

Kent Academic Repository

Full text document (pdf)

Citation for published version

Barde, Sylvain (2014) Back to the Future: Economic Self-Organisation and Maximum Entropy Prediction. *Computational Economics*, 45 (2). pp. 337-358. ISSN 0927-7099.

DOI

<http://doi.org/10.1007/s10614-014-9422-2>

Link to record in KAR

<http://kar.kent.ac.uk/41872/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Back to the future: economic self-organisation and maximum entropy prediction

Sylvain Barde

Received: date / Accepted: date

Abstract This paper shows that signal restoration methodology is appropriate for predicting the equilibrium state of certain economic systems. A formal justification for this is provided by proving the existence of finite improvement paths in object allocation problems under weak assumptions on preferences, linking any initial condition to a Nash equilibrium. Because a finite improvement path is made up of a sequence of systematic best-responses, backwards movement from the equilibrium back to the initial condition can be treated like the realisation of a noise process. This underpins the use of signal restoration to predict the equilibrium from the initial condition, and an illustration is provided through an application of Maximum Entropy signal restoration to the Schelling model of segregation.

Keywords Information entropy · self-organisation · potential function · Schelling segregation

JEL C02 · C11 · C63 · D80

Copyright information

This is a self-archived copy of the author's accepted manuscript, allowed under Springer's self-archiving policy. The final version of the article is published in *Computational Economics*, and is available at Springer via:

<http://dx.doi.org/10.1007/s10614-014-9422-2>

Sylvain Barde
School of Economics, Keynes College, University of Kent, Canterbury, CT2 7NP, United Kingdom
Tel.: + 44 1 227 82 40 92
Fax: + 44 1 227 82 78 50
E-mail: s.barde@kent.ac.uk

Affiliate, Observatoire Français des Conjonctures Economiques

1 Introduction

The presence of self-organisation in economic systems is one of the earliest features identified in economic thought. This phenomenon, referred to initially as the “invisible hand” by Adam Smith, then as “spontaneous order” following the work of Hayek (1945), is now often referred to as “emergence” by the agent-based computational economic literature, for instance in Tesfatsion (2006). Two common elements in all three of these interpretations are the notion that (a) economic mechanisms are not planned *ex ante* but appear as the result of interactions between agents, and (b) these mechanisms maintain the organisation of economic activity over time. To some extent, this is analogous to the Schrödinger (1967) definition of living organisms as entities able to maintain or increase their internal order, thus maintaining or reducing their internal entropy.

This paper argues that a consequence of such self-organisation is that under the right conditions, the evolution over time of economic systems should be predictable by signal restoration algorithms. The intuition for this is simple: signal restoration algorithms are designed to reconstruct signals that have been degraded by noise or by distortion. The standard interpretation of this, in a system where entropy (i.e. disorder) increases with time is that the non-degraded signal existed in the past. However, in a self-organising system, the arrow of time is reversed and the non-degraded signal exists in the future. In practical terms, the paper establishes this result by obtaining a strong formal statement relating to self-organisation in allocation problems, and then illustrates it by using Maximum Entropy (MaxEnt) signal restoration to successfully predict the evolution of a simple self-organising agent-based model.

The rest of the paper is structured as follows. The conceptual structure of the argument and the various strands of literature required to obtain it are first detailed in section 2, and section 3 explains how the use of MaxEnt in signal restoration is related to its existing use in economics. Section 4 then shows the equivalence of allocation problems and congestion games under reasonable assumptions on preferences, opening up the use of signal restoration algorithms due to the presence of systematic improvement paths. This possibility is then illustrated in section 5 by applying MaxEnt to the Schelling model of segregation. Section 6 discusses the implications of these findings and concludes.

2 Review of the relevant literature

The main contribution of the paper is that as a consequence of the self-organising nature of economic systems, signal restoration algorithms, which are normally designed to retrieve the past

states of physical systems, can be used instead to look into the future. Several distinct strands of literature are combined to demonstrate this concept both theoretically and with a practical application.

The first strand is taken from the extensive economics and operational research literature on object allocation, initiated amongst others by Hurwicz (1973), Harris and Raviv (1981) or Myerson (1981). The specific type of allocation problem used in the paper is based on the knapsack problem, which is a well known combinatorial optimisation problem where one has a set of objects with given values and weights and the objective is to pick the combination of objects with the highest value, without exceeding fixed a weight limit, i.e. the capacity of the knapsack. The specific variant used here is the multichoice multi-dimensional knapsack problem (MMKP), in which several groups of objects are available, with each object providing a specific value and requiring a particular subset of several distinct sets of resources. The objective is to pick a single object from each of the groups, maximising their aggregate value while ensuring the multi-dimensional resource constraint is met. As will be show in section 4.1, this offers enough flexibility to model a very general economic allocation problem. In fact, the MMKP has already been used in the operational research literature to model problems such as the allocation of nurses with different skills and time preferences to different types of shifts (Dowland and Thompson, 2000), or the allocation of distinct computing resources such as memory and CPU cycles to several networked users with different session preferences (Khan et al, 2002).

In contrast to the object allocation literature, however, the aim of modelling object allocation with the MMKP is not to provide a practical solution to or an optimal design for the general allocation problem, but rather to establish an important property relating to self-organisation. This is achieved by showing that the MMKP allocation problem is related to the literature on congestion and potential games. Congestion games were introduced by Rosenthal (1973) as a way of modelling resource allocation on a network, with Rosenthal's initial example of road users attempting to select routes so as to minimise the congestion they experience providing the name to this class of games. The paper then relies on the fundamental result provided by Monderer and Shapely (1996), who prove that every congestion game is isomorphic to a potential game, where a single potential function encodes all the payoffs of the game. This in turn ensures the existence of the finite improvement property (FIP), which implies that every initial condition is linked to a Nash equilibrium by a finite number of improvement steps. As will be shown in section 4, this chain of reasoning allows us to make a strong statement of self-organisation for general allocation problems.

A third important strand of literature used in the paper relates to the use of the MaxEnt methodology both in existing economic research and in signal restoration literature. Because the MaxEnt method forms the base of the proposed algorithm, the formal relation between the two approaches is detailed separately in section 3, however it is important to point out the track-record that MaxEnt already possess in economic theory. Versions of this methodology have been used by Foley (1994) and Toda (2010) to prove the existence of a statistical market equilibrium when agents have offer sets of transactions they are willing to accept and meet in a random fashion. Applied investigations using the MaxEnt methodology include Castaldi and Milaković (2007), which investigates the distribution of wealth using information on turnover in portfolios and Alfarano and Milaković (2008) which similarly explored the origin of the Laplace distribution of firm growth rates. A second important reason for choosing the MaxEnt signal restoration algorithm is that because it uses Shannon (1948) information entropy as its objective function, it offers a natural information-theoretical intuition of why signal restoration in the presence of self-organisation is in fact a prediction on the future, in keeping with the Schrödinger (1967) definition of self-organisation mentioned above.

Finally, as stated in the introduction, a practical demonstration of the methodology is provided by applying the proposed methodology to the Schelling (1969, 1971) model of segregation. The model is chosen because of its simplicity,¹ and crucially because the literature establishes that its strong self-organisation is underpinned by presence of a potential function. Indeed, in the physical analog to the Schelling model proposed by Vinkovic and Kirman (2006), particles on a lattice systematically rearrange themselves to reduce the internal energy of their configuration, and the overall energy of the system provides the potential function. Very recent analysis of the model by Grauwin et al (2012) confirms that it possess a potential function for continuous neighbourhoods under some parameterisations. The existence of the a potential function for the Schelling model ensures the presence of the FIP, making it ideally suited as a test bed for illustrating the theoretical argument proposed in the paper.

3 Maximum entropy as a Bayesian signal restoration method

The use of MaxEnt as a signal restoration methodology is formally equivalent to its use in the economics literature, with two main differences that will be explained below. The first is that the methodology has an explicitly Bayesian setting, with maximum entropy being used to assign priors

¹ Marks (2007) shows that the functional complexity of the Schelling model is relatively low, with most of it coming from the random initial condition. This makes the model a desirable benchmark, as the proposed methodology is useless if it cannot function on even the simplest cases.

for a degenerate maximisation. The second is that the entropy function used for signal restoration is more general and thus offers more flexibility than the one used in the literature mentioned above.

As an illustration, let us assume that a message made up of N characters, each one drawn from an alphabet of length C is transmitted to an observer over a noisy channel. This general structure encompasses a wide range of possible messages, such a string of N letters, intensities in a N -pixel image, or some data measurements for N agents, etc. As a result of the noisy transmission channel, the data d received by the observer will be different from the original message. Supposing that the observer wishes to identify the best possible reconstruction of the signal h from the imperfect data d , a reasonable criterion for doing so is to maximise the posterior probability of h being the true signal given the data available, i.e. $\arg \max_h p(h|d)$. Bayes' rule states that this posterior can be expressed as the product of a prior on h and the likelihood $p(d|h)$, normalised by the evidence $p(d)$:

$$p(h|d) = \frac{p(h)p(d|h)}{p(d)} \quad (1)$$

As is standard with Bayesian methods, rewriting using the log-likelihood $\ell(d|h)$, simplifies the maximisation of the posterior with respect to h to:

$$\arg \max_h p(h|d) \Leftrightarrow \arg \max_h [p(h) \exp(-\ell(d|h))] \quad (2)$$

At this point one might think that solving (2) is a matter of maximising the likelihood, effectively leaving aside the prior probability of the hypothesis $p(h)$. However, in this case, because h and d have the same size and the same support, there are as many 'parameters' to determine in h as there are 'data points' in d . As a result the maximisation problem is degenerate and the choice of prior will have a strong effect on the prediction. The central tenet of MaxEnt is that in such cases, the prior probability of a hypothetical distribution h should be exponentially proportional to the information entropy of the distribution.

$$p(h) = \frac{\exp(\alpha S(h))}{Z_S} \quad (3)$$

Here α is a regularisation parameter and $S(h)$ is the Shannon (1948) information entropy of the candidate reconstruction h . Using $i \in N$ to index locations in the message and $c \in C$ to index characters in the alphabet, the following information entropy measures the uncertainty of the observer as to the content of a proposed reconstruction h :

$$S(h) = - \sum_i \sum_c h_i^c \ln h_i^c \quad (4)$$

The justification for this, provided by Jaynes (1957) and Foley (1994), is that $S(h)$ is in fact the logarithm of the multiplicity of h , i.e the number of ways the distribution can be realised. This implies that the prior probability (3) of a candidate reconstruction h is simply proportional to its multiplicity. This allows the methodology to be “maximally noncommittal with regard to missing information” (Jaynes, 1957, p. 623). A formal proof of this is provided by Shore and Johnson (1980), who prove that it is the only method of selecting priors that satisfies three key axioms: the prior leads to a unique prediction; it is independent of the coordinate system used; and it does not depend on whether information about independent systems or sub-systems is accounted for jointly or separately.

In practice, the observer often has knowledge of a pre-existing structure in the message or of some properties of the noise process in the transmission. The observer’s uncertainty about h , which (4) measures in absolute terms, must therefore be corrected so that it is measured relative to this pre-existing knowledge, modeled by m . This is achieved by using the following relative entropy, which is equal to minus the Kullback-Leibler (KL) divergence from m to h , and measures the similarity between two distributions.² It reaches a global maximum of zero for $m = h$ and is strictly negative for $m \neq h$.

$$S(h|m) = - \sum_i \sum_c h_i^c \ln \frac{h_i^c}{m_i^c} \quad (5)$$

In the case where m is a uninformative uniform distribution, the two forms of entropy (4) and (5) are equivalent in a maximisation problem. This corresponds to a situation where the observer has no useful prior knowledge about the possible reconstructions, and measuring uncertainty in absolute terms is equivalent to measuring it relative to an uninformative m .

Given the entropic prior (3) and relative entropy (5), the posterior probability (1) can now be expressed as:

$$p(h|d) = \frac{\exp(\alpha S(h|m) - \ell(d|h))}{p(d) Z_S Z_\ell} \quad (6)$$

One can see from (6) that identifying the candidate distribution h with the highest a posteriori probability involves maximising the information entropy of h subject to a set of constraints provided by the observations received in d :

$$\arg \max_h p(h|d) \Leftrightarrow \arg \max_h (\alpha S(h|m) - \ell(d|h))$$

² Formally, the KL divergence measures how many bits of information are gained by learning that the true distribution is h rather than m . As explained by Cover and Thomas (1991) it is often used as a measure of the distance between two distributions, and its additive inverse is therefore a measure of similarity.

The general prediction is given by the expression below, where $1/\alpha$ plays the role of a Lagrange multiplier.

$$h_i^c = m_i^c \exp\left(\frac{1}{\alpha} \frac{\partial \ell(d|h)}{\partial h_i^c}\right) \quad (7)$$

If one assumes a uniform value for m , this structure nests the existing approaches of Foley (1994) or Toda (2010). For instance, using a uniform m and a likelihood of the form $\ell(d|h) = \sum_i \sum_c c h_i^c - Nd$ recovers the exponential endowment distributions in a pure exchange economy found by Foley (1994), where c is the individual endowment level and d is the aggregate endowment. The advantage of using relative entropy (5) rather than the Shannon entropy (4) is that it allows additional flexibility through the integration of a model m , which can incorporate any prior knowledge of existing structure in the reconstruction h which might be missing from the observed data d . The usefulness of this extra flexibility when making predictions about economic systems will be demonstrated when applying the methodology in section 5. We first need to formally show, however, why it is appropriate to use MaxEnt signal restoration as a prediction methodology in self-organising economic systems.

4 Finite improvement paths in object allocation

4.1 Object allocation as a knapsack problem

As explained in section 2, we show that the object allocation problem facing the social planner can be modeled using a MMKP. Formally, there are N agents in the economy, labeled $i \in \{1, 2, \dots, N\}$, and the social planner has to allocate Q different units amongst those agents. Although this does not influence the general problem, it will be convenient in the discussion to distinguish K types of commodities, labeled $k \in \{1, 2, \dots, K\}$ for which quantities $q_k \in \mathbf{N}$ are available, in which case $Q = \sum_k q_k$. The allocation problem can be solved, in principle, with the following four steps.

- Step 1: The social planner labels all the possible bundles that can be built with the Q units available and lists them in a $2^Q \times Q$ binary table B , shown in table 1. Each $\{0, 1\}$ cell of the table states whether the j^{th} unit is included in the b^{th} bundle or not. The binary string formed by each row therefore provides not only the corresponding bundle's composition but also a unique identifier for the bundle.
- Step 2: The social planner sends the B -table to the N agents who, assuming completeness, rank the 2^Q bundles according to their preferences. The rankings are returned to the social planner

Table 1 Binary bundle identifiers

B	$j = 1$	$j = 2$	$j = 3$	$j = 4$...	$j = Q$
$b = 1$	0	0	0	0	...	0
$b = 2$	1	0	0	0	...	0
$b = 3$	0	1	0	0	...	0
...
$b = 2^Q$	1	1	1	1	...	1

Table 2 Bundle preference ranking

U	$i = 1$	$i = 2$	$i = 3$	$i = 4$...	$i = N$
$b = 1$	1	1	1	1	...	1
$b = 2$
...
$b = 2^Q$	2^Q	2^Q	2^Q	2^Q	...	2^Q

who then builds a $2^Q \times N$ ranking table U , shown in table 2. Under the usual assumptions of transitivity and monotonicity, all agents will rank the full bundle highest and the empty bundle lowest.

- Step 3: The social planner must pick a bundle for each agent, using a $2^Q \times N$ choice matrix X , where the choice variables are $X_{b,i} \in \{0, 1\}$. Importantly, each agent only receives a single bundle, i.e. $\sum_{b=1}^{2^Q} X_{b,i} = 1 \quad \forall i \in N$.³ The goal of the social planner is to maximise the sum of the ranks over agents while remaining within the resource constraint. Formally, this can be expressed as the following MMKP:

$$\begin{aligned} \max \operatorname{tr}(UX') \\ \text{s.t. } :B'X1_N = 1_Q \end{aligned} \tag{8}$$

Here 1_N and 1_Q are the N and Q -length unit vectors respectively. Choosing an objective function for the MMKP is directly related to the problem of choosing a social welfare function. The standard approach of knapsack problems is to maximise the sum of the values of the objects chosen, which in this case means maximising the sum of the individual rankings. This choice is equivalent to a Benthamite social welfare function, where overall social welfare is simply the sum of individual orderings.⁴ The constraint ensures that the sum of the binary identifiers for each selected bundle equals the unit vector, i.e. each unit in Q is selected only once. Expressed in scalar notation, this corresponds to the standard MMKP as presented by Hifi et al (2004); Sbihi (2007). The only differences compared to the more general framework in the operational

³ One can see that even if the agent is allocated two bundles a and b from B , then $a + b$ is also a bundle in B .

⁴ This use of the term Benthamite as a reference to a simple sum of individual rankings is borrowed from Ng (1975). This may seem overly simplistic, however given that a utility function is never uniquely defined, it is in fact possible to modify the choice of social welfare function to some extent within the linear sum framework of the MMKP by first applying monotonic transformations to the rankings expressed by the agents in table U.

Table 3 Allocation table

A^*	$i = 1$	$i = 2$	$i = 3$	$i = 4$...	$i = N$
$j = 1$	0	0	1	0	...	0
$j = 2$	1	0	0	0	...	0
$j = 3$	0	1	0	0	...	0
...
$j = Q$	0	0	0	1	...	0

research literature is that the resource requirement per bundle in B is the same for all i agents and the available capacity is restricted to one for all dimensions in Q :

$$\begin{aligned} \max \quad & \sum_{i=1}^N \sum_{b=1}^{2^Q} U_{b,i} X_{b,i} \\ \text{s.t.} \quad & \sum_{i=1}^N \sum_{b=1}^{2^Q} B_{b,j} X_{b,i} = 1 \quad \forall j \in Q \end{aligned} \quad (9)$$

- Step 4: Once the optimal choice table X^* is obtained, the social planner can build a $Q \times N$ allocation table $A^* = B'X^*$, shown in table 3. This table uniquely assigns every unit in Q to an agent in N , and can therefore be used for the purpose of selecting goods one by one and dispatching them to their allocated owner.

In theory all four steps of the MMKP are feasible and A^* exists. The problem is not tractable in practice, however, and one of the main advantages of the framework is that it neatly separates the types of hurdles facing a social planner. The first is the choice of the correct social welfare function, followed, as pointed out by Hayek (1945), by a high and potentially unfeasible informational requirement (Step 2) and by a large and computationally complex combinatorial optimisation (Step 3).⁵ As a result, although it exists, the optimal allocation A^* is unknown to the social planner. We now move to proving a strong statement about the emergence of a stable decentralised allocation in such a system.

4.2 Knapsacks, congestion games and improvement paths

As explained in section 2, the operational research literature has used the MMKP to model resource allocation on a network, a setting similar to that of the congestion game in Rosenthal (1973). We now show that the two are in fact equivalent under standard assumptions on preferences, which has important implications in terms of self-organisation.

The allocation of Q goods over N agents with preferences given by U is modeled as a road congestion game where $q_k \in \mathbf{N}$ users of K different types have to choose a route i in an N -

⁵ The knapsack problem is known to be NP-complete, in other words solutions to the problem can be verified efficiently (in polynomial time), but there is no known algorithm for calculating the solutions efficiently in the first place.




Fig. 1 Multigraph congestion game

edge multigraph between start point s and finish point f , in order to maximise their payoff $V_{b,i}^k$. Elaborating on Rosenthal (1973), one could imagine that the K different types represent different categories of vehicles, such as cars, trucks, etc. who each generate different congestion costs. This choice of graph as, illustrated by Figure 1, implies that distinct routes follow separate edges i.e. there are no externalities between routes, where the benefit of a user choosing a route might depend on the number of users choosing another route.

As is standard in congestion games, the payoffs $V_{b,i}^k$ for choosing an edge are a function of the number of users already on the edge. These will be derived from the ranking information returned in table 2. Specifically, let us define $\Delta_k U_{b,i} = U_{b,i} - U_{b\{-k\},i}$ as the change in ranks at the margin, following the addition of the last k -type good to bundle b . In terms of notation, $b\{-k\}$ is the bundle obtained by removing a k -type good from bundle b . Similarly, in the following, $b\{+k\}$ will refer to the bundle obtained by adding a k -type good to bundle b .⁶ This allows the derivation of a $2^Q \times N \times K$ payoff array V which is set as $V_{b,i}^k = \Delta_k U_{b,i}$.

The following two assumptions on the rankings in the table 2 are required in order to show equivalence between the MMKP and the congestion game frameworks:

Monotonicity: $\Delta_k U_{b,i} > 0 \quad \forall k, i, b$.

Concavity: Given two bundles a and b , $\forall i \in N$ if $U_{b,i} > U_{a,i}$ then $\Delta_k U_{b,i} < \Delta_k U_{a,i}$.

Monotonicity, which was already implicitly assumed in the description of the rankings table 2, ensures that all the congestion game payoffs in V are strictly positive. Concavity intuitively means that the bundle rankings exhibit decreasing marginal values, as adding extra units of k -type goods to a bundle, keeping the rest of the bundle constant, will bring successively smaller increases in the ranking. Together with monotonicity, this is required in order to ensure that the payoff of choosing a particular edge is decreasing with the number of users on that edge, as in the basic congestion

⁶ $\Delta_k U_{b,i}$ is of course undefined for the empty bundle and whenever bundle b contains no k -type units.

game framework of Rosenthal (1973).⁷ We now prove that if the rankings table U satisfies these two assumptions, the MMKP and congestion game formulations are equivalent.

Proposition: *If the ranking table U displays monotonicity and concavity, the optimal solution to the MMKP problem is a Nash equilibrium for the corresponding multigraph congestion game based on payoff table V .*

Proof: By contradiction. Let X^* be the decision table that satisfies the MMKP (8) and $A^* = B'X^*$ the corresponding allocation of the Q users over the N edges of the multigraph. Let us assume that A^* is not a Nash equilibrium for the Q users. If b and a are the bundles allocated to edges i and x respectively by X^* , this implies:

$$\exists i, x, k : V_{a\{+k\},x}^k > V_{b,i}^k$$

Let Y be the decision table resulting from the switch of the k -type agent from edge i to x . Y is identical to X^* , except for edges i and x , which receive bundles $b\{-k\}$ and $a\{+k\}$ respectively. Using the definition of $V_{b,i}^k$:

$$\begin{aligned} \Delta_k U_{a\{+k\},x} &> \Delta_k U_{b,i} \\ U_{a\{+k\},x} + U_{b\{-k\},i} &> U_{a,x} + U_{b,i} \\ \text{tr}(UY') &> \text{tr}(UX^{*'}) \end{aligned}$$

X^* does not satisfy the MMKP, which is a contradiction. ■

Corollary: *If the ranking table U displays monotonicity and concavity, the objective function of the MMKP is an exact potential function for the corresponding multigraph congestion game based on payoff table V .*

Proof: Immediate from the previous proof and the definition of the payoffs $V_{b,i}^k$. The change in payoff to a k -type user for switching from edge i to x is $V_{a\{+k\},x}^k - V_{b,i}^k$. Given that $V_{b,i}^k = \Delta_k U_{b,i}$ one has:

$$V_{a\{+k\},x}^k - V_{b,i}^k = \Delta_k U_{a\{+k\},x} - \Delta_k U_{b,i} = \text{tr}(UY') - \text{tr}(UX^{*'})$$

The objective function of the MMKP is an exact potential for the congestion game based on the corresponding $V_{b,i}^k$ payoffs. ■

⁷ In the standard framework of Rosenthal (1973), congestion costs on an edge are increasing with the numbers of users on the edge, and the aim of the network users is to choose the edge with the lowest cost. This is equivalent to the framework used here, where the benefit of using an edge falls with the numbers of users on the edge, and users choose the edge with the highest benefit.

As was shown by Monderer and Shapely (1996) and explained in section 2, the equivalence of the MMKP and congestion game frameworks and the existence of a potential function implies the existence of the finite improvement property (FIP). Starting from any initial state a simple myopic best-response path will lead to a Nash equilibrium in a finite number of steps. For the MMKP, an Edgeworth process where agent pairs meet randomly and trade goods with low marginal utility $\Delta_k U_{b,i}$ against goods with higher marginal utility would satisfy the requirement, as the potential function $\text{tr}(UX')$ would increase following such trades.⁸ The central implication of this result is that if the rankings expressed by the N agents are monotonic and concave, then even though the MMKP cannot be solved centrally, the social planner can be confident that the system will eventually reach a decentralised allocation.

The existence of a convergence process to a set of decentralised allocations is no surprise and only duplicates the existing findings of the objet allocation literature mentioned in section 2. Instead, the result of interest is the presence of the FIP, as this is what underpins the proposed use of the MaxEnt signal restoration methodology. Indeed, under the FIP any initial state I is linked to a final Nash equilibrium F by a finite sequence of intermediate states $I \rightarrow F$ on which transitions are the result of agents making best-response, welfare-increasing trades. If, however, the same sequence of states is viewed in reverse, $F \rightarrow I$, transitions now reflect a systematic sequence of random errors. This is because a reverse transition will involve at least one agent moving from a unique and optimal situation to one of many possible sub-optimal ones. The key consequence of this is that the initial condition I can be treated as the realisation of a particular noise process applied to F . The implication is thus that the problem of predicting F from I is equivalent to the signal restoration problem presented in section 3. The following section illustrates this argument by applying the analysis the Schelling model, which is known to possess the FIP.

5 Application to the Schelling model of segregation

5.1 FIP in the Schelling model of segregation

In the standard setting of the Schelling model two types of agents live in a city made up of discrete locations, and each type has a preference for living in a local neighbourhood composed of agents of the same type. When agents are allowed to move, segregated neighborhoods will emerge from an integrated initial condition as agents relocate to unoccupied locations in the city that are more attractive. The attractiveness of a location to an agent is a function of the number of similar

⁸ Clearly, the switching process used in the proof is simplistic: one does not expect goods to choose their owners in order to maximise a payoff! Edgeworth trading, however, can be broken down into a sequence of such switches: a k -type good moves from agent i to agent x , immediately followed by a another commodity switching from x to i .




Fig. 2 Emergence of segregation in the Schelling model

agents in the vicinity, which usually determined by counting the number of similar agents within a neighbourhood of given width. If B is a $N \times N$ binary matrix which identifies the neighbours for all N locations, and l^c is the binary vector for the location of c -type agents, this similarity for each location i is given by:

$$(B \times l^c)_i = \sum_j B_{i,j} l_j^c \quad (10)$$

As is the case with the work of Grauwil et al (2012), it is assumed that the space occupied by the city is toroidal, so that the top/bottom and left/right edges are in contact. This simplification allows the neighbourhood matrix B to be encoded as a circulant matrix, which greatly facilitates the analysis. A further assumption used here is that the utility of an agent is directly proportional to the number of similar neighbors. This is contrast to original Schelling (1969, 1971) model,

where utility is a unimodal function of similarity, initially increasing with similarity, peaking for a balanced neighbourhood composed of 50% of agents of each type, then declining as similarity increases further. Indeed, Grauwin et al (2012) show that in the case of continuous neighborhoods such as (10), the existence of a potential function - critical to presence of the FIP discussed in section 4.2 - requires that utility functions be linear functions of the number of similar agents.

Because it is assumed that relocation opportunities arrive randomly (typically as a Poisson process), simulation is usually the method of choice for investigating this model. Grauwin et al (2012) themselves point out that most analyses of this model rely on agent-based simulations and lack analytical solutions. Such a simulation is therefore provided as a point of reference for the MaxEnt prediction methodology.⁹

The sequence of images in Figure 2 (a)→(i) provides a time-lapse of the simulation process, and illustrates the concept of a finite improvement path. The random initial state is given in Figure 2 (a), while Figure 2 (i) represents the state of the city after 44841 individual moves have occurred. The final state in 2 (i), which exhibits the segregated outcomes typical of the Schelling model, is stable as no further utility-improving relocations exist.

5.2 MaxEnt prediction

The finite improvement path visible in Figure 2 provides a simple yet effective illustration for the argument made in section 4.2. Indeed, viewing the finite improvement path in reverse, (i)→(a), provides a situation where an initially well-defined and coherent image gradually picks up noise, eventually obscuring the initial information. Thus, the problem of predicting 2 (i) from the initial condition 2 (a) is equivalent to the problem of reconstructing a signal 2 (i) from noisy data 2 (a). We show that predictions of the Schelling model from the initial condition can indeed be obtained using a modified version of the MaxEnt image reconstruction approach discussed in section 3, specifically the algorithm of Skilling and Gull (1991).

Within the setting described above in 5.1, let h_i^c be the probability that the i^{th} location is occupied by an agent of the c^{th} colour, with $c \in \{R, G, W\}$ and $\sum_c h_i^c = 1$. Given this, relative entropy (5) measures the expected information content of a message revealing the final state of a randomly picked location, relative to any prior information on the location of agents:

$$S(h|m) = -\frac{1}{N} \sum_i \sum_c h_i^c \ln \frac{h_i^c}{m_i^c} \quad (11)$$

⁹ The parameters for the benchmark simulation are as follows: the city is 200 pixels across and each pixel represents a location, so there are $N = 200^2 = 40000$ locations. There are $N^R = N^G = 16000$ red and green agents and $N^W = 8000$ free spaces. The continuous neighborhood agents consider when assessing the desirability of a given location is a 7×7 square area centered on that location.

As outlined in Section 3, relative entropy (11) encodes prior information through the underlying model m_i^c . In the Schelling model, however, agent satisfaction does not depend on absolute location, but on location relative to other agents. As a result, there is no prior information regarding the probability of a single location being occupied by a particular type of agent, and m_i^c in expression (11) is not particularly useful. This is dealt with by following Skilling and Gull (1987) and considering the expected information content of a message revealing the state $\{c, d\}$ of two randomly picked locations $\{x, y\}$.¹⁰ Using expression (12) enables the integration of a two-dimensional model $m_{i,j}^{c,d}$ which can contain knowledge of correlations across locations.¹¹ This is better suited to the prior information provided in the Schelling model, in which one expects neighbouring locations to have a relatively high probability of being occupied by similar agents.

$$S(h_x, h_y | m_{x,y}) = \frac{1}{N} \left(-2 \sum_i \sum_c h_i^c \ln h_i^c + \frac{1}{N} \sum_{i,j} \sum_{c,d} h_i^c h_j^d \ln m_{i,j}^{c,d} \right) \quad (12)$$

The second important piece of information required for the MaxEnt prediction is the initial condition of the system, which provides the data entering the likelihood in (6). With the reversed FIP, where Figure 2 (i) decays to a noisy state in Figure 2 (a), this represents the information that has not been wiped out in the decayed image. Within the Schelling setting, this intuitively represents the key stable locations that are initially most attractive and are not modified as the segregated outcome emerges. This information is revealed by taking the convolutions of the initial state in order to determine the initial attractiveness (10) for each type of population, shown in Figure 3.

As a further simplification we assume, following the standard image restoration literature, that the divergence between the prediction h^c and the initial attractiveness data d^c is normally distributed. This implies that the likelihood $p(d^c | h^c) = \exp(-\ell(d^c | h^c)) / Z_\ell$ is gaussian, and therefore the log-likelihood $\ell(d^c | h^c)$ is directly related to the chi-squared deviation between the initial data available and the prediction, where $(\sigma^c)^2$ is the variance of the d^c data.

$$\ell(d^c | h^c) = \sum_i \frac{((B * h^c)_i - d_i^c)^2}{(\sigma^c)^2} = \frac{\chi^2(h^c)}{2} \quad (13)$$

The the information theoretic problem is therefore to maximise the ignorance of an observer (12) subject to the information provided by the likelihood (13), normalising to ensure that the predicted number of agents of each colour equals the initial amount N^c . As pointed out by Skilling

¹⁰ The derivation of the double entropy specification is detailed in appendix A.

¹¹ This structure also allows correlations across agent types, for example if agents were to evaluate the attractiveness of a location not only by the number of similar agents but also by the number of agents of a different type. This is not the case here as in the basic Schelling model, agents only consider their own type in their location decision, in other words $m_{i,j}^{c,d} = 0 \quad \forall d \neq c$.

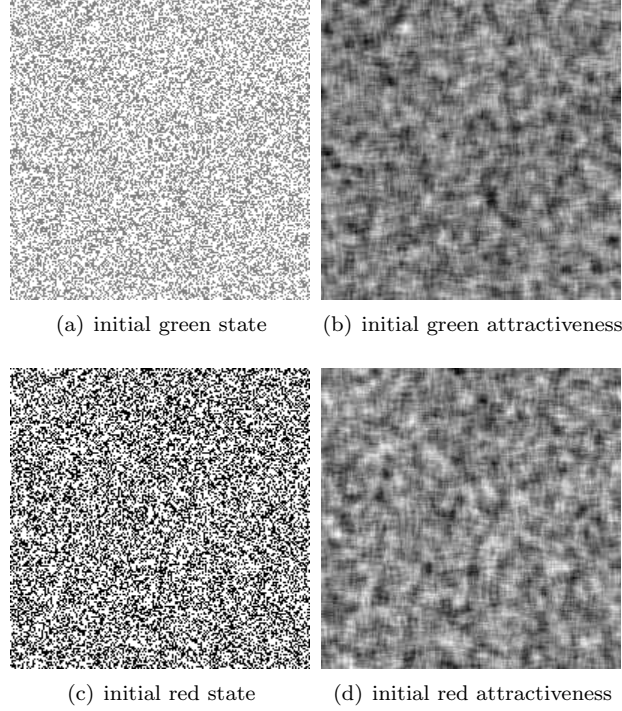


Fig. 3 Initial condition information

and Gull (1991), the value of the implicit Lagrange parameter α is used to constrain the noise level measured by (13) to be equal to number of degrees of freedom controlled by noise, i.e. the overall number of locations N minus the number of good locations I^c in the initial data.¹²

The first order condition of the problem directly provides the best prediction for the distribution of agents over the locations:

$$h_i^c = \frac{\mu_i^c}{Z^c} \exp\left(-\frac{1}{2\alpha^c} \frac{\partial \chi^2(h^c)}{\partial h_i^c}\right) \quad (14)$$

The effective model μ_i^c and the normalisation parameter Z^c are given by:

$$\mu_i^c = \exp\left(\frac{1}{2N} \sum_j \sum_{c,d} h_j^d \ln m_{i,j}^{c,d}\right) \quad \text{and} \quad Z^c = \frac{1}{N^c} \sum_i \mu_i^c \exp\left(-\frac{1}{2\alpha^c} \frac{\partial \chi^2(h^c)}{\partial h_i^c}\right)$$

One can see that the effective model for a location μ_i^c is simply the geometric mean of the individual correlations $m_{i,j}$, weighted by the probability vector. As pointed out by Skilling and Gull (1987), this is effectively a convolution of the reconstruction h^c with the logarithm of the $N \times N$ model matrix, similar to (10).¹³ It is important to point out that expression (14) only

¹² The relation between α and the number of noisy degrees of freedom as well as the calculation of I^c are explained in appendix B.

¹³ In practice the convolution used in the prediction algorithm is different: instead of calculating $h_j^d \ln m_{i,j}^{c,d}$ the algorithm uses $m_{i,j}^{c,d} \ln h_j^d$. This is done for computational reasons. Most of the entries in the model M are vanishingly small as one expects the correlations across locations to exist only over short distances. As a result they are truncated




Fig. 4 MaxEnt predictions vs. Monte Carlo frequencies

provides an implicit solution for the probability distribution h_i^c as both the model term μ_i^c and noise term $\chi^2(h^c)$ are themselves functions of h_i^c . The predicted distributions are therefore obtained using a gradient-based algorithm, outlined in appendix B.

A Monte-Carlo (MC) analysis was carried out in addition to the MaxEnt reconstruction in order to assess the predictive power of the proposed methodology. Colour-specific frequencies f^c were obtained by running 1000 Monte-Carlo (MC) iterations of the Schelling model on the same initial condition. Intuitively, f_i^c indicates the percentage of simulations that resulted in the i^{th} location being occupied by an c -coloured agent, and these can directly be compared to h_i^c , the probability that location i is of colour c in the MaxEnt reconstruction. Figures 4 (a) and (d) are the colour-specific results of the simulation shown in Figure 2, while figures 4 (c) and (f) provide the MaxEnt prediction (14) given the information from the initial condition in Figure 3. The MC frequencies f^c are shown in Figures 4 (b) and (e).

Table 4 provides two main measures of goodness of fit for the reconstruction h^c compared to the MC frequencies f^c for both types of agents. The first is the Spearman rank correlation between the two vectors, which is large and highly significant. A second measure is the relative mean square error (MSE) of h^c as a predictor of f^c , which is calculated as the sum of squared deviations between

out of the matrix, which can be stored efficiently as a sparse matrix with many zero elements. Taking the logarithm of this sparse matrix thus becomes problematic, therefore in practice it is easier to take the logarithm of the h^c .

Table 4 Goodness of fit tests, Monte Carlo vs. MaxEnt

	Spearman's ρ	p-value	MSE/ σ^2	Standardised MSE/ σ^2
Green	0.910	0	0.9879	0.2292
Red	0.8190	0	0.9888	0.4546

the h^c and f^c relative to the variance of the MC frequency f^c , i.e. $(f-h)'(f-h)[(f-\bar{f})'(f-\bar{f})]^{-1}$. The MSE values just below unity suggests that h^c is a better predictor of f^c than the expectation $\bar{f}^c = \bar{h}^c = N^C/N$, but only by a small margin. This is because in practice the reconstruction h^c is very flat, with a standard deviations of 0.0023 and 0.0026 for the greens and reds respectively, which means that h^c never moves very far from its expected value $\bar{h}^c = N^C/N$. Intuitively, there is so much noise in the initial information of Figure 3, as measured by the chi-squared deviation (13), that the degrees of freedom available for the description of the reconstructed image (14), measured by $\Gamma^c = N - \chi^2(h^c)$, are severely limited.

Once the relative flatness of the reconstruction is controlled for by standardising both vectors h^c and f^c , the MSE measure falls significantly, which is shown in the final column of Table 4. This standardised MSE supports the good predictive power of h^c with regards to the location of red/green agents, which is what is revealed by visual comparison in Figure 4. One can conclude from this that although the initial information in Figure 3 is very noisy and therefore seems of limited use, and even though the power of the MaxEnt methodology is limited in terms of predicting the amplitude of the absolute frequencies f^c , it nevertheless provides a reliable prediction for the relative locations of the two sets of agents.

6 Discussion and Conclusion

The formal justification provided for transposing the signal restoration interpretation of MaxEnt into economics is the structural similarity between object allocation problems, as modeled by the MMKP, and congestion games frameworks. In fact, the only requirement for the optimal MMKP allocation to also be a Nash equilibrium in a corresponding congestion game is concavity in the bundle rankings, allowing the MMKP objective function to become an exact potential for the game. Assuming this is the case, the system displays the FIP, i.e. from any initial state there exists a finite path to a Nash equilibrium under even the simplest adjustment dynamics. It is the presence of the FIP and the corresponding improvement path that then provides the key motivation for the use of the MaxEnt signal restoration methodology in such a system. Indeed, the sequence of states forming the improvement path can be interpreted in two ways, depending the direction in which it is viewed.

If the improvement path is viewed forward, starting at the initial state and finishing at the equilibrium, the picture one has is of a system that gradually self-organises as agents perform systematic welfare-increasing trades. This corresponds to the self-organisation concept mentioned in introduction and, as illustrated with Figure 2, this interpretation corresponds to the biological entropy reduction insight of Schrödinger (1967): the amount of information required to describe the state of the overall system falls over time as it organises itself. If, however, the same improvement path is played backwards, i.e. starting at the equilibrium and traveling back towards the initial condition, then a physical analogy is more appropriate. Agents now appear to be making systematic welfare-decreasing trades, in other words systematic errors. With this reverse view, the initially ordered state gradually decays as noise is introduced. This corresponds to the direct physical interpretation of entropy increases: An initially ordered system, say an ice cube in a glass of water, which gradually decays into a disordered thermal equilibrium. In this case the increasing information entropy measures ignorance as to the amount and type of noise that has been introduced.

The central finding is therefore that if a system possesses the FIP, then predicting its equilibrium distribution from a known initial state is formally equivalent to reconstructing an unobserved clean signal out of an observed noisy one. This is illustrated by the application to the Schelling model, where such an algorithm is shown to be able to predict the emergence and location of segregated neighbourhoods with a good level of accuracy. Not only does this provide a stronger justification for the use of MaxEnt and information-theoretic methods in economics, but it also suggests that signal reconstruction algorithms, designed specifically for the purpose of recovering a degraded signal, could become useful tools in agent-based computational economics. As pointed out by Fagiolo et al (2007) as part of a special issue of *Computational Economics* the empirical validation in agent-based models is still an open question, and the methodology outlined here might be helpful in that regard.

There are two important considerations going forwards, however. The first is that the use of the FIP as a formal definition of self-organisation is probably too strong. The FIP implies both systematic improvements in transitions and a finite horizon, i.e. the existence of a stable final state. While this helps to underpin the use of MaxEnt, it is easy to imagine richer self-organising systems where neither holds, for example if forward-looking agents accept short run welfare reductions in order to benefit in the longer run or simply if agents occasionally make mistakes. It is therefore interesting to examine the performance of signal reconstruction methods in such weaker environments. A second consideration is that the MaxEnt algorithm is chosen because versions of it have already been used in economics, and because it offers an intuitive illustration of

the arrow-of-time interpretation of the mechanism. However, while it seems to perform adequately in the case presented here, it is only one of many available signal reconstruction algorithms, and may not be the most powerful or the most adapted to the specific economic case.

Acknowledgements The author wishes to acknowledge the suggestions received at the ESHIA 2011 and CEF 2012 conferences, as well as thank the GREQAM seminar participants for suggesting the Schelling application for the MaxEnt methodology. Particular thanks goes to Jagjit Chadha, Alan Kirman, Mishael Milaković, Lionel Nesta and Alexis Akira Toda for their helpful advice, and especially to Sonia Moulet for tirelessly providing a sounding board for the author's ideas. Any errors are the author's.

References

- Alfarano S, Milaković M (2008) Does classical competition explain the statistical features of firm growth? *Economics Letters* 101:272–274
- Castaldi C, Milaković M (2007) Turnover activity in wealth portfolios. *Journal of Economic Behavior and Organization* 63:537–552
- Cover TM, Thomas JA (1991) *Elements of Information Theory*. John Wiley & Sons
- Dowland KA, Thompson JM (2000) Solving a nurse scheduling problem with knapsacks, networks and tabu search. *Journal of the Operational Research Society* 51:825–833
- Fagiolo G, Moneta A, Windrum P (2007) A critical guide to empirical validation of agent-based models in economics: Methodologies, procedures, and open problems. *Computational Economics* 30:195–226
- Foley DK (1994) A statistical equilibrium theory of markets. *Journal of Economic Theory* 62:321–345
- Grauwin S, Goffette-Nagot F, Jensen P (2012) Dynamic models of residential segregation: An analytical solution. *Journal of Public Economics* 96:124–141
- Harris M, Raviv A (1981) Allocation mechanisms and the design of auctions. *Econometrica* 49:1477–1499
- Hayek F (1945) The use of knowledge in society. *American Economic Review* 35:519–530
- Hifi M, Michrafy M, Sbihi A (2004) Heuristic algorithms for the multiple-choice multidimensional knapsack problem. *Journal of the Operational Research Society* 55:1323–1332
- Hurwicz L (1973) The design of mechanisms for resource allocation. *American Economic Review* 63:1–30
- Jaynes ET (1957) Information theory and statistical mechanics i. *The Physical Review* 106:620–630

-
- Khan S, Li KF, Manning EG, Akbar M (2002) Solving the knapsack problem for adaptive media systems. *Studia Informatica* 2:161–182
- Marks RE (2007) Validating simulation models: A general framework and four applied examples. *Computational Economics* 30:265–290
- Monderer D, Shapely LS (1996) Potential games. *Games and Economic Behaviour* 14:124–143
- Myerson RB (1981) Optimal auction design. *Mathematics Of Operations Research* 6:58–73
- Ng YK (1975) Bentham or bergson? finite sensibility, utility functions and social welfare functions. *The Review of Economic Studies* 42:545–569
- Rosenthal RW (1973) A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2:65–67
- Sbihi A (2007) A best first search exact algorithm for the multiple-choice multidimensional knapsack problem. *Journal of Combinatorial Optimisation* 13:337–351
- Schelling TC (1969) Models of segregation. *American Economic Review* 59:488–493
- Schelling TC (1971) Dynamic models of segregation. *Journal of Mathematical Sociology* 1:143–186
- Schrödinger E (1967) *What Is Life? with Mind and Matter and Autobiographical Sketches*, Cambridge University Press, chap What Is Life?
- Shannon CE (1948) A mathematical theory of communication. *The Bell System Technical Journal* 27:379–423
- Shore JE, Johnson R (1980) Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory* 26:26–37
- Skilling J, Gull SF (1987) Maximum-Entropy and Bayesian spectral analysis and estimation problems, Kluwer, chap Prior knowledge must be used
- Skilling J, Gull SF (1991) Bayesian maximum-entropy image reconstruction. *Spatial Statistics and Imaging* 20:341–367
- Tesfatsion L (2006) *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, North-Holland, chap Agent-Based Computational Economics: A Constructive Approach to Economic Theory, pp 831–880
- Toda AA (2010) Existence of a statistical equilibrium for an economy with endogenous offer sets. *Economic Theory* 45:379–415
- Vinkovic D, Kirman A (2006) A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences* 103:19,261–19,265

A Information-theoretic framework

The key difference between the standard relative information content (11) typically used in the image processing literature and the specification (12) used here is the use of the double space entropy suggested by Skilling and Gull (1987) to integrate prior knowledge of *relative* rather than *absolute* positions of agents. Formally, the relative entropy is the same as (11), except that it encodes the information content of a message revealing the colours $\{c, d\}$ of a randomly chosen pair of locations $\{x, y\}$, relative to what would be expected given prior knowledge of correlations $m_{x,y}$:

$$S(h_x, h_y | m_{x,y}) = -\frac{1}{N^2} \sum_{i,j} \sum_{c,d} h_{i,j}^{c,d} \ln \left(\frac{h_{i,j}^{c,d}}{m_{i,j}^{c,d}} \right) \quad (\text{A-1})$$

$$S(h_x, h_y | m_{x,y}) = -\frac{1}{N^2} \sum_{i,j} \sum_{c,d} h_{i,j}^{c,d} \ln h_{i,j}^{c,d} + \frac{1}{N^2} \sum_{i,j} \sum_{c,d} h_{i,j}^{c,d} \ln m_{i,j}^{c,d} \quad (\text{A-2})$$

Treating the joint probability as the product of the marginal probabilities $h_{i,j}^{c,d} = h_i^c h_j^d$, and recognising that $\sum_i \sum_c h_i^c \ln h_i^c = \sum_j \sum_d h_j^d \ln h_j^d$, one obtains the specification used in equation (12). Although the existence of correlations in the model $m_{i,j}^{c,d}$ means that the probabilities are not in fact independent, this assumption allows the relative entropy to measure the extra information required to treat probabilities h_i^c and h_j^d as independent when they are in fact related by the underlying model.

Given the specifications for the entropy (12) and the the likelihood (13), maximising the posterior distribution involves solving the following maximum entropy problem. It is assumed that the α^c parameter integrates the multiplicative $2/N$ term in (12).

$$\arg \max_{h_i^c} \left(\alpha^c S(h_x, h_y | m_{x,y}) - \frac{\chi^2(h^c)}{2} \right) \quad (\text{A-3})$$

This leads to the following first order condition with respect to h_i^c :

$$\alpha^c (-\ln h_i^c - 1 + \ln \mu_i^c) - \frac{1}{2} \frac{\partial \chi^2(h^c)}{\partial h_i^c} = 0 \quad (\text{A-4})$$

$$h_i^c \propto \mu_i^c \exp \left(-\frac{1}{2\alpha^c} \frac{\partial \chi^2(h^c)}{\partial h_i^c} \right) \quad (\text{A-5})$$

Because N^c , the total number of agents of a particular colour, is given in the initial condition and does not change over time, it is possible to derive a partition function Z^c which serves to normalise the distribution over locations:

$$\sum_i h_i^c = N^c \Rightarrow Z^c = \frac{1}{N^c} \sum_i \mu_i^c \exp \left(-\frac{1}{2\alpha^c} \frac{\partial \chi^2(h^c)}{\partial h_i^c} \right) \quad (\text{A-6})$$

B Maximum entropy algorithm

The algorithm used to obtain the probability distribution (14) follows from Skilling and Gull (1991).¹⁴ The initial probability and model vectors are given by the uniform distribution $h_i^c = m_i^c = N^c/N$. Prior to running the algorithm, the initial conditions are processed in order to extract the relevant data for calibrating the model constraints:

¹⁴ The code for the Schelling simulation and the MaxEnt reconstruction algorithm is available from the author on request, as well as the initial condition matrix required for replicating the figures shown here.

- The initial attractiveness data vector d^c is calculated as a convolution of initial state vector l_0^c , i.e. $d^c = B \times l_0^c$.
- The initially most attractive locations G are determined as those where $d_i^c \gtrsim \bar{d}^c \pm 2\sigma^c$. Because these good locations are clustered, the number of distinct clusters Γ^c is obtained by convolving the initial attractive locations G with B a second time to identify those which most attractive because located closest to each other. This provides $\Gamma^R = 14$ and $\Gamma^G = 11$.
- Finally the expected radius of a cluster $b = \sqrt{G/(\Gamma^c * \pi)}$ is calculated. This is used to calibrate the model M^c , which is assumed to be a circulant matrix containing a gaussian convolution of standard deviation b .

B.1 Newton method iteration

The iterative algorithm uses the Newton method to find the most probable reconstruction h . In terms of notation, assuming that v is a $N \times 1$ column vector, $[v]$ refers to a $N \times N$ diagonal matrix with the entries of v on the main diagonal and zeros off the main diagonal. Referring to Q^c as the argument of the maximisation in (A-3), the Jacobian vector and Hessian matrix are given by:

$$\begin{cases} \nabla Q^c = \alpha^c \nabla S(h_x, h_y | m_{x,y}) - \nabla \ell(d^c | h^c) \\ \nabla \nabla Q^c = \alpha^c \nabla \nabla S(h_x, h_y | m_{x,y}) - \nabla \nabla \ell(d^c | h^c) \end{cases} \quad (\text{A-7})$$

The step change in the probability vector at each iteration can be calculated using the standard Newton method:

$$\Delta h^c = -(\nabla \nabla Q^c)^{-1} \cdot \nabla Q^c \quad (\text{A-8})$$

Given that the Hessian matrix $\nabla \nabla Q^c$ is symmetric by construction, calculation of the iteration step (A-8) can be carried out efficiently by using the Preconditioned Conjugate Gradient method (PCG) to solve $-\nabla \nabla Q^c \cdot \Delta h^c = \nabla Q^c$ without inverting the Hessian $\nabla \nabla Q^c$. Once this is done, the prediction is updated: $h^c + \Delta h^c$. The model is also updated at this point using $\Delta \mu^c = [\mu^c] [h^c]^{-1} M^c \Delta h^c$.

B.2 Control and termination

Two related control issues must be solved as the Newton iterations proceed. First of all, the value of the α^c parameter has to be determined and adjusted, and secondly the iteration must be terminated at some point. The main advantage of the Skilling and Gull (1991) approach is that it is the optimal value of α^c which both controls the iteration process and provides this termination condition. By integrating α^c into the hypothesis space of the Bayesian problem, they show that the most probable $\hat{\alpha}^c$ satisfies:

$$-2\alpha^c S(h_x, h_y | m_{x,y}) = \text{tr} \left((\alpha^c I + L^c)^{-1} L^c \right) \quad \text{where} \quad \alpha^c I + L^c = [h^c]^{\frac{1}{2}} \nabla \nabla Q^c [h^c]^{\frac{1}{2}} \quad (\text{A-9})$$

If λ_i^c are the eigenvalues of L^c , then $\text{tr} \left((\alpha^c I + L^c)^{-1} L^c \right) = \sum_i \lambda_i^c / (\alpha^c + \lambda_i^c)$. The trace term is therefore a measure of the number of good observations in the data, i.e. the number of dimensions for which $\lambda_i^c \gg \alpha^c$, and the role of α^c is to identify the number of good observations and hence the amount of noise, as $-2\alpha^c S + \chi^2 = N$. If r is a $N \times 1$ vector of $N(0, 1)$ errors, then the trace term can be estimated by:

$$\text{tr} \left((\alpha^c I + L^c)^{-1} L^c \right) = \left\langle r' [h^c]^{-\frac{1}{2}} (\nabla \nabla Q^c)^{-1} \cdot [h^c]^{-\frac{1}{2}} L^c r \right\rangle \quad (\text{A-10})$$

This implies that the $\text{tr} \left((\alpha^c I + L^c)^{-1} L^c \right)$ term can be calculated by using PCG to solve $\nabla \nabla Q^c Y = [h^c]^{-\frac{1}{2}} L^c r$, then calculating $r' [h^c]^{-\frac{1}{2}} Y$. Given the similarity of (A-10) and the step-size problem (A-8), this is carried out in parallel to the main iteration at very little extra cost. This provides control by providing a target value $\tilde{\alpha}^c = -\text{trace}/(2S)$ towards which the α^c parameter can be adjusted at each iteration.

In the original Skilling and Gull (1991) algorithm, equation (A-9) also provides the following termination condition for the algorithm, which is satisfied when $\Omega \approx 1$.

$$\Omega = -\frac{\text{tr} \left((\alpha^c I + L^c)^{-1} L^c \right)}{2\alpha^c S(h_x, h_y | m_{x,y})} \quad (\text{A-11})$$

Given that the number of distinct good locations Γ^c is known in advance, (A-9) and (A-11) are modified to take this into account, by rescaling α^c with a free parameter θ , shown below. This parameter ensures that when the $\Omega \approx 1$ termination condition is reached $\hat{\alpha}^c = \tilde{\alpha}^c$. More importantly, it also ensures $2(\hat{\alpha}^c \theta) S = \text{tr}((\hat{\alpha}^c I + L^c)^{-1} L^c) = \Gamma^c$ and $\chi^2(\hat{h}^c) = N - \Gamma^c$.

$$\theta = -\frac{2\alpha^c S(h_x, h_y | m_{x,y})}{\Gamma^c} \quad (\text{A-12})$$

$$\alpha^c = -\theta \frac{\text{tr} \left((\alpha^c I + L^c)^{-1} L^c \right)}{2S(h_x, h_y | m_{x,y})} \quad (\text{A-13})$$

$$\Omega = -\theta \frac{\text{tr} \left((\alpha^c I + L^c)^{-1} L^c \right)}{2\alpha^c S(h_x, h_y | m_{x,y})} \quad (\text{A-14})$$