act taking into account beliefs about how others will respond). The degree of persistence of expectations is shown to be increasing in the number of competitors observed.

Idiosyncratic shocks (that is, those that affect only individual agents), which in many models have no effect on aggregate variables, are shown to influence the hierarchy of aggregate beliefs. Even when idiosyncratic shocks last only one period, these effects are also shown to be persistent, lasting for several periods. The paper therefore contributes to a new field of research demonstrating that aggregate volatility may emerge from idiosyncratic shocks.

Because of the focus on a setting with an underlying state that evolves over time and the way the hierarchy of average expectations evolves, those interested in exploring models of this type are able to determine the aggregate effects of network learning without a need to simulate individual agents' decisions. This makes the model particularly amenable to nesting within broader general equilibrium models of the economy that take account of all the interactions within and between different sectors of the economy – in other words, macroeconomic models.



1 Introduction

This paper attempts a partial bridging of three strands of research, namely those of *network learning* in the microeconomic theory literature; *dispersed information* in macroeconomic modelling; and more recent work identifying *idiosyncratic origins of aggregate volatility*. Although models at this intersection are not generally solvable, with a plausible restriction on the observability of the network structure I am able to derive a law of motion for the full hierarchy of expectations and demonstrate that a finite approximation of the same may be made arbitrarily accurate. I further show that network learning causes aggregate beliefs to be more persistent than the shocks that cause them, so that average expectations *overshoot* the truth. The observation network also provides a channel for idiosyncratic shocks to cause aggregate volatility even when all agents are identically sized and do not trade with each other. The model may be applicable to a variety of financial and macroeconomic settings, including firms' price-setting, search and matching models and financial asset pricing.

The *dispersed information* literature, which started with [Woodford \(2003\)](#), is one of three strands of research (the others being *sticky information* and *rational inattention*¹) that seek to reintroduce the ideas of [Lucas \(1972\)](#) and [Phelps \(1984\)](#) – that information frictions are crucial to explaining the dynamics of aggregate variables following a shock.

In particular, Woodford invoked the central insight of [Townsend \(1983\)](#) – that with heterogeneous information and strategic interaction, rational agents become interested in an infinite regress of higher-order beliefs – and demonstrated that because of the sluggish response of higher-order expectations, aggregate rigidity broadly equivalent to that produced by [Calvo \(1983\)](#) pricing may be replicated in a model with fully flexible price-setting firms observing independent and unbiased signals of nominal GDP. A flurry of further work has ensued,² but three key attributes are typically seen as essential features of such models, beyond the simple fact of a hidden dynamic state and heterogeneous information:

1. agents act repeatedly;
2. agents update their beliefs in a Bayesian and model-consistent (i.e. rational) manner;³ and
3. agents act strategically, with their payoffs a function of other players' actions.

These requirements are standard and largely uncontroversial features of macroeconomic modelling. The first is a defining feature of any dynamic model, the second is considered necessary to address the [Lucas \(1976\)](#) critique and the third is both analytically necessary to generate the higher-order expectations of [Townsend \(1983\)](#) and widely observed in a variety of financial and macroeconomic settings, as listed above.

The addition of *network learning* – whereby individual agents observe the actions of specific competitors in order to learn about a hidden state – would appear a natural extension to the dispersed information literature, particularly when acquiring comprehensive information would be prohibitively

¹The sticky information literature derives from [Mankiw and Reis \(2002, 2006, 2007\)](#), while the rational inattention literature dates to [Sims \(2003\)](#).

²See, for example, [Nimark \(2008\)](#), [Lorenzoni \(2009\)](#), [Angeletos and La'O \(2009, 2010\)](#), [Graham and Wright \(2010\)](#), [Graham \(2011a,b\)](#) and [Melosi \(2014\)](#).

³The *learning* literature (see, e.g., [Evans and Honkapohja, 2001](#)) explores settings under which non-model-consistent expectations converge to model-consistent expectations. A model by [Graham \(2011b\)](#) demonstrates that such dynamics at the aggregate level are dominated by dispersed information (i.e. that learning converges quickly to model-consistent expectations relative to the persistence of dispersed expectation errors).

costly. For example, a firm that experiences shocks to its demand or its marginal costs but does not know whether these are common to all firms can partially infer the average by observing the price of a competitor. A trader wanting to learn the level of demand for a given asset can improve their standing by speaking to other traders. Households that experience complementarity in their consumption will improve their welfare by observing the choices of their neighbours.

A model of network learning that possesses all of the above attributes is notoriously difficult to study, however. Solving such a model, let alone simulating it or nesting it within a broader model of the economy, has typically been thought to be sufficiently great as to preclude comprehensive analysis in anything other than trivially small networks (Jackson, 2008). As such, the literature to date has proceeded by avoiding one or more of the above assumptions (see below for a brief review).




Figure 1: An example of an observation network. Arrows indicate the direction of observation (agent 1 observes agent 2, etc). It is “regular” in that all agents observe, and are observed by, the same number of competitors. The adjacency matrix (G) is such that $G_{ij} = 1$ if node i observes node j and zero otherwise.

In contrast to earlier work, the present paper is able to embrace all three of these assumptions by combining them with a fourth: network opacity. By denying agents knowledge of the exact topology of the network (the network is *opaque*) and instead supposing that they know only the distribution from which observation targets are drawn, I derive the law of motion for the full hierarchy of agents’ expectations and show that *network shocks* (weighted sums of agents’ idiosyncratic shocks) enter at an aggregate level.⁴ The researcher is therefore able to simulate the aggregate effects of network learning without a need to simulate individual agents’ decisions. This makes the model particularly amenable to nesting within broad general equilibrium models of the economy.

With an opaque network, agents switch from considering their competitors’ *individual* beliefs to instead contemplating a sequence of *weighted averages* of all agents’ beliefs. Since agents’ learning is recursive, this allows the curse of dimensionality to be overcome in practice, as an arbitrarily accurate approximation of the full solution can be found by selecting a sufficiently high cut-off for the number of weighted-averages to include, together with the standard cut-off for the number of higher-orders of expectation.

The imposition of an opaque network is both intuitive and appealing. It is not plausible, for example, to suppose that every business knows to whom every other business speaks, just as nobody knows the identity of all of their friends’ friends. From the researcher’s perspective, this ignorance of topology makes it particularly challenging when attempting to consider the aggregate effects of network learning. But by recognising that not only the researcher but also the economic agents

⁴With the underlying (and unobservable) state following an AR(1) process, the full hierarchy of expectations about it will follow an ARMA(1,1) process, with network shocks entering both contemporaneously and with a lag.

themselves are ignorant of the network structure, the researcher can identify laws of motion for the agents' aggregate beliefs, even if they can never pin down the path of any individual's expectation.

The second requirement – that agents not learn about the structure of the network over time – may be thought of in two ways. First, one might consider a setting in which the network is dynamic, changing every period. *In extremis*, this would involve the network being destroyed and redrawn each period, so that agents are not *able* to learn about the network as it does not persist over time. Alternatively, one might suppose that the network was drawn once, at time zero, but agents are boundedly rational in that they do not attempt to learn about it beyond the common knowledge of the distribution from which it was drawn. In this setting, agents' decisions are perhaps best described as *conditionally rational*, in that conditional on the structure of the network, they are rational in their processing of the information they gain from it.

For sufficiently irregular networks – i.e. where some agents' actions are disproportionately observable – I show that network shocks do not converge to zero and therefore add aggregate volatility to the system, even when all agents are the same size (e.g. even when all firms contribute the same share of aggregate production). Despite idiosyncratic shocks being purely transitory, the aggregate volatility they induce through the network is also shown to exhibit (endogenous) persistence.




Figure 2: An irregular observation network. Agent 1 observes agent 2, while all others observe agent 1.

This paper therefore adds to the burgeoning literature on deriving aggregate volatility from agents' idiosyncratic shocks, within which the most closely related work is that of [Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Saleh \(2012\)](#), who examine a static model of firms operating within an intersectoral supply network (focussing on the cumulative effect of the network as a transmission mechanism of real shocks). By contrast, I present a dynamic model that considers the evolution of higher-order beliefs over each round of learning. Acemoglu *et al*'s emphasis on what they call “higher-order interconnectivity” in the network is captured and given an explicit dynamic role here. Finally, the observation network explored here may also clearly be different to the trading network of an economy.

In another vein, [Gabaix \(2011\)](#) demonstrates how aggregate volatility can emerge from idiosyncratic shocks when the distribution of firm sizes exhibits fat tails, even when those firms do not trade directly with each other. Each of these share with the current paper an emphasis on unequal, or fat-tailed, distributions. In the model of [Gabaix \(2011\)](#), aggregate volatility arises because the largest firms contribute disproportionately to aggregate production, while in that by [Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Saleh \(2012\)](#) it emerges through those firms whose output is most extensively used as an intermediate good by other firms. By contrast, the current paper demonstrates a granular origin of aggregate volatility even for identically sized agents that do not trade with each other by focussing, instead, on firms' signal extraction problem.

The model generates average expectations that are more persistent than the shocks that cause them, so that following an innovation to the underlying state, impulse responses overshoot the underlying state in a cascade effect that combines the herding exhibited in both [Banerjee \(1992\)](#), where agents observe others' actions, but have no strategic motive; and [Morris and Shin \(2002\)](#), where agents have a strategic motive, but do not observe each others' actions. The degree of persistence is shown to be increasing in the number of agents observed.

The intuition for this is as follows. Since the network is opaque, agents cannot know who their observee is watching. Common knowledge of the distribution from which observees are drawn, together with the linearity of the model, means that all agents treat all other agents *as though* they observe the same weighted average of everybody's action. In a model with strategic complementarity and dispersed information, public signals represent a source of herding, as famously shown by [Morris and Shin \(2002\)](#). With network learning over an opaque network, the observation of a competitor's action is therefore a signal of a hypothetical public signal that, to the best of each agent's knowledge, everyone else effectively observes. Bayesian updating therefore causes them to place extra weight on observations of other agents' actions.

Methodologically, this paper expands on the work of [Nimark \(2008, 2011a,b\)](#), who in turn extended that of [Woodford \(2003\)](#). While Woodford only granted agents signals of the underlying state, Nimark also permitted agents to observe, with a lag, aggregate variables that depend on the entire hierarchy of expectations. This addition required the development of a new solution methodology that I here extend to the idea of agents observing the previous-period actions of specific competitors.

Although the models of Woodford and Nimark focus on firms' price-setting behaviour, the model developed here is context free and may be applied to any general setting with strategic interaction and network learning. The conclusion considers a number of examples of such applications.

The remainder of this paper is organised as follows. The remainder of this introduction first provides a brief survey of previous models of network learning. Section 2 then provides some preliminary definitions related to graph theory, hierarchies of expectations and asymptotically non-uniform distributions. Section 3 next presents the general model, together with a characterisation of the solution and a methodology for finding it. Section 4 provides an illustrative example of the model in action, applying it to the commonly-used decision rule examined by [Morris and Shin \(2002\)](#) and [Calvó-Armengol and de Martí \(2007\)](#). Section 5 concludes.

1.1 Existing literature on network learning⁵

As mentioned above, literature on network learning has, to date, proceeded by avoiding one or more of the three assumptions that (a) agents are rational; (b) agents act simultaneously and repeatedly over many periods; and (c) agents' optimal decisions include consideration of strategic complementarity. Early work in observational learning, for example, focussed on *sequential learning*, with each agent making a single, irrevocable decision in an exogenously defined order, typically after observing the actions of all, or a well-defined subset, of their predecessors. In such a setting, it is well known that agents can rationally (in the Bayesian sense) exhibit "herding", or "information cascades", whereby their private signals regarding the unknown state are swamped by the weight of past actions (see, for example, [Banerjee, 1992](#); [Lee, 1993](#); and [Smith and Sørensen, 2000](#)).

⁵[Acemoglu and Ozdaglar \(2011\)](#) also provide a recent review.

More recently, work in sequential learning has examined situations where the observation neighbourhood of each agent is determined stochastically. [Banerjee and Fudenberg \(2004\)](#), for example, demonstrate that convergence will occur if the sampling of earlier players' beliefs is "unbiased" in the sense that it is representative of the population as a whole and at least two earlier players are sampled. More generally, [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#) characterise the (Bayesian) equilibrium of a sequential learning model for a general stochastic sampling process. They demonstrate that so long as no group of agents is excessively influential, there will be asymptotic learning of the truth when private beliefs are unbounded⁶ and characterise some settings under which asymptotic learning still emerges when private beliefs are bounded.

Although this more recent work carries the flavour of network learning in that agents observe the actions of only a subset of their competitors,⁷ they do not meet the popular conception of network learning in which agents undertake repeated, simultaneous actions in an environment of strategic interaction. Tackling such a problem, however, is notoriously difficult. The presence of strategic interaction introduces the need to consider the infinite hierarchy of higher-order (average) beliefs. When agents exist in an observation network, it is also necessary for each of them to consider the specific belief held by their observation target and, in turn, the belief of their target's target and so forth. As the number of agents in the network expands, this causes an explosion in the size of the state vector quite apart from the presence of higher-order expectations (see section 2.2 below for more detail), thereby subjecting the problem to the famous curse of dimensionality.

In order to analyse learning in a repeated, simultaneous action environment, the literature has therefore most commonly chosen to abandon the assumption of Bayesian updating. Non-Bayesian learning over a network is typically modelled in the style of [DeGroot \(1974\)](#), with agents applying a constant weight to their observations of competitors' actions. For example, [DeMarzo, Vayanos, and Zwiebel \(2003\)](#) explore situations where agents assume that signals they receive from observing each other contain *entirely new information*. In a setting where a finite number of agents wish to estimate an unknown, but fixed state $\theta \in \mathbb{R}^L$, they suppose that agents each receive a single, conditionally independent and unbiased signal of the state and then communicate their beliefs over multiple rounds. Imposing the assumption that agents update their beliefs via a simple and constant weighted sum greatly simplifies analysis, but introduces what the authors label "persuasion bias" from the agents' failure to properly discount the repetition of information they receive. [Calvó-Armengol and de Martí \(2007\)](#) extend this setting to provide an assessment of the welfare losses from "unbalanced," or irregular⁸ networks.

[Golub and Jackson \(2010\)](#) likewise study learning in a setting where agents "naïvely" update their beliefs regarding a fixed state of the world by taking weighted averages of their neighbour's opinions. In contrast to earlier work, they are able demonstrate that with such heuristic learning, individual beliefs converge to the truth for a broad variety of networks (provided they are sufficiently large) and provide upper and lower bounds on the rate of convergence.

In the area of what might be called "true" Bayesian network learning (repeated simultaneous actions with agents engaged in Bayesian updating), there has been remarkably little work to date. [Gale and Kariv \(2003\)](#) examine Bayesian network learning in a setting with *connected* networks⁹

⁶That is, where agents may receive arbitrarily strong signals so that the support of their posterior belief that the state is equal to a given possibility is $[0, 1]$.

⁷Indeed, [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#) refer to their model as one of learning over a social network.

⁸A regular network is one in which all nodes have the same number of inbound and outbound links.

⁹In this context, a connected network is one in which information is able to flow *from* any agent *to* any other agent.

and in which agents' payoffs depend only on the proximity of their expectation to the state (i.e. without any strategic interaction). They note that the “computational difficulty of solving the model is massive even in the case of three persons.” [Mueller-Frank \(2013\)](#) details a formal structure for Bayesian learning over an undirected social network (i.e. with pairwise sharing), allowing for a choice correspondence from information to actions (and general strategies for the selection between indifferent options) as opposed to outright decision rules, but notes the extreme practical difficulties of actually implementing such a rule, both for the agents in principle and the researcher more generally.

Furthermore, both [Gale and Kariv \(2003\)](#) and [Mueller-Frank \(2013\)](#) step away from consideration of strategic interaction in agents' decision-making, so that when observing any competitor, every agent knows that their action is driven entirely by their belief regarding the underlying state.

2 Preliminary definitions

2.1 Network terminology

Only terms necessary for this paper are provided here. Readers interested in a more comprehensive introduction are directed to more general texts on networks in economics.¹⁰

A *network*, or *graph*, is a collection of nodes and edges.¹¹ In this paper, each agent will be a node in a network. A network's *adjacency matrix*, G , is such that $G_{ij} = 1$ if node i is connected to node j and zero otherwise. An *undirected* network is one in which $G_{ij} = G_{ji} \forall i, j$, while a *directed* network permits $G_{ij} \neq G_{ji}$. Figure 1 illustrates an example of a directed network comprised of six nodes.

For a directed network, the *out-degree* of a node is the number of edges originating at that node ($d_i^{out} = \sum_{j=1}^n G_{ij}$), while the *in-degree* of a node is the number of edges arriving at it ($d_j^{in} = \sum_{i=1}^n G_{ij}$).

A *regular* network is one in which all nodes have exactly the same out-degree and in-degree, so that $d_i^{out} = d_i^{in} = \bar{d} \forall i$. An *irregular* network is one in which this is not the case. The network shown in figure 1 is regular, while that shown in figure 2 is irregular.

In the model of this paper, all nodes will have the same out-degree in that every agent will observe the same number of other agents ($d_i^{out} = q \forall i$). I will therefore interchangeably refer to agents' in-degree as simply their *degree* (d_i). The *degree sequence* of a network is the set $\{d_1, d_2, \dots, d_n\}$. Without loss of generality, I assume that nodes (agents) are arranged such that $d_1 \geq d_2 \geq \dots \geq d_n$.

A network is *connected* if it is undirected and a path exists from any node to any other node. A network is *strongly connected* if it is directed and a route (a directed path) exists from any node to any other node. A network is *complete* if every node is directly connected to every other node so that all possible edges exist.

A *cycle* of length k is a sequence of nodes, starting and ending with the same node, $\{i_1, \dots, i_{k-1}, i_k, i_1\}$ such that each pair of consecutive nodes are connected ($G_{i_j i_{j+1}} = 1$). A network is *aperiodic* if the greatest common divisor of the lengths of its cycles is one. In other words, a network is aperiodic if at least one node has a link to itself ($\exists i : G_{ii} = 1$).

In an environment of social learning where agents share their beliefs truthfully, a network being aperiodic implies that at least one agent places non-zero weight on their own prior when updating

¹⁰See, for example, [Goyal \(2007\)](#) or [Jackson \(2008\)](#).

¹¹Some papers refer to a network as a *weighted graph* – i.e. a graph with a weight associated with each edge.

their belief following the observation of their neighbours. In such a setting, if the network is both aperiodic and strongly connected, it is well known that all agents' beliefs will converge to the same values (Kemeny and Snell, 1960).

In the current paper, all agents will receive private, unbiased signals in every period. With Bayesian updating, every *sub-graph* of the network will therefore be aperiodic, so that convergence is assured and uninteresting. Instead, this paper focusses on the dynamics of aggregate beliefs over time.

2.2 Higher-order expectations

The near-ubiquitous treatment of higher-order expectations in economic literature to date¹² has considered only the hierarchy of *simple average* expectations. That is, to consider settings where agents are interested only in the sequence of objects $\{x_t, \bar{E}_t[x_t], \bar{E}_t[\bar{E}_t[x_t]], \dots\}$ where $\bar{E}_t[\cdot] \equiv \int_0^1 E_t(i) [\cdot] di$.¹³

This is a modelling choice only, however, made for analytical convenience. In particular, it is not appropriate for a model of learning over a network where economic agents must (in principle, at least) form opinions regarding the beliefs of every other agent in the network and know that they will each, in turn, do the same. To model these fully, it is necessary to work with a more generalised definition of a hierarchy of expectations.

Definition 1. A *compound expectation* is a weighted average of all agents' expectations.

For example, let \mathbf{x}_t be an $(m \times 1)$ vector of random variables, $E[\mathbf{x}_t | \mathcal{I}_t(i)]$ be the expectation of \mathbf{x}_t conditioned on the period t information set of agent i and let \mathbf{w} be an $(n \times 1)$ vector of weights such that $w_i \in [0, 1]$ and $\sum_{i=1}^n w_i = 1$. The compound expectation $E_{\mathbf{w},t}[\mathbf{x}_t]$ is given by:

$$E_{\mathbf{w},t}[\mathbf{x}_t] \equiv \begin{bmatrix} E[\mathbf{x}_t | \mathcal{I}_t(1)] & E[\mathbf{x}_t | \mathcal{I}_t(2)] & \dots & E[\mathbf{x}_t | \mathcal{I}_t(n)] \end{bmatrix} \mathbf{w} = \sum_{i=1}^n w_i E_t[\mathbf{x}_t | \mathcal{I}_t(i)] \quad (1)$$

Note that this nests both simple, or unweighted, average expectations (e.g. $\mathbf{w}_A = [\frac{1}{n} \dots \frac{1}{n}]'$) and individual expectations (e.g. $\mathbf{w}_B = [\mathbf{0}' \ 1 \ \mathbf{0}']'$). With the usual notation that the 0^{th} -order expectation of a variable is the variable itself, we next define:

Definition 2. A *hierarchy of expectations*, from order 0 to k , is defined recursively as:

$$\mathbb{E}_t^{(0:k)}[\mathbf{x}_t] = \begin{bmatrix} \mathbf{x}_t \\ E_{\mathbf{w}_A,t} \left[\mathbb{E}_t^{(0:k-1)}[\mathbf{x}_t] \right] \\ E_{\mathbf{w}_B,t} \left[\mathbb{E}_t^{(0:k-1)}[\mathbf{x}_t] \right] \\ \vdots \end{bmatrix} \quad (2)$$

This is not simply the stacking of each order of expectations on top of each other. For example, if \mathbf{x}_t is scalar and there are two compound expectations, the hierarchies $(0 : 1)$ and $(0 : 2)$ are given by:

¹²While modern macroeconomic literature on higher-order expectations dates to Townsend (1983), the general idea has been known since, at least, the famous “beauty contest” argument of Keynes (1936).

¹³One recent exception is Kohlhas (2013), who examines the value of central bank disclosure in a model with two compound expectations – that of the central bank and the average expectation of the private sector. Kohlhas extends the solution methodology of Nimark (2008, 2011a) in a similar manner to the current paper.

$$\mathbb{E}_t^{(0:1)}[\mathbf{x}_t] = \begin{bmatrix} \mathbf{x}_t \\ E_{w_A,t}[\mathbf{x}_t] \\ E_{w_B,t}[\mathbf{x}_t] \end{bmatrix} \quad \mathbb{E}_t^{(0:2)}[\mathbf{x}_t] = \begin{bmatrix} \mathbf{x}_t \\ E_{w_A,t} \begin{bmatrix} \mathbf{x}_t \\ E_{w_A,t}[\mathbf{x}_t] \\ E_{w_B,t}[\mathbf{x}_t] \end{bmatrix} \\ E_{w_B,t} \begin{bmatrix} \mathbf{x}_t \\ E_{w_A,t}[\mathbf{x}_t] \\ E_{w_B,t}[\mathbf{x}_t] \end{bmatrix} \end{bmatrix}$$

The benefit of depicting hierarchies in this recursive manner is that it becomes both conceptually and computationally simple to extract sub-hierarchies comprised of a single compound expectation. For example, if $E_{w_A,t}[\mathbf{x}_t] = \bar{E}_t[\mathbf{x}_t] = \bar{\mathbf{x}}_{t|t}^{(1)}$ is the simple-average expectation, then the sub-hierarchy of $\bar{\mathbf{x}}_t^{(0:k)}$ may be extracted as: $\bar{\mathbf{x}}_t^{(0:k)} = \begin{bmatrix} I & 0 \end{bmatrix} \mathbb{E}_t^{(0:k)}[\mathbf{x}_t]$. This recursive formulation of the expectation hierarchy is also a necessary feature of the solution methodology developed below.

Size of the expectation hierarchy

When \mathbf{x}_t contains m elements and there are p compound expectations of interest, the set of k^{th} -order expectations will contain mp^k distinct elements. However, it does not generally follow that the hierarchy $\mathbb{E}_t^{(0:k^*)}[\mathbf{x}_t]$ will contain $m \left(\sum_{k=0}^{k^*} p^k \right)$ unique elements. This is because if one of the compound expectations, say $E_{w_B}[\cdot]$, is formed from a single information set then the law of iterated expectations implies that $E_{w_B,t}[E_{w_B,t}[\mathbf{x}_t]] = E_{w_B,t}[\mathbf{x}_t]$. In general, when $q (\leq p)$ is the number of individual expectations, the number of unique elements in the hierarchy $\mathbb{E}_t^{(0:k^*)}[\mathbf{x}_t]$ will be:¹⁴

$$N(m, p, q, k^*) = m \left(p^{k^*} + \sum_{k=0}^{k^*-1} \left(p^k - q \sum_{s=0}^k p^s \right) \right) \quad (3)$$

with $N(m, p, 0, k^*) = m \left(\sum_{k=0}^{k^*} p^k \right)$. Even when $q = p$, though, it should be readily apparent that the size of an expectation hierarchy explodes in both p (the number of compound expectations) and k^* (the highest order in expectations), as figure 3 shows.

Approximations of model solutions must therefore be found by limiting attention to a finite subset of the full state. The vast majority of models in the dispersed information literature have $p = 1$ and place decreasing weight on higher-order expectations (i.e. the weight is decreasing in k). Provided that the variance of higher-order expectations remains bounded from above, these models can be approximated to an arbitrary degree of accuracy by imposing a limit, k^* , on the number of orders of expectation and including all orders from zero up to that cut-off.

In contrast, increasing the number of relevant compound expectations can be more problematic as there is rarely, if ever, an obvious reason for weighting them differently. In network learning, in particular, p will be equal to the number of nodes when the network is strongly connected.¹⁵

14

$$\begin{aligned} & m \left(\underbrace{[1]}_{0^{th} \text{ order}} + \underbrace{[p]}_{1^{st} \text{ order}} + \underbrace{[p^2 - q]}_{2^{nd} \text{ order}} + \underbrace{[p * (p^2 - q) - q]}_{3^{rd} \text{ order}} + \underbrace{[p * (p * (p^2 - q) - q) - q]}_{4^{th} \text{ order}} + \dots \right) \\ & = m \left(\left(\sum_{k=0}^{k^*} p^k \right) - q \left(\sum_{k=0}^{k^*-1} \sum_{s=0}^k p^s \right) \right), \text{ which rearranges to the equation in the text} \end{aligned}$$

¹⁵Since in a strongly connected observation network information will flow from all nodes to all nodes, the construction of agents' priors requires that they consider the beliefs of all other agents.

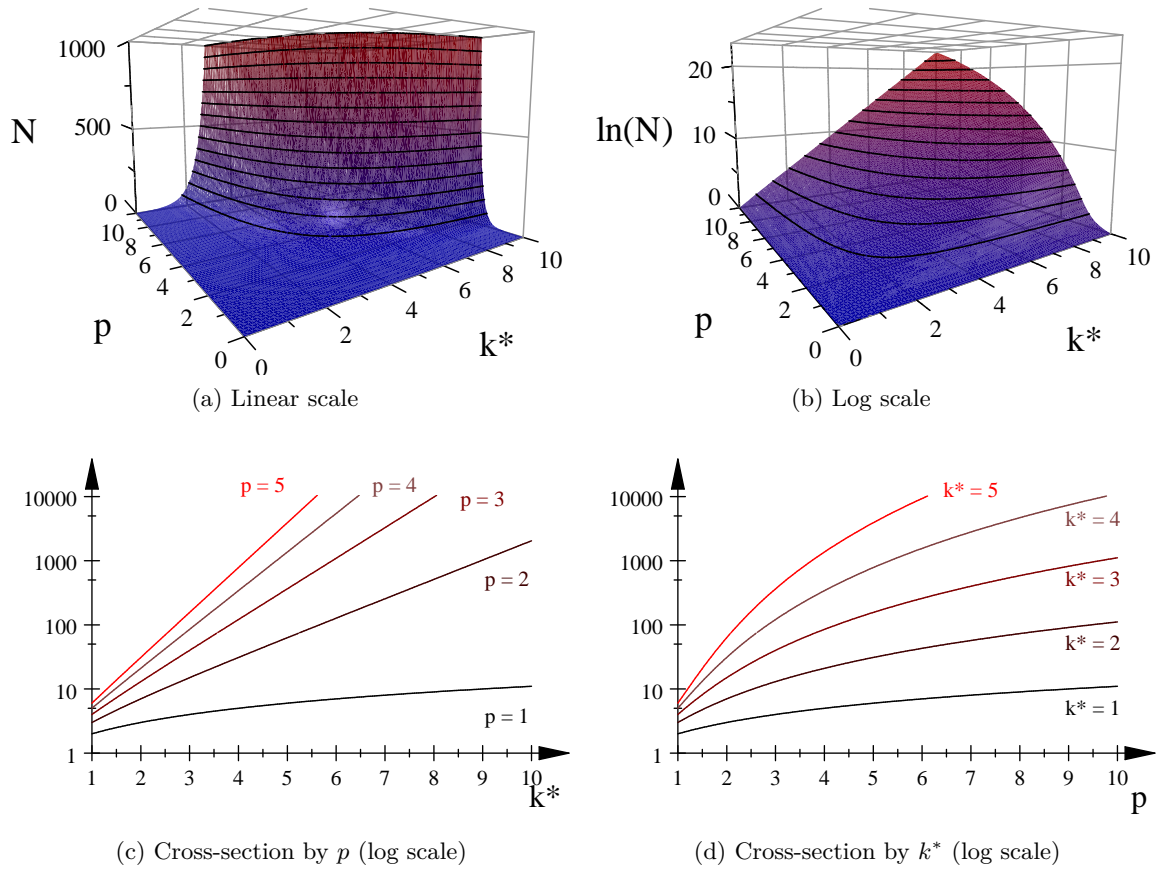


Figure 3: The number of elements in an expectation hierarchy ($q = 0, m = 1$)

2.3 Asymptotically non-uniform distributions

This paper develops a novel channel through which idiosyncratic shocks need not “wash out” and may, instead, induce aggregate volatility. This emerges because of agents’ need to consider *weighted* averages of other agents’ idiosyncratic shocks. In particular, a sufficient condition for such weighted sums to not converge to zero as $n \rightarrow \infty$ is to suppose that the weights are asymptotically non-uniform:

Definition 3. Let Φ_n be a discrete distribution with corresponding p.d.f.¹⁶ $\phi_n(i)$ and let $\zeta_n \equiv \sum_{i=1}^n \phi_n(i)^2$ be the Herfindahl index of the same. Φ_n is **asymptotically non-uniform** if:

- $\lim_{n \rightarrow \infty} \phi_n(i) = 0 \forall i$; and
- $\lim_{n \rightarrow \infty} \zeta_n = \zeta^*$ where $\zeta^* \in (0, 1)$.

To illustrate the emergence of aggregate volatility, suppose that each agent receives an independent, mean zero shock drawn from a common Gaussian distribution – $\mathbf{v}(i) \sim N(\mathbf{0}, \Sigma_{vv}) \forall i$ – and consider a setting where a weighted average of agents’ shocks is economically significant:

$$\tilde{\mathbf{v}}_n \equiv \sum_{i=1}^n \mathbf{v}(i) \phi_n(i) \text{ where } \phi_n(i) \in (0, 1) \text{ and } \sum_{i=1}^n \phi_n(i) = 1$$

¹⁶Strictly, for a discrete distribution, it is a probability *mass* function. But since this paper is concerned only with the limiting case of $n \rightarrow \infty$ assumes that they are indexed uniformly from zero to one so that the distribution becomes continuous, I stick with the conventional nomenclature.

Since \tilde{v}_n is a linear combination of mean zero Gaussian variables, it must itself be Gaussian with a mean of zero. Its variance will then be given by:

$$Var[\tilde{v}_n] = \sum_{i=1}^n Var[v(i)\phi_n(i)] di = \sum_{i=1}^n \Sigma_{vv}\phi_n(i)^2 di = \zeta_n \Sigma_{vv}$$

where the first equality relies on the shocks' independence to ignore covariance terms. The limiting variance as $n \rightarrow \infty$ is $\zeta^* \Sigma_{vv}$ and, hence, so long as $\zeta^* \neq 0$, the law of large numbers does not apply.

The set of asymptotically non-uniform distributions is quite broad, but in particular it includes the discrete power law distribution:

$$\phi_n(i) = c_n i^{-\gamma}; \text{ where } c_n = \left(\sum_{i=1}^n i^{-\gamma} \right)^{-1} \text{ and } \gamma > 1$$

and its equivalent for infinite n , the Zeta distribution. The shape parameter, $\gamma > 1$, governs the scaling of the distribution's tail, with larger values of γ corresponding to greater non-uniformity. Figure 4 plots the values of ζ^* for a range of values of γ for the Zeta distribution.¹⁷




Figure 4: A plot of ζ^* for power law distributions with shape parameter γ

A great many observed networks, from links between pages on Wikipedia to established relationships in social networks, have been shown to be well approximated by power law degree sequences (i.e. the networks are *scale free*). See, for example, the work of [Albert and Barabasi \(2002\)](#), [Jackson and Rogers \(2007\)](#) or [Clauset, Shalizi, and Newman \(2009\)](#). It is important to appreciate, though, that I do not generally assume any particular distribution, only that it remains non-uniform (in the sense of definition 3) as the support of that distribution grows arbitrarily large.

3 The Model¹⁸

3.1 The general setting

There is a countably infinite number of agents,¹⁹ indexed in a continuum between zero and unity.²⁰ The *underlying state* follows a vector autoregressive process:

$$\mathbf{x}_t = A\mathbf{x}_{t-1} + P\mathbf{u}_t \tag{4}$$

¹⁷Strictly, these are calculated for Zipf distributions with $n = 10^8$.

¹⁸Unless otherwise indicated, all proofs are provided in the appendix.

¹⁹An infinite number of agents is assumed to allow an appeal to relevant laws of large numbers when considering simple averages of zero-mean shocks.

²⁰The assumption of indexing agents from zero to one is innocuous and made only to simplify the calculation of averages (e.g. $\bar{g}_t = \int_0^1 g_t(i) di$).

where \mathbf{u}_t is a vector of shocks with mean zero, while A and P are fixed and publicly known. Agents do not observe the value of \mathbf{x}_t and must instead form beliefs about it. I define X_t as the hierarchy of expectations regarding \mathbf{x}_t , in the sense of definition 2, and refer to it as the *state vector of interest*.

$$X_t \equiv \mathbb{E}_t^{(0:\infty)} [\mathbf{x}_t] \quad (5)$$

At a minimum, X_t contains \mathbf{x}_t and the hierarchy of at least one compound expectation. For illustrative purposes, I will assume that agents' primary concern is with the hierarchy of *simple-average* expectations, so that

$$\bar{\mathbf{x}}_{t|t}^{(0:\infty)} \in X_t \text{ where } \bar{\mathbf{x}}_{t|t}^{(0:\infty)} \equiv \begin{bmatrix} \mathbf{x}_t' & \bar{E}_t [\mathbf{x}_t]' & \bar{E}_t [\bar{E}_t [\mathbf{x}_t]]' \dots \end{bmatrix}'$$

but it will be shown below that X_t must also include a variety of other compound expectations.

Agents' decision rule

Agents' actions are determined simultaneously and according to a common linear decision rule:

$$g_t(i) = \boldsymbol{\lambda}_1' E_t(i) [X_t] + \boldsymbol{\lambda}_2' \mathbf{x}_t + \boldsymbol{\lambda}_3' \mathbf{v}_t(i) \quad (6)$$

where $E_t(i) [\cdot] \equiv E[\cdot | \mathcal{I}_t(i)]$ is agent i 's (first-order) expectation conditioned on information available to her in period t (defined below); and $\mathbf{v}_t(i)$ is a transitory, mean zero shock specific to agent i in period t (defined below).

Non-zero elements in $\boldsymbol{\lambda}_1$ against higher-order average expectations capture strategic considerations in agents' actions. Note that the terms in \mathbf{x}_t and $\mathbf{v}_t(i)$ are included here to make the model as general as possible. They allow for the possibility that components of i 's signal vector may have direct economic significance in addition to their informational role. Note, too, that although \mathbf{x}_t may be included in agents' decision rule, it is *not* directly observed.²¹

Equation (6) nests a wide array of commonly studied settings and its derivation will invariably be context-specific. For example, in the model of Morris and Shin (2002), (6) would be written as

$$g_t(i) = (1 - \beta) \begin{bmatrix} 1 & \beta & \beta^2 & \dots \end{bmatrix} E_t(i) \begin{bmatrix} \bar{\mathbf{x}}_t^{(0:\infty)} \end{bmatrix} \quad (7)$$

where β is the weight placed on average actions. This example is explored further in section 4 below. Another example could be the setting of firms' prices, with strategic interaction arising from firms' demand schedules being a function of their relative prices.

Agents' information

Agents possess common knowledge of joint rationality, in the sense of Nimark (2008), so that they are aware of the structure and the coefficients of the system. Their information sets then evolve as:

$$\mathcal{I}_0(i) = \{\Omega, \Phi\} \quad \mathcal{I}_t(i) = \{\mathcal{I}_{t-1}(i), \mathbf{s}_t(i)\} \quad (8)$$

where Ω is the set of all system coefficients, $\mathbf{s}_t(i)$ is the signal vector received each period and $\Phi : [0, 1] \rightarrow [0, 1]$ is the (cumulative) distribution from which agents' observation targets in the

²¹For example, a firm may privately observe their productivity, which includes both aggregate and idiosyncratic components, but not their separate values. If their decision rule relies directly on their productivity, it will include a term in the aggregate productivity even though firms do not observe it directly.

network are drawn, assumed to be identical and independent, both across agents and across time. $\Phi(i)$ is absolutely continuous over the range $[0, 1]$ and has p.d.f. $\phi(i)$.

Each agent's signal vector is made up of two, distinct components. First, a combination of *public and private signals* based on the current underlying state or the lagged full state. These are essentially identical to those used in existing incomplete information work such as Nimark (2008). Second, each agent receives a *social signal* vector derived from observing competitors' actions over the network with a one-period lag:

$$\mathbf{s}_t(i) = \begin{bmatrix} \mathbf{s}_t^p(i) \\ \mathbf{s}_t^s(i) \end{bmatrix} \quad (9)$$

$$\mathbf{s}_t^p(i) = D_1 \mathbf{x}_t + D_2 X_{t-1} + R_1 \mathbf{v}_t(i) + R_2 \mathbf{e}_t + R_3 \mathbf{z}_{t-1}$$

$$\mathbf{s}_t^s(i) = \mathbf{g}_{t-1}(\boldsymbol{\delta}_{t-1}(i))$$

Public and private signals may include both current and lagged information and are noisy, including three sources of uncertainty:

- $\mathbf{v}_t(i)$ is a vector of transitory shocks specific to agent i in period t , drawn from independent and identical Gaussian distributions with mean zero and variance Σ_{vv} . These may simply be noise in agents' private signals or may carry economic significance, depending on the context.
- \mathbf{z}_t is a vector of *network shocks* (see equation 17 below), comprised of weighted sums of all agents' idiosyncratic shocks.
- \mathbf{e}_t is a vector of transitory "noise" shocks to public signals, drawn from an independent Gaussian distribution with mean zero and variance Σ_{ee} .

Although agents may observe signals based on the current underlying state (\mathbf{x}_t), they do not observe signals based on the current hierarchy of expectations about the state (X_t). This is because to do so would involve agents observing a signal based on their beliefs before they have formed them!

Terms in X_{t-1} and \mathbf{z}_{t-1} are permitted (instead of just \mathbf{x}_{t-1}) to allow agents to observe aggregate variables with a lag,²² and thus the past effect of their network learning.

Social signals are observations of the previous-period actions of specific agents, with the function $\boldsymbol{\delta}_t$ mapping each agent onto their observation targets:

$$\boldsymbol{\delta}_t : [0, 1] \rightarrow [0, 1]^q \quad (10)$$

where q is the number of agents observed. In other words, $\boldsymbol{\delta}_t(i)$ is the result of i 's q separate draws from Φ for period t . For presentational simplicity, I will typically assume that $q = 1$ (i.e. that all agents observe a single other agent) and simply write $j = \delta_t(i)$ to mean that agent j 's period- t action will be observed by agent i (in period $t + 1$) so that on the i^{th} row of the network adjacency matrix, we will have $G_{i,\delta_t(i)} = 1$.

To speak of the observee of an observee, one may write $\delta_s(\delta_t(i))$: the identity of the agent whose period- s action is observed by the agent whose period- t action is observed by agent i .

²²For example, if allocations are functions of the entire hierarchy of beliefs, then the publication by a national statistical organisation of an estimate of the previous period's GDP would be a function of both X_{t-1} and \mathbf{z}_{t-1} .

With agent i observing the previous-period action of a single competitor, their social signal is therefore given by

$$\begin{aligned} \mathbf{s}_t^s(i) &= g_{t-1}(\delta_{t-1}(i)) \\ &= \boldsymbol{\lambda}'_1 E_{t-1}(\delta_{t-1}(i)) [X_{t-1}] + \boldsymbol{\lambda}'_2 \mathbf{x}_{t-1} + \boldsymbol{\lambda}'_3 \mathbf{v}_{t-1}(\delta_{t-1}(i)) \end{aligned} \quad (11)$$

3.2 The observation network

Because agent i 's social signal is based on her observee's expectation, Bayesian updating then requires that i include $E_t(\delta_t(i)) [X_t]$ in her own state vector of interest. However, knowing that agent $\delta_t(i)$ is himself considering $\delta_t(\delta_t(i))$ then requires that i also maintain an estimate of $E_t(\delta_t(\delta_t(i))) [X_t]$, and so forth. This is the explosion of the state vector in p (the number of compound expectations) described in section 2.2 above. In order to make the problem tractable, I make two key assumptions:

Assumption 1. *The network is stochastic and opaque, in that:*

- *all agents observe the same number of other agents;*
- *observees are drawn from identical, fully independent distributions with p.d.f. $\phi(i)$;*
- *agents know the identities of the other agents they observe;*
- *agents do not know who they are observed by; and*
- *agents do not learn about the network topology over time.*

To obtain this last point, I suppose that agents make a fresh draw of whom to observe every period, in which case nothing *could* be learned about the network topology (since it changes every period).

Assumption 2. *The network is asymptotically irregular, in that its degree sequence is asymptotically non-uniform (see definition 3).*

As shown in section 2.3 and expanded on below, assumption 2 is sufficient to ensure that idiosyncratic shocks do not “wash out” in aggregation, but will instead enter into agents' average beliefs.

Note, too, that the unconditional expected (in) degree of agent i in a network of n agents will be $E_n[d_i] = q\phi_n(i)$, so that $E_n[d_i] \rightarrow 0$ as $n \rightarrow \infty$.

3.3 Agents' signal extraction problem

It will be shown below that the hierarchy of agents' expectations obeys the following vector ARMA(1,1) law of motion:

$$X_t \equiv \mathbb{E}_t^{(0:\infty)}[\mathbf{x}_t] = F X_{t-1} + G_1 \mathbf{u}_t + G_2 \mathbf{z}_t + G_3 \mathbf{e}_t + G_4 \mathbf{z}_{t-1} \quad (12)$$

where \mathbf{z}_t is a vector of transitory *network shocks*, derived as weighted sums of agents' idiosyncratic shocks. The exact statistical properties of \mathbf{z}_t are derived below in proposition 1.

The system described here is not in the form of a classic state space problem, however, both because of the presence of the lagged state in agents' signals (9) and because of the moving average component of the law of motion (12). The most common approach to addressing these features is

to stack the state vector with both its own lag and the lag of the shock with the moving average component, thus creating a combined state that follows an VAR(1) process:

$$\begin{bmatrix} X_t \\ \mathbf{z}_t \\ X_{t-1} \end{bmatrix} = \begin{bmatrix} F & G_4 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ I & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} X_{t-1} \\ \mathbf{z}_{t-1} \\ X_{t-2} \end{bmatrix} + \begin{bmatrix} G_1 & G_2 & G_3 \\ \mathbf{0} & I & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}_t \\ \mathbf{z}_t \\ \mathbf{e}_t \end{bmatrix}$$

and then to express agents' signals in terms of this combined state and estimate the system as a classic filtering problem. This approach more than doubles the size of the state vector, though, which may present problems when simulating the system with finite computing resources (and particularly so in the present setting with multiple compound expectations). Fortunately, the following lemma grants us that it is not necessary here to include \mathbf{z}_t in the state vector of interest.²³

Lemma 1. *Agents' contemporaneous expectations of the network shocks are zero:*

$$E_t(i) [\mathbf{z}_t] = \mathbf{0} \forall i, t \quad (13)$$

Since all individual agents' expectations of the network shock are zero, it must be the case that all average expectations (simple or weighted) of the network shock are also zero and since agents are jointly rational, this must be common knowledge. There is therefore no need to include any expectation of \mathbf{z}_t within the state vector to be estimated.

Because of the linearity of the system, the best linear estimator in the sense of minimising the mean squared error²⁴ will be a Kalman filter.²⁵

$$E_t(i) [X_t] = E_{t-1}(i) [X_t] + K \{ \mathbf{s}_t(i) - E_{t-1}(i) [\mathbf{s}_t(i)] \} \quad (14)$$

where K is a time-invariant projection matrix (the Kalman gain). As in other models of imperfect common knowledge, since X_t includes $\bar{\mathbf{x}}_{t|t}^{(0:\infty)}$, it must be that (a) the state vector to be estimated is of infinite dimension; and (b) the Kalman filter serves a dual role, both as estimator and as part of the law of motion for the state vector.

In the context of firms' price-setting behaviour, [Nimark \(2008\)](#) allowed agents to observe an aggregate signal (the average price) from the previous period in addition to their private signals. This means that each agent's signal vector includes a linear combination of the *entire hierarchy* of previous-period expectations. As a result, the solution must be found for all higher-order expectations simultaneously and the state vector of interest expands to include $\bar{\mathbf{x}}_{t-1|t-1}^{(0:\infty)}$ so that $X_t = \begin{bmatrix} \bar{\mathbf{x}}_{t|t}^{(0:\infty)'} & \bar{\mathbf{x}}_{t-1|t-1}^{(0:\infty)'} \end{bmatrix}'$.

An alternative to including $\bar{\mathbf{x}}_{t-1|t-1}^{(0:\infty)}$ in the state vector of interest is to retain the current signal vector and instead to modify the Kalman filter:

$$E_t(i) \left[\bar{\mathbf{x}}_{t|t}^{(0:\infty)} \right] = K \mathbf{s}_t(i) + (F - K(D_1 F + D_2)) E_{t-1}(i) \left[\bar{\mathbf{x}}_{t-1|t-1}^{(0:\infty)} \right]$$

This approach was first developed by [Nimark \(2008, 2011b\)](#) and is also used in the current paper to avoid the need to stack the state vectors of interest.

²³Unless otherwise stated, the proof of this and all further propositions may be found in the appendix.

²⁴With all shocks drawn from Gaussian distributions, it will be the best such estimator, linear or otherwise.

²⁵A derivation of the standard Kalman filter may be found in most texts on dynamic macroeconomics (e.g. [Ljungqvist and Sargent \(2004\)](#)) or time series analysis (e.g. [Hamilton \(1994\)](#)).

3.4 Peering into the mist

In implementing the Kalman filter (14), each agent must form a prior expectation of the signal(s) they will receive in the next period. Stepping equation (11) forward by one period, it is easily seen that it is necessary for agent i to construct $E_t(i) [E_t(\delta_t(i)) [X_t]]$ as part of her prior for period $t+1$:

$$\begin{aligned} E_t(i) [s_{t+1}^s(i)] &= E_t(i) [g_t(\delta_t(i))] \\ &= E_t(i) [\lambda_1' E_t(\delta_t(i)) [X_t] + \lambda_2' x_t + \lambda_3' v_t(\delta_t(i))] \end{aligned}$$

Constructing $E_t(i) [E_t(\delta_t(i)) [X_t]]$ requires, in turn, that agent i take a view regarding who $\delta_t(i)$ is observing: that is, the action of $\delta_{t-1}(\delta_t(i))$.

Lemma 2. *Given assumption 1 and common knowledge of rationality, agents' use of a linear estimator implies that all agents treat all other agents as though they observe a common, weighted average of previous-period actions, with the weights given by the distribution ϕ .*

From equation (6), it follows that the weighted-average action, \tilde{g}_t , is given by:

$$\begin{aligned} \tilde{g}_t &\equiv \int_0^1 g_t(j) \phi(j) dj \\ &= \lambda_1' \tilde{E}_t [X_t] + \lambda_2' x_t + \lambda_3' \tilde{v}_t \end{aligned} \quad (15)$$

It is not possible, in general, to make use of some law of large numbers to disregard the effect of idiosyncratic shocks in the weighted-average action – that is, one cannot assume that $\tilde{v}_t \equiv \int_0^1 v_t(j) \phi(j) dj$ will be equal to zero – because the weights applied to each agent may not be sufficiently close to equal. As an extreme example, if all agents were to observe agent 1 and nobody else (i.e. $\phi(1) = 1$ and $\phi(i) = 0 \forall i \neq 1$), then $\tilde{v}_t = v_t(1)$, which in any given period will be non-zero, almost surely.

Equation (15) is used for consideration of agents that are one step away in the observation network, but it is also necessary to consider the actions of agents that are two or more steps away:

Definition 4. *Let $\delta_t(i)$ be a period- t mapping from agent i to their observation target, drawn from the distribution Φ_n . The p^{th} -weighted average of agents' idiosyncratic shocks is given by*

$${}^{p:\sim} \mathbf{v}_{n,t} \equiv \frac{1}{n} \sum_{i=1}^n \mathbf{v}_t \left(\underbrace{\delta_t(\dots(\delta_t(i)))}_p \right) \quad {}^{p:\ddot{\sim}} \mathbf{v}_{n,t} \equiv \frac{1}{n} \sum_{i=1}^n \mathbf{v}_t \left(\underbrace{\delta_t(\dots(\delta_t(i)))}_{p-1} \right) \phi_n(i) \quad (16)$$

For example, the 1^{st} -weighted average is ${}^{1:\sim} \mathbf{v}_{n,t} \equiv \frac{1}{n} \sum_{i=1}^n \mathbf{v}_t(\delta_t(i))$ and represents the average idiosyncratic shock over all agents' observees. The 2^{nd} -weighted average is ${}^{2:\sim} \mathbf{v}_{n,t} \equiv \frac{1}{n} \sum_{i=1}^n \mathbf{v}_t(\delta_t(\delta_t(i)))$ and represents the average idiosyncratic shock over all agents' observees' observees.

The following proposition then demonstrates an equivalence between ${}^{p:\sim} \mathbf{v}_{n,t}$ and ${}^{p:\ddot{\sim}} \mathbf{v}_{n,t}$ as the number of agents approaches infinity and characterises the resultant distributions:

Proposition 1. *Suppose that $\mathbf{v}_t(i) \sim i.i.d. N(\mathbf{0}, \Sigma_{vv}) \forall i, t$ and assumptions 1 and 2 hold. Then in the limit (as $n \rightarrow \infty$):*

1. ${}^{p:\sim} \mathbf{v}_{n,t} \xrightarrow{d} {}^{p:\sim} \mathbf{v}_t \quad \forall p \geq 1$ where ${}^{p:\sim} \mathbf{v}_t \sim N(\mathbf{0}, \Sigma_{vv}^{\{p\}})$ and $\Sigma_{vv}^{\{p\}} = (1 - (1 - \zeta^*)^p) \Sigma_{vv}$
2. $Cov({}^{p:\sim} \mathbf{v}_t, {}^{r:\sim} \mathbf{v}_t) = \Sigma_{vv}^{\{p\}} \quad \forall p < r$

$$\mathcal{J} \cdot \overset{p:\sim}{\mathbf{v}}_{n,t} \xrightarrow{\mathcal{L}^2} \overset{p:\sim}{\mathbf{v}}_t \quad \forall p \geq 1$$

It may also be worth noting that the variance of the p^{th} -weighted average may also be expressed recursively as $\Sigma_{vv}^{\{p\}} = \zeta^* \Sigma_{vv} + (1 - \zeta^*) \Sigma_{vv}^{\{p-1\}}$. The following two corollaries then trivially follow:

Corollary 1. $\Sigma_{vv} \geq \dots \geq \Sigma_{vv}^{\{3\}} \geq \Sigma_{vv}^{\{2\}} \geq \Sigma_{vv}^{\{1\}}$ where \geq is in the sense that the difference between the two is a positive-definite matrix.

Corollary 2. $E \left[\overset{p:\sim}{\mathbf{v}}_t \mid \overset{1:\sim}{\mathbf{v}}_t = \mathbf{a} \right] = \mathbf{a} \quad \forall p \geq 2$

The first of these is a necessary component of approximating the full solution with a finite state vector (see section 3.6 below) and the latter is used when simulating the effects of network learning.

The increasing variance of higher-weighted averages captures the effect of what Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Saleh (2012) refer to as the p^{th} -order *interconnectivity* of the network, with each step back in the network more and more likely to arrive at the most highly connected agents. I avoid their nomenclature, however, and reserve the word “order” to refer to higher-order expectations.

Definition 5. The vector of **network shocks**, \mathbf{z}_t , is the infinite sequence of higher-weighted averages of agents’ idiosyncratic shocks:

$$\mathbf{z}_t \equiv \begin{bmatrix} \overset{1:\sim}{\mathbf{v}}_t \\ \overset{2:\sim}{\mathbf{v}}_t \\ \overset{3:\sim}{\mathbf{v}}_t \\ \overset{4:\sim}{\mathbf{v}}_t \\ \vdots \end{bmatrix} \sim N(\mathbf{0}, \Sigma_{zz}) \quad \Sigma_{zz} = \begin{bmatrix} \Sigma_{vv}^{\{1\}} & \Sigma_{vv}^{\{1\}} & \Sigma_{vv}^{\{1\}} & \Sigma_{vv}^{\{1\}} & \dots \\ \Sigma_{vv}^{\{1\}} & \Sigma_{vv}^{\{2\}} & \Sigma_{vv}^{\{2\}} & \Sigma_{vv}^{\{2\}} & \dots \\ \Sigma_{vv}^{\{1\}} & \Sigma_{vv}^{\{2\}} & \Sigma_{vv}^{\{3\}} & \Sigma_{vv}^{\{3\}} & \dots \\ \Sigma_{vv}^{\{1\}} & \Sigma_{vv}^{\{2\}} & \Sigma_{vv}^{\{3\}} & \Sigma_{vv}^{\{4\}} & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad (17)$$

Including these higher weighted averages is necessary because of the recursive nature of agents’ learning through the Kalman filter: it will be shown below that $\overset{1:\sim}{E}_t[X_t]$ is a function of $\overset{1:\sim}{\mathbf{v}}_t$ and $\overset{2:\sim}{E}_{t-1}[X_{t-1}]$, while $\overset{2:\sim}{E}_t[X_t]$ is a function of $\overset{2:\sim}{\mathbf{v}}_t$ and $\overset{3:\sim}{E}_{t-1}[X_{t-1}]$, etc.

3.5 Social learning over an opaque, irregular network

I am now in a position to present the main result of this paper.

Theorem 1. Given the broad setting described above and assumptions 1 and 2, the hierarchy of agents’ aggregate expectations will obey the following VARMA(1,1) law of motion:

$$X_t \equiv \begin{bmatrix} \mathbf{x}_t \\ \overline{E}_t[X_t] \\ \overset{1:\sim}{E}_t[X_t] \\ \overset{2:\sim}{E}_t[X_t] \\ \vdots \end{bmatrix} = F X_{t-1} + G_1 \mathbf{u}_t + G_2 \mathbf{z}_t + G_3 \mathbf{e}_t + G_4 \mathbf{z}_{t-1}$$

$$\begin{aligned} \text{where } \overline{E}_t[\cdot] &= \int_0^1 E_t(i) [\cdot] di \\ \overset{1:\sim}{E}_t[\cdot] &= \int_0^1 E_t(\delta_t(i)) [\cdot] di \\ \overset{2:\sim}{E}_t[\cdot] &= \int_0^1 E_t(\delta_t(\delta_t(i))) [\cdot] di \\ &\vdots \end{aligned}$$

and with each successively higher-weighted average expectation having a smaller effect on the simple-average expectation.

Although the complete derivation is provided in the appendix, an outline of the agents' learning process may be of interest. To begin, I define the general notation that $\theta_{t|q}^{\text{err}}(i)$ represents the *error* in agent i 's period- q expectation regarding θ_t . In particular, the following will be used:

$$\begin{aligned} \mathbf{s}_{t|t-1}^{\text{err}}(i) &\equiv \mathbf{s}_t(i) - E_{t-1}(i)[\mathbf{s}_t(i)] && : \text{signal innovation} \\ X_{t|t-1}^{\text{err}}(i) &\equiv X_t - E_{t-1}(i)[X_t] && : \text{prior expectation error} \\ X_{t|t}^{\text{err}}(i) &\equiv X_t - E_t(i)[X_t] && : \text{contemporaneous expectation error} \end{aligned}$$

The filter

As with a standard Kalman filter, the Kalman gain is calculated as:

$$K_t = \text{Cov}(X_t, \mathbf{s}_{t|t-1}^{\text{err}}(i)) \left[\text{Var}(\mathbf{s}_{t|t-1}^{\text{err}}(i)) \right]^{-1} \quad (18)$$

where $\mathbf{s}_{t|t-1}^{\text{err}}(i)$ is the agent's signal innovation (the portion of their signal that was not forecastable). Under a classic filtering problem with no network learning, the agents' signal innovation is a function of their expectation error from the previous period and current period shocks:

$$\mathbf{s}_{t|t-1}^{\text{err}}(i) = M_1 X_{t-1|t-1}^{\text{err}}(i) + N_1 \mathbf{u}_t + N_2 \mathbf{v}_t(i) + N_3 \mathbf{e}_t$$

In contrast, with network learning the signal innovation is expressed as follows (the additional terms are shown in red):

$$\begin{aligned} \mathbf{s}_{t|t-1}^{\text{err}}(i) &= M_1 X_{t-1|t-1}^{\text{err}}(i) + M_2 X_{t-1|t-1}^{\text{err}}(\delta_{t-1}(i)) + M_3 X_{t-1} \\ &\quad + N_1 \mathbf{u}_t + N_2 \mathbf{v}_t(i) + N_3 \mathbf{e}_t + N_4 \mathbf{v}_{t-1}(\delta_{t-1}(i)) + N_5 \mathbf{z}_{t-1} \end{aligned} \quad (19)$$

Note that innovation in i 's signal includes not only a term in their own previous-period expectation error but also a term in their *observee*'s expectation error. As such, both the covariance and variance terms in the Kalman gain (18) will therefore include terms in both the variance of i 's expectation error, $V_{t-1|t-1} \equiv E[X_{t-1|t-1}^{\text{err}}(i) X_{t-1|t-1}^{\text{err}}(i)']$, and the *covariance* between any two agents' errors, $W_{t-1|t-1} \equiv E[X_{t-1|t-1}^{\text{err}}(i) X_{t-1|t-1}^{\text{err}}(j)']$.

The variance in agents' own expectation errors then updates in via the familiar Riccati equation, but since the variance-covariance of $\mathbf{s}_{t|t-1}^{\text{err}}(i)$ includes terms in $W_{t-1|t-1}$, a corresponding expression must also be found for updating the covariance between agents' errors.

The law of motion

The law of motion starts from the basic form of the Kalman filter:

$$E_t(i)[X_t] = F E_{t-1}(i)[X_{t-1}] + K_t \mathbf{s}_{t|t-1}^{\text{err}}(i)$$

Equation (19) is substituted in for $\mathbf{s}_{t|t-1}^{\text{err}}(i)$ and a simple average is taken to obtain $\bar{E}_t[X_t]$. Since the signal innovation includes a term in $E_{t-1}(\delta_{t-1}(i))[X_{t-1}]$ (from the observee's expectation error), taking the simple average over i turns this into a term in $\bar{E}_{t-1}[X_{t-1}]$, thereby introducing the need to also determine the (first) weighted-average expectation.

Taking the weighted average of the filter to obtain $\bar{E}_t[X_t]$ then produces a term in $\bar{E}_{t-1}[X_{t-1}]$, thus requiring that the 2nd-weighted average expectation be included. The 2nd-weighted average expectation subsequently produces a term in the 3rd-weighted average expectation, and so forth.

3.6 Working with a finite approximation

The full state vector of interest and, hence, the transition matrices in the law of motion and the filter variances in the Kalman filter are all of infinite dimension. The full solution therefore cannot be found in practice and must be approximated with a truncated state.

Proposition 2. *An arbitrarily accurate approximation of the full solution implied by theorem 1 may be obtained by defining cut-off on the number of higher orders of expectation, k^* , and the number of higher-weighted compound expectations, p^* ; and including all weights and all orders from zero up to these cut-offs.*

Note that the size of the state vector can still be very large even when operating with few state variables and quite low choices of k^* and p^* . Table 1 lists the sizes that emerge for a variety of choices.

m	k^*	No network (standard ICK)	With network learning ($p^* = 3$)
1	4	$5 \times 5 : 200 \text{ B}$	$121 \times 121 : 114.4 \text{ KB}$
1	6	$7 \times 7 : 392 \text{ B}$	$1093 \times 1093 : 9.1 \text{ MB}$
4	4	$16 \times 16 : 2.0 \text{ KB}$	$484 \times 484 : 1.8 \text{ MB}$
4	6	$28 \times 28 : 6.1 \text{ KB}$	$4372 \times 4372 : 145.8 \text{ MB}$

Table 1: Size (each) of F , U , V and W , assuming use of double-precision.

Given the size of the matrices involved, problems of *numerical instability* must be considered when implementing the model. When iterating a large system over many steps, round-off errors that necessarily occur with floating-point operations on computers can accumulate and magnify to the extent that the system does not converge. Such a problem is, regrettably, relatively common in the implementation of larger Kalman filters and typically first appears as a failure of symmetry or positive definiteness in the variance matrices of the Ricatti equation.

Arguably the most robust (to roundoff error) implementations of Kalman filters are those that factor the relevant variance matrices, with a modified Cholesky decomposition (a “UD decomposition”) the most commonly used.²⁶ Using this technique for a regular Kalman filter, the algorithm for implementing the temporal update of the filter (from $V_{t-1|t-1}$ to $V_{t|t-1}$) was developed by Thornton (1976) and that for the observational update (from $V_{t|t-1}$ to $V_{t|t}$) by Bierman (1977).

Unfortunately, although the model developed here is amenable to use of the Thornton temporal update, the Bierman observational update algorithm is not applicable. This is because the inclusion of social signals introduces the need to consider the covariance of agents’ expectation errors so that, when calculating the Kalman gain, the covariance between the state (X_t) and the signal innovation ($\mathbf{s}_{t|t-1}^{\text{err}}(i)$) can no longer be expressed in the form

$$\text{Cov} \left(X_t, \mathbf{s}_{t|t-1}^{\text{err}}(i) \right) = V_{t|t-1} H$$

which is required for Bierman’s factorisation. A successful UD implementation of the current model would therefore require the derivation of a new algorithm in the style of Bierman that accounted for the more complex structure of the Kalman gain found here. This is left for future research.

²⁶A UD decomposition breaks V into UDU' with U unit upper triangular (i.e. with ones on the leading diagonal) and D diagonal. By working exclusively on these component matrices, the implied variance matrices remain well defined.

4 An illustrative example

I here present a simplified example to illustrate some of the results that emerge from adding network learning to a setting of strategic complementarity. Key simplifying assumptions include:

- A univariate underlying state
- No public signals
- No lagged signals (except the social signal through the network)
- Private signals serve an information role only
- Agents are myopic, in that they optimise on a period-by-period basis

Section 4.4 below extends this model to inclusion of a lagged public signal, while section 4.5 illustrates how to apply the model to dynamic settings where agents are forward looking in their decision rule.

4.1 The simplified model

There exists only a single hidden state that follows an AR(1) process and about which agents each observe a single, unbiased private signal

$$\begin{aligned} x_t &= \rho x_{t-1} + u_t & u_t &\sim N(0, \sigma_u^2) \\ s_t^p(i) &= x_t + v_t(i) & v_t(i) &\sim N(0, \sigma_v^2) \end{aligned}$$

with u_t and $v_t(i)$ being fully independent for all i and t . Agents face quadratic losses from mismatch between their action, a single hidden state and the average action of others:²⁷

$$u_i(\mathbf{g}_t, x_t) = -(1 - \beta) [(g_t(i) - x_t)^2] - \beta [(g_t(i) - \bar{g}_t)^2] \quad \beta \in (0, 1)$$

With agents maximising their expected payoff without explicitly knowing the state or the average action that other agents will take, their optimal action is given by

$$g_t(i) = (1 - \beta) E_t(i) [x_t] + \beta E_t(i) [\bar{g}_t]$$

Taking the simple average of this and repeatedly substituting it back in, we obtain

$$g_t(i) = (1 - \beta) \begin{bmatrix} 1 & \beta & \beta^2 & \dots \end{bmatrix} E_t(i) [\bar{x}_t^{(0:\infty)}]$$

With each agent observing the previous-period action of q competitors, theorem 1 then grants that the following laws of motion emerge:

$$\begin{aligned} \mathbf{x}_t &= \rho \mathbf{x}_{t-1} & + & \mathbf{u}_t \\ \bar{E}_t[X_t] &= B \mathbf{x}_{t-1} + C \bar{E}_{t-1}[X_{t-1}] + D \overset{1:\sim}{E}_{t-1}[X_{t-1}] + H \mathbf{u}_t \\ \overset{1:\sim}{E}_t[X_t] &= B \mathbf{x}_{t-1} + C \overset{1:\sim}{E}_{t-1}[X_{t-1}] + D \overset{2:\sim}{E}_{t-1}[X_{t-1}] + H \mathbf{u}_t + Q \overset{1:\sim}{v}_t \\ \overset{2:\sim}{E}_t[X_t] &= B \mathbf{x}_{t-1} + C \overset{2:\sim}{E}_{t-1}[X_{t-1}] + D \overset{3:\sim}{E}_{t-1}[X_{t-1}] + H \mathbf{u}_t + Q \overset{2:\sim}{v}_t \\ &\vdots \end{aligned}$$

²⁷This utility function is quite common in the network literature (see, e.g., [Calvó-Armengol and de Martí, 2007](#)). An alternative utility function described by [Morris and Shin \(2002\)](#) presents the strategic complementarity as being a zero-sum game, but produces the same optimal decision rule for individual agents (although not for a social planner).

$$\begin{aligned}
\text{where } B &= \mathbf{k}_p \rho & H &= \mathbf{k}_p \\
C &= F - BS_x - DT_{w_1} & Q &= q \mathbf{k}_p \\
D &= q \mathbf{k}_s \boldsymbol{\lambda}'_1
\end{aligned}$$

with \mathbf{k}_p being the Kalman gain applied to the private signal and \mathbf{k}_s the Kalman gain applied to each social signal, while S_x and T_{w_1} select x_t and $\overset{1:\sim}{E}_t[X_t]$ respectively from X_t . The transition matrix for the full state therefore takes the following form:

$$F = \begin{bmatrix}
\rho & 0 & 0 & 0 & \cdots \\
B & C & D & 0 & \\
B & 0 & C & D & \\
B & 0 & 0 & C & \ddots \\
\vdots & & & & \ddots
\end{bmatrix}$$

The p^{th} -weighted expectation is given by:

$$\overset{p:\sim}{E}_t[X_t] = (\mathbf{k}_p \rho S_x + (F - \mathbf{k}_p \rho S_x) T_{w_p}) X_{t-1} + q \mathbf{k}_s \boldsymbol{\lambda}'_1 (T_{w_{p+1}} - T_{w_1} T_{w_p}) X_{t-1} + shocks$$

When considering the expectations of agents p levels deep in the network, the component derived from consideration of agents $p+1$ levels deep is captured in the term $q \mathbf{k}_s \boldsymbol{\lambda}'_1 (T_{w_{p+1}} - T_{w_1} T_{w_p}) X_{t-1}$.

In order to simulate the model, it is necessary to form a finite approximation of the solution. For example, if we impose that $k^* = 2$ and $p^* = 2$, then the state vector will have eight elements:

$$X_t = \mathbb{E}_t^{(0:2)}[x_t] = \begin{bmatrix} x_t & \overline{E}_t[x_t] & \overline{E}_t[\overline{E}_t[x_t]] & \overline{E}_t[\overset{1:\sim}{E}_t[x_t]] & \overset{1:\sim}{E}_t[x_t] & \overset{1:\sim}{E}_t[\overline{E}_t[x_t]] & \overset{1:\sim}{E}_t[\overset{1:\sim}{E}_t[x_t]] \end{bmatrix}'$$

It can be readily shown that $Cov(\overline{E}_t[X_t], \mathbf{s}_{t|t-1}^{\text{err}}(i)) = Cov(\overset{p:\sim}{E}_t[X_t], \mathbf{s}_{t|t-1}^{\text{err}}(i)) \forall p$, which implies the following repetitive structures for the Kalman gains:

$$\mathbf{k}_p = \begin{bmatrix} \kappa_1^p \\ \kappa_2^p \\ \kappa_3^p \\ \kappa_3^p \\ \kappa_2^p \\ \kappa_3^p \\ \kappa_3^p \end{bmatrix} \quad \mathbf{k}_s = \begin{bmatrix} \kappa_1^s \\ \kappa_2^s \\ \kappa_3^s \\ \kappa_3^s \\ \kappa_2^s \\ \kappa_3^s \\ \kappa_3^s \end{bmatrix}$$

and the following coefficients in the law of motion:

$$B = \begin{bmatrix} \kappa_1^p \rho \\ \kappa_2^p \rho \\ \kappa_2^p \rho \end{bmatrix} \quad C = \begin{bmatrix} (1 - \kappa_1^p) \rho & 0 & -q \kappa_1^s (1 - \beta) \\ (\kappa_2^p - \kappa_1^p) \rho & (1 - \kappa_1^p) \rho & q (\kappa_1^s - \kappa_2^s) (1 - \beta) \\ (\kappa_2^p - \kappa_1^p) \rho & 0 & (1 - \kappa_1^p) \rho + q (\kappa_1^s - \kappa_2^s) (1 - \beta) \end{bmatrix} \quad D = \begin{bmatrix} q \kappa_1^s (1 - \beta) & 0 & 0 \\ q \kappa_2^s (1 - \beta) & 0 & 0 \\ q \kappa_2^s (1 - \beta) & 0 & 0 \end{bmatrix}$$

For the simulations that follow, I suppose the following parameters:

4.2 Aggregate beliefs following a shock to the underlying state

Figure 5 plots impulse responses for the hierarchy of simple-average expectations²⁸ following a one standard deviation shock to the hidden state, both with and without network learning.

²⁸So $k = 0$ denotes the time path of x_t , $k = 1$ the time path of $\overline{E}_t[x_t]$, $k = 2$ the time path of $\overline{E}_t[\overline{E}_t[x_t]]$ and so on.

Parameter	Value	Description
β	0.5	The relative importance of strategic complementarity
ρ	0.6	The persistence of shocks to the hidden state
σ_v^2/σ_u^2	5.0	The relative innovation variance
ζ^*	0.1	The degree of irregularity in the network

Table 2: Baseline parameterisation

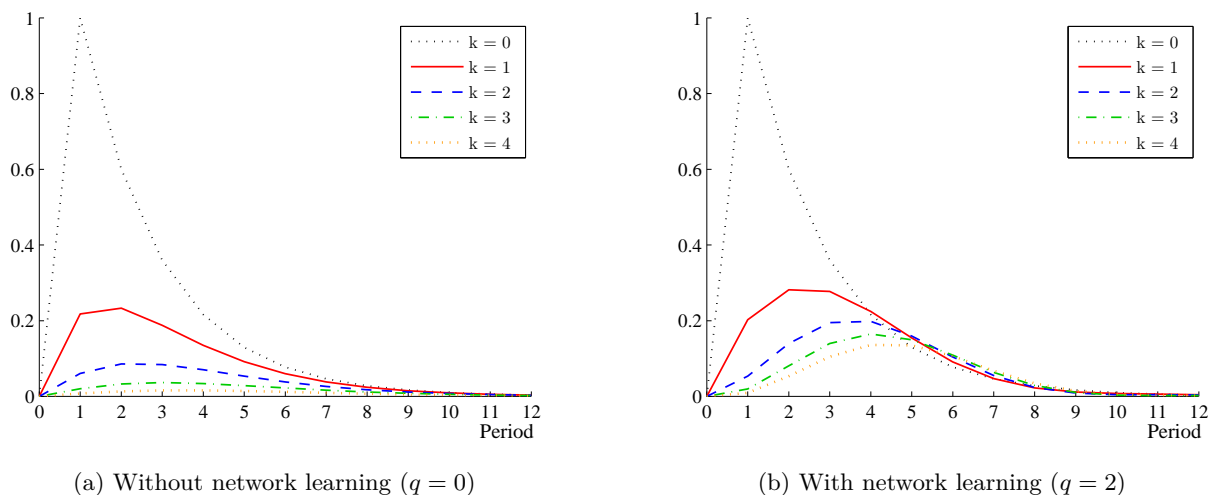


Figure 5: The hierarchy of simple-average expectations ($\bar{x}_{t|t}^{(0:\infty)}$) following a one standard deviation shock to the underlying state.

Figure 5a shows a standard scenario in the dispersed information literature, with agents only having access to their private signals. Although all agents' signals are unbiased, the presence of noise ensures that they attribute some of their signal to idiosyncratic factors, so the average expectation responds by less than the truth. Since each agent knows this (common knowledge of rationality), each successive order of expectation responds by less than its predecessor. All orders of expectation remain below the underlying state, so the average expectation error ($\bar{x}_{t|t}^{\text{err}} \equiv x_t - \bar{E}_t[x_t]$) remains strictly positive. The hierarchy of beliefs subsequently decays back to zero with the underlying shock.

Figure 5b then plots the equivalent impulse responses when, in addition to observing their private signals, each agent observes the previous-period action of two competitors. On impact, there is very little difference because social signals are received with a lag (the observation of competitors' actions having been zero in the pre-impact period lowers the beliefs fractionally). In the near term, agents' average expectations are improved relative to the no-network case, with observations of their peers' actions reinforcing their own private signals that an aggregate shock has occurred. In the longer term, however, as the underlying state decays back to zero, agents' beliefs tend to *overshoot* the truth, so the average expectation error ($\bar{x}_{t|t}^{\text{err}}$) becomes negative.

This is herding in the broad sense of Banerjee (1992), but with an amplification from Morris and Shin (2002)-style strategic complementarity. First and most simply, by observing that their competitors' actions were high yesterday, agents infer that the state may be high today. As a result, they partially attribute their low private signals to idiosyncratic noise, consequently choosing a high action themselves. However, although there is no public signal available, by effectively assuming that their competitors all observe the same weighted average action, agents' social observations act as private signals about a public signal that they themselves cannot observe but which they assume is

seen by everybody else.²⁹ For any given agent, their social observation therefore acts as a coordination device for addressing their strategic complementarity concerns. When the underlying state is falling, this therefore acts as a kind of upward bias in social signals for signal extraction purposes.

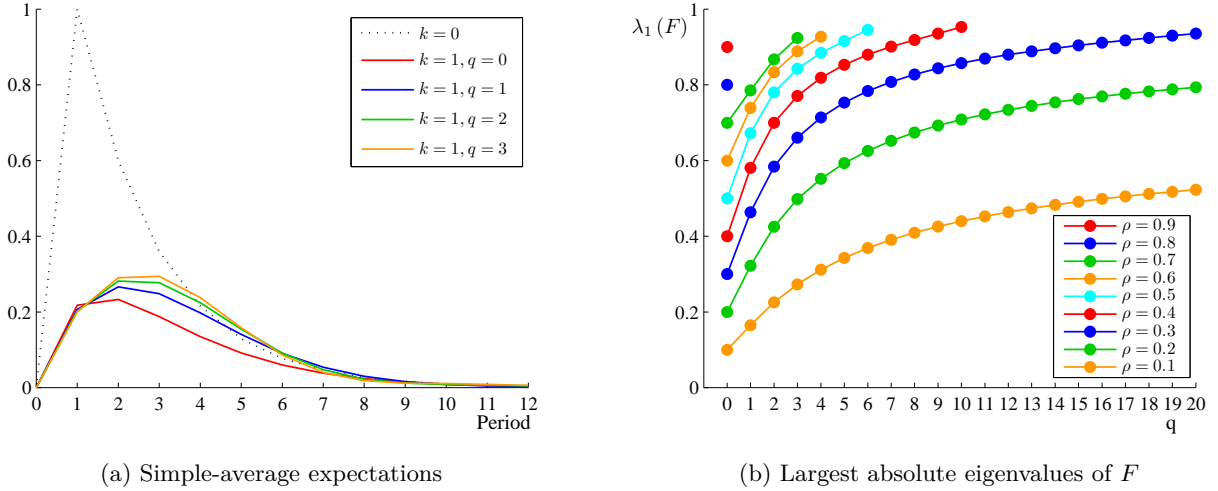


Figure 6: Varying the number of other agents observed (q)

Figure 6 illustrates the increase in persistence obtained when agents increase the number of competitors they each observe. Figure 6a first plots the truth ($k = 0$) and the first-order simple-average expectations ($k = 1$) for a variety of the number of competitors observed. Figure 6b presents a broader picture, showing the largest absolute eigenvalue of F when agents each observe q competitors for different degrees of persistence in the underlying state (ρ). When there is no network learning ($q = 0$) the full hierarchy of expectations exhibits the same persistence as the underlying state, but as the number of observees increases, the overall persistence of agents' higher-order beliefs rises.

This is **not** to say that the (simple) average expectation becomes arbitrarily persistent. The full hierarchy also includes weighted-average expectations and higher-order expectations of both and it is the higher-weighted, higher-order expectations that exhibit the greatest persistence. Nevertheless, the first-order simple-average expectation, and the hierarchy of simple-average expectations above it, do appear to exhibit greater persistence than the underlying shock in a form of rational herding.

Solutions are not currently able to be found for all combinations of q and ρ , as for more persistent systems (with either high ρ or high q), the filter becomes susceptible to numerical instability.³⁰ Nevertheless, figure 6b is suggestive of the following conjecture.

Conjecture 1. *Let $\lambda_1(A)$ denote the largest eigenvalue of the transition matrix for the underlying state and let $\lambda_1(F, q)$ be the largest absolute eigenvalue of the transition matrix for the full hierarchy of aggregate expectations when each agent observes q competitors. Then*

- $\lambda_1(F, 0) = \lambda_1(A)$
- $1 > \lambda_1(F, q) > \lambda_1(A) \quad \forall q \geq 1$

²⁹Strictly speaking, agents do not assume that their competitors observe a public signal. Rather, their Bayes-rational signal extraction problem is mathematically equivalent to making the assumption.

³⁰Numerical instability is the result of round-off errors due to the limitation of performing floating-point operations on a computer. In Kalman filters, it most often (and, indeed, here) enters predominantly via the Riccati equation updating the variance in expectation errors. The round-off errors cause these variance-covariance matrices to become non-symmetric or non-positive-definite (a mathematical impossibility), which causes the system to explode.

- $\lambda_1(F, q)$ is strictly increasing in q and approaches an asymptote at one

Proof of this conjecture is left for future work, although it is interesting to note that in contrast to work demonstrating long memory with directed networks (see, for example, [Schennach, 2013](#)), the autocorrelation function appears to be absolutely summable here.

Varying parameters

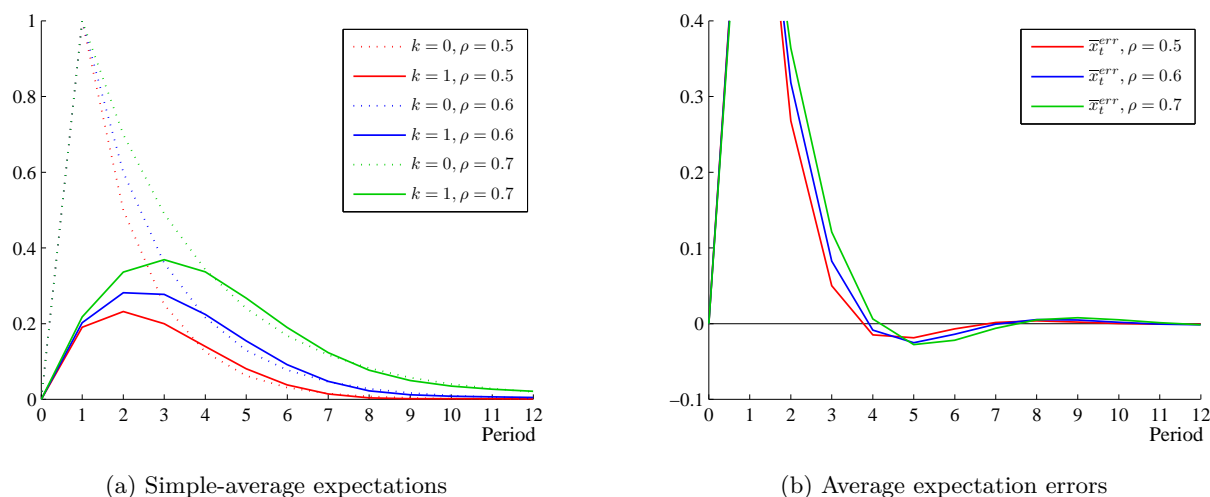


Figure 7: Varying underlying persistence (ρ)

Figure 7 shows the impulse responses of first-order simple-average expectations and the corresponding average expectation errors for different values of ρ . Larger values of ρ cause not only larger movements in average expectations, but renders the errors in those expectations larger for longer. In other words, the presence of network learning introduces a persistence multiplier effect so that the persistence of average beliefs increases by more than that of the state.

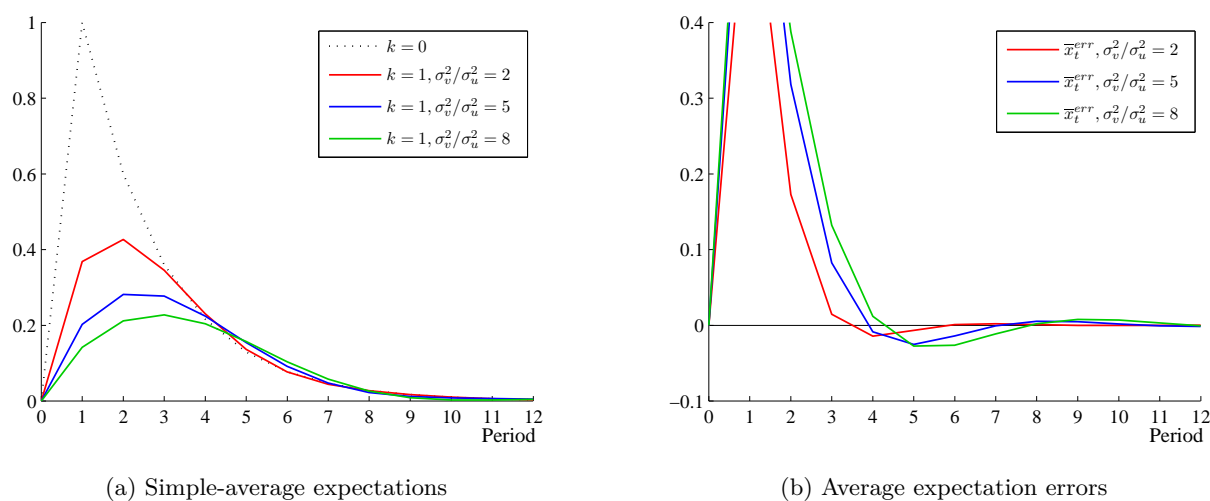


Figure 8: Varying the relative innovation variance (σ_v^2/σ_u^2)

Figure 8 then presents equivalent plots for a variety of values for σ_v^2/σ_u^2 . Lowering the signal-to-noise ratio of agents' private signals³¹ worsens the value of those signals, causing them to rely more

³¹That is, raising the relative variance of idiosyncratic shocks.

heavily on the social signals. This reduces agents' average performance shortly after a shock and produces a stronger overshoot.

4.3 Aggregate beliefs following a network shock

In addition to shocks to the underlying state, the irregularity of the observation network gives rise to the possibility of aggregate *network shocks*: a set of idiosyncratic shocks such that prominent agents happen to draw innovations in one direction (say, positive) while obscure agents draw innovations in the opposite direction. With a continuum of agents, the law of large numbers ensures that the simple average innovation is zero, but weighted averages (with weights given by the probability of being observed) will be non-zero.





Figure 9: The hierarchy of simple-average expectations ($\bar{x}_{t|t}^{(0:\infty)}$) following a one standard deviation network shock (a one standard deviation shock to $\tilde{v}_t^{1:}$ and the corresponding conditional expected value for higher-weighted averages) with agents each observing two competitors ($q = 2$).


Figure 9 plots the hierarchy of simple-average expectations regarding the hidden state following a one standard deviation network shock – strictly, a one standard deviation shock to $\tilde{v}_t^{1:}$ plus the corresponding (conditionally) expected value for higher weighted averages – when agents each observe two competitors ($q = 2$). Note that the underlying state remains at zero throughout. Unlike with a shock to the state, there is no movement in aggregate beliefs on impact because the law of large numbers does apply: all agents receive the same social signal from the pre-impact period and movements in the expectations of prominent and obscure agents balance out. In the second period, the average expectation rises as people observe the positive movement in prominent agents' actions from period one and largely ignore the opposite movements by obscure agents. Consequently in period two, despite the average private signal being zero, not just prominent agents but *all* agents, on average, choose positive actions. Aggregate beliefs then gradually decay back to zero as agents continue to receive average private signals of zero but continue to place weight on the previous actions of others.

Overall, the scale of movements in average beliefs is roughly one order of magnitude smaller than those following a true shock to the underlying state. This scale is controlled by the relative variance of the network shocks (recall that $Var(\tilde{v}_t^{1:}) = \zeta^* \sigma_v^2$), but also by the persistence of underlying state shocks and the degree of strategic complementarity, as shown in figure 10.


Figure 10a first shows the IRFs of simple-average expectations for different degrees of irregularity




(a) Varying the degree of network irregularity (ζ^*)



(b) Varying the relative innovation variance (σ_v^2/σ_u^2)



(c) Varying underlying persistence (ρ)



(d) Varying strategic complementarity (β)

Figure 10: IRFs of simple-average expectations ($\bar{E}_t[x_t]$) following a network shock for a variety of parameters ($q = 2$ for all).

in the observation network. At one extreme ($\zeta^* \rightarrow 0$), the distribution of links is sufficiently close to uniform that the law of large numbers applies, meaning that network shocks have no effect. At the other extreme, as the probability of being observed approaches unity for a single agent and zero for everybody else ($\zeta^* \rightarrow 1$), that sole agent's idiosyncratic shocks come to play a significant role in shaping average beliefs. Although varying ζ^* changes the magnitude of any movement in agents' average expectations, the profile and persistence of that movement is unchanged.

Figure 10b next shows the effect of network shocks when varying the relative variance of agents' idiosyncratic shocks. As with increasing ζ^* , an increase in σ_v^2/σ_u^2 increases the magnitude of the average expectation's response, but in addition, as seen for shocks to the underlying state above, the increased uncertainty also increases the persistence of the shocks' effects. The profile of the responses also changes, with IRFs become more predominantly hump-shaped for higher relative variances.

Figure 10c then considers the different responses to a network shock for various degrees of persistence in the underlying state. A more persistent hidden state causes a larger and more hump-shaped response to a network shock: (lagged) social signals are more informative about the underlying state when that state is more persistent (higher ρ). Despite the change in profile, the persistence of average beliefs following a network shock does not appear to change with persistence in the underlying state.

Figure 10d finally plots the responses for different degrees of strategic complementarity. Greater strategic complementarity (higher β) dampens the response in period 2 and produces a more pronounced later peak and more persistent response, as agents seek greater confirmation that other agents are adjusting their actions before acting themselves and this flows into the collective signal extraction problem.

4.4 Adding a (lagged) public signal

It is straightforward to add a public signal to this simplified setting. I consider two cases: one in which the signal is based on the hierarchy of simple-average expectations only and one based on the entire hierarchy (i.e. including weighted-average expectations).

$$\text{Scenario 1: } s_t^{pub} = \mathbf{1}' \bar{x}_{t-1|t-1}^{(0:\infty)} + e_t \quad (20)$$

$$\text{Scenario 2: } s_t^{pub} = \mathbf{1}' X_{t-1} + e_t \quad (21)$$

Figure 11 then plots the IRFs of first-order simple-average expectations following a shock to the underlying state, a network shock, or a noise shock in the public signal. A public signal reduces

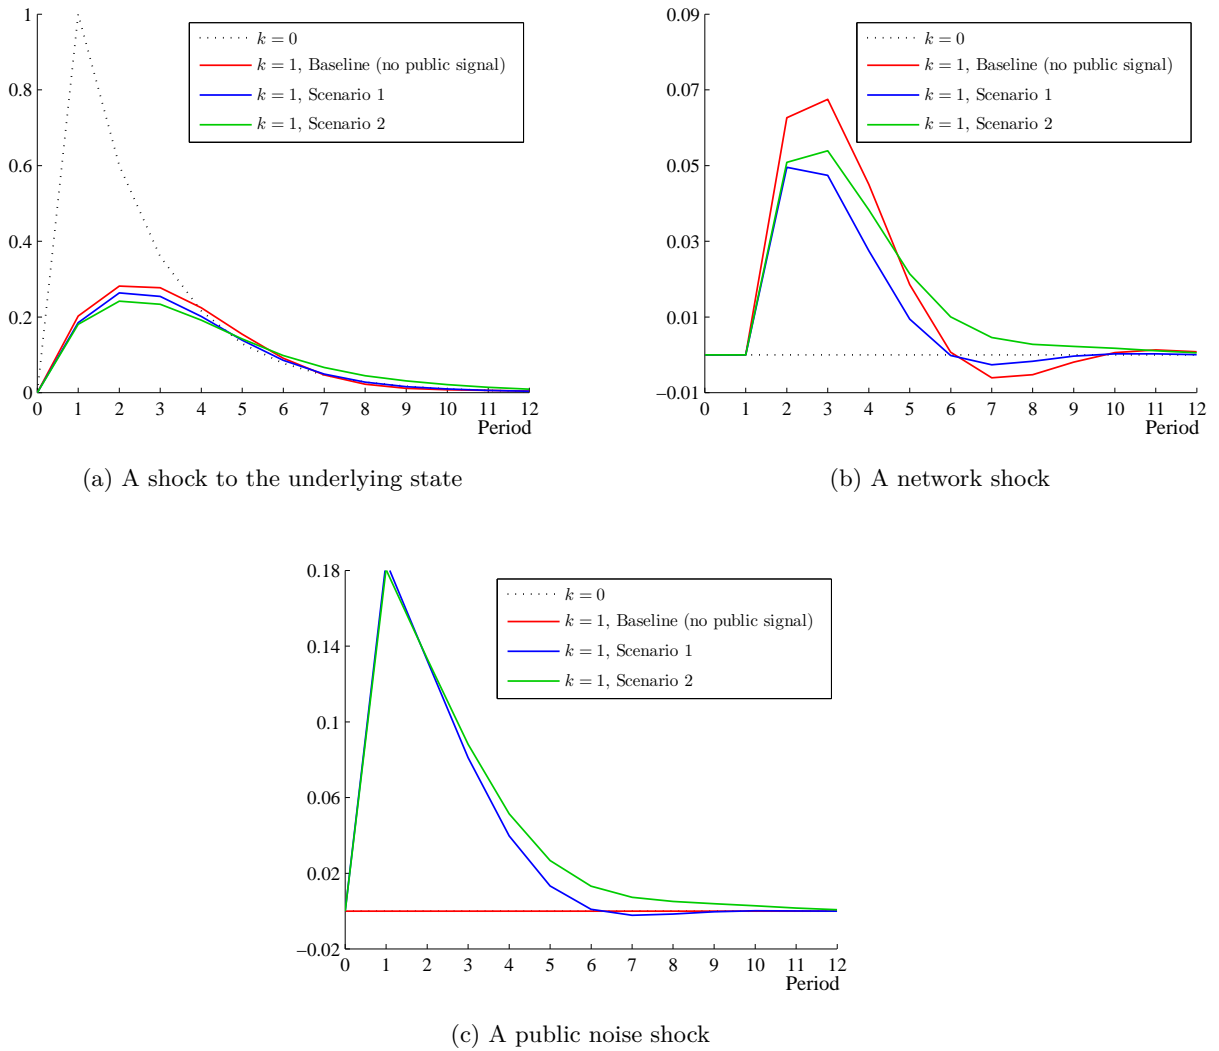


Figure 11: IRFs of simple-average expectations ($\bar{E}_t[x_t]$) when agents have access to a noisy, lagged public signal ($q = 2$ for all).

the near-term response of simple-average expectations following both shocks to the underlying state

and network shocks. For a public signal based only on the hierarchy of simple-average expectations, the overshoot following a shock to the underlying state is reduced (and removed entirely if the signal is perfect). A public signal that is influenced by weighted-average expectations, however, can increase the persistence of simple-average expectations.

4.5 Dynamic actions

The example above presented only a repeated static problem, with agents' optimal actions in period t being only a function of period t variables. This can be extended to a dynamic setting, however. As a simple example, suppose that agents' private signals are given by:

$$\mathbf{s}_t^p(i) = B\mathbf{x}_t + Q\mathbf{v}_t(i)$$

and that the linearised first-order conditions of agents' optimisation problems are given by:

$$g_t(i) = \boldsymbol{\alpha}' \mathbf{s}_t^p(i) + \boldsymbol{\eta}'_x E_t(i) [X_t] + \eta_y E_t(i) [\bar{g}_t] + \eta_z E_t(i) [\bar{g}_{t+1}]$$

so that agents' period- t action is a function of their expectation of the average action in both period t and period $t + 1$. I show in the appendix that this may be expressed as

$$g_t(i) = \underbrace{(\boldsymbol{\eta}'_x + \eta_y \mathbf{a}' + \eta_z \mathbf{a}' F) E_t(i) [X_t]}_{\lambda'_1} + \underbrace{\boldsymbol{\alpha}' B \mathbf{x}_t}_{\lambda'_2} + \underbrace{\boldsymbol{\alpha}' Q \mathbf{v}_t(i)}_{\lambda'_3}$$

where

$$\mathbf{a}' \equiv (\boldsymbol{\alpha}' B S + \boldsymbol{\eta}'_x T_s) (I - \eta_y T_s)^{-1} \left(I - \eta_z F T_s (I - \eta_y T_s)^{-1} \right)^{-1}$$

which is clearly in the form of equation (6). Further extensions to consideration of an infinite sum of forward-looking variables in agents' decision rule are straightforward.

5 Conclusion

This paper has introduced and solved a model of social learning over an exogenous directed network with a continuum of agents that satisfies the three requirements that (a) agents are rational; (b) agents act simultaneously and repeatedly over many periods; and (c) agents' optimal decisions include consideration of strategic complementarity. To avoid the curse of dimensionality that ordinarily prevents analysis of large networks, I introduce the idea of *network opacity* – that agents know who they observe, but not who anybody else observes. Instead, I suppose that agents know only the (common) distribution from which those observees are drawn.

This assumption grants that an arbitrarily accurate simulation may be performed by selecting a cut-off, k^* , on the number of higher-order expectations and a cut-off, p^* , on the number of compound expectations to consider. The first of these arises from the standard assumption that agents place decreasing weight on higher-order expectations. The second emerges from (a) the opacity of the network (so that agents are interested in a sequence of weighted-average expectations); (b) the recursive nature of the Kalman filter (so that each weighted-average expectation depends on the next-higher weighted average from the previous period); and (c) the AR process of the underlying state (so that older shocks are of decreasing importance to the current state).

Theorem 1 demonstrates that when the underlying state follows a VAR(1) process, the full hierarchy of relevant aggregate expectations will follow a VARMA(1,1) process with *network shocks* – weighted sums of agents' idiosyncratic shocks – entering both contemporaneously and with a lag.

A number of broad consequences of the model emerge directly from theorem 1. First, it is possible to simulate the effects of network learning without having to simulate the network explicitly: the network shocks together represent a sufficient statistic for the effect of the network on agents' aggregate beliefs. This makes the model particularly amenable to nesting within broad general equilibrium models of the economy.

Second, impulse responses of average expectations following shocks to the underlying state will exhibit greater persistence than the state itself, increasing in the number of agents observed. This is a form of rational herding behaviour that combines the herding exhibited in both Banerjee (1992), where agents observe others' actions, but have no strategic motive; and Morris and Shin (2002), where agents have a strategic motive, but do not observe others' actions.

Third, when the network is asymptotically irregular (i.e. has a distribution of links that is sufficiently far from uniform), mean zero idiosyncratic shocks do not wash out in aggregation, thereby leading to a network-based source of aggregate volatility, independent of "true" aggregate shocks to the hidden state. The scale of this additional volatility depends on the degree of irregularity in the network, which is captured simply in a single parameter: ζ^* .

Finally, because of the recursive nature of agents' learning, the aggregate effects of idiosyncratic shocks are *persistent*, even though the shocks themselves are entirely transitory.

The model would appear to be applicable to a variety of problems in macroeconomic research, including, for example, firms' price-setting decisions, labour search-and-matching models and asset pricing problems.

References

- ACEMOGLU, D., V. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEH (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80(5), 1977–2016.
- ACEMOGLU, D., M. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): "Bayesian Learning in Social Networks," *Review of Economic Studies*, 78(4), 1201–1236.
- ACEMOGLU, D., AND A. OZDAGLAR (2011): "Opinion Dynamics and Learning in Social Networks," *Dynamic Games and Applications*, 1, 3–49.
- ALBERT, R., AND A.-L. BARABASI (2002): "Statistical mechanics of complex networks," *Reviews of Modern Physics*, 74(1), 47–97.
- ANGELETOS, G., AND J. LA'O (2009): "Incomplete Information, Higher-Order Beliefs and Price Inertia," *NBER Working Papers*, No. 15003.
- (2010): "Noisy Business Cycles," in *NBER Macroeconomics Annual 2009*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, vol. 24, chap. 5, pp. 319–378. University of Chicago Press, Chicago.
- BANERJEE, A. (1992): "A simple model of herd behavior," *The Quarterly Journal of Economics*, 107(3), 797–817.
- BANERJEE, A., AND D. FUDENBERG (2004): "Word-of-mouth learning," *Games and Economic Behavior*, 46(1), 1–22.

- BIERMAN, G. J. (1977): *Factorization Methods for Discrete Sequential Estimation*, Mathematics in Science and Engineering v. 128. Academic Press, New York.
- CALVÓ-ARMENGOL, A., AND J. DE MARTÍ (2007): “Communication Networks: Knowledge and Decisions,” *American Economic Review*, 97(2), 86–91.
- CALVO, G. (1983): “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12(3), 383–398.
- CLAUSET, A., C. R. SHALIZI, AND M. E. J. NEWMAN (2009): “Power-Law Distributions in Empirical Data,” *SIAM Review*, 51(4), 661–703.
- DEGROOT, M. (1974): “Reaching a Consensus,” *Journal of the American Statistical Association*, 69(345), 118–121.
- DEMARZO, P., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence and Unidimensional Opinions,” *The Quarterly Journal of Economics*, 118(3), 909–968.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- GABAIX, X. (2011): “The Granular Origins of Aggregate Volatility,” *Econometrica*, 79(3), 733–772.
- GALE, D., AND S. KARIV (2003): “Bayesian learning in social networks,” *Games and Economic Behavior*, 45(2), 329–346.
- GOLUB, B., AND M. JACKSON (2010): “Naïve Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2(1), 112–149.
- GOYAL, S. (2007): *Connections: an introduction to the economics of networks*. Princeton University Press, Princeton.
- GRAHAM, L. (2011a): “Individual rationality, model-consistent expectations and learning,” CDMA Working Paper Series 201112, Centre for Dynamic Macroeconomic Analysis.
- (2011b): “Learning, information and heterogeneity,” CDMA Working Paper Series 201113, Centre for Dynamic Macroeconomic Analysis.
- GRAHAM, L., AND S. WRIGHT (2010): “Information, heterogeneity and market incompleteness,” *Journal of Monetary Economics*, 57(2), 164–174.
- HAMILTON, J. (1994): *Time Series Analysis*. Princeton University Press, 1 edn.
- JACKSON, M. (2008): *Social and Economic Networks*. Princeton University Press, Princeton.
- JACKSON, M., AND B. ROGERS (2007): “Meeting Strangers and Friends of Friends: How Random Are Social Networks?,” *American Economic Review*, 97(3), 890–915.
- KEMENY, J., AND J. L. SNELL (1960): *Finite Markov Chains*. Van Nostrand, Princeton.
- KEYNES, J. M. (1936): *The General Theory of Employment, Interest and Money*. Palgrave Macmillan.

- KOHLHAS, A. (2013): “Learning-by-sharing: Monetary Policy and the Information Content of Aggregate Variables,” Job market paper, Cambridge.
- LEE, I. H. (1993): “On the Convergence of Informational Cascades,” *Journal of Economic Theory*, 61(2), 395–411.
- LJUNGQVIST, L., AND T. SARGENT (2004): *Recursive Macroeconomic Theory*. MIT Press, 2 edn.
- LORENZONI, G. (2009): “A Theory of Demand Shocks,” *American Economic Review*, 99(5), 2055–2084.
- LUCAS, R. (1972): “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, 4(2), 103–124.
- (1976): “Econometric Policy Evaluation: A Critique,” in *The Phillips Curve and Labor Markets*, ed. by K. Brunner, and A. Meltzer, pp. 19–46. American Elsevier, New York.
- MANKIW, G., AND R. REIS (2002): “Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve,” *The Quarterly Journal of Economics*, 117(4), 1295–1328.
- (2006): “Pervasive Stickiness,” *American Economic Review*, 96(2), 164–169.
- (2007): “Sticky Information in General Equilibrium,” *Journal of the European Economic Association*, 5(2-3), 603–613.
- MELOSI, L. (2014): “Estimating Models with Dispersed Information,” *American Economic Journal: Macroeconomics*, 6, 1–31.
- MORRIS, S., AND H. S. SHIN (2002): “Social Value of Public Information,” *American Economic Review*, 92(5), 1521–1534.
- MUELLER-FRANK, M. (2013): “A General Framework for Rational Learning in Social Networks,” *Theoretical Economics*, Forthcoming(1), 1–40.
- NIMARK, K. (2008): “Dynamic pricing and imperfect common knowledge,” *Journal of Monetary Economics*, 55(2), 365–382.
- (2011a): “Dynamic Higher Order Expectations,” Universitat Pompeu Fabra Economics Working Papers No 1118.
- (2011b): “A low dimensional Kalman Filter for systems with lagged observables,” mimeo.
- PHELPS, E. (1984): “The trouble with ‘rational expectations’ and the problem of inflation stabilization,” in *Individual Forecasting and Aggregate Outcomes: ‘Rational Expectations’ Examined*, ed. by R. Frydman, and E. Phelps, chap. 2, pp. 31–46. Cambridge University Press.
- SCHENNACH, S. (2013): “Long memory via networking,” CeMMAP working papers CWP13/13, Centre for Microdata Methods and Practice, Institute for Fiscal Studies, UCL.
- SIMS, C. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3), 665–690.
- SMITH, L., AND P. SØRENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68(2), 371–398.

- THORNTON, C. L. (1976): “Triangular Covariance Factorizations for Kalman Filtering,” Ph.D. thesis, University of California at Los Angeles.
- TOWNSEND, R. M. (1983): “Forecasting the Forecasts of Others,” *Journal of Political Economy*, 91(4), 546–587.
- WOODFORD, M. (2003): “Imperfect Common Knowledge and the Effects of Monetary Policy,” in *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford. Princeton University Press.