Spatial Mobility in the Formation of Agent-Based Economic Networks

Camillia Zedan
Institute for Complex Systems Simulation
University of Southampton, UK
Email: c.zedan@soton.ac.uk

Seth Bullock
Institute for Complex Systems Simulation
University of Southampton, UK
Email: sgb@ecs.soton.ac.uk

Antonella Ianni
School of Social Sciences
University of Southampton, UK
Email: a.ianni@soton.ac.uk

***Abstract*—We extend the model of spatial social network formation of Johnson and Gilles (Review of Economic Design, 2000, 5, 273-299) by situating each economic agent within one of a set of discrete spatial locations and allowing agents to maximise the utility that they gain from their direct and indirect social contacts by relocating, in addition to forming or breaking social links. This enables the exploration of scenarios in which agents are able to alter the distance between themselves and other agents at some cost. Agents in this model might represent countries, firms or individuals, with the distance between a pair of agents representing geographical, social or individual differences. The network of social relationships characterises some form of self-organised persistent interaction such as trade agreements or friendship patterns. By varying the distance-dependent costs of relocation and maintaining social relationships we are able to identify conditions that promote the formation of spatial organisations and network configurations that are pairwise stable and efficient. We also examine the associated patterns in individual and aggregate agent behaviour. We find that both relative location and the order in which agents are allowed to act can drastically affect individual utility. These traits are found to be robust to random perturbations.**[1]

***Keywords*—Spatial social networks, endogenous network formation, pairwise stability, network efficiency, agent mobility.**

***JEL Classification Codes*—A14, C63, D85.**

I. Introduction

In response to the damage caused by the recent global economic crisis, the UK government announced plans to restructure the way in which the financial markets are regulated [2]. It is believed that suitable regulation could enable the early detection and possible prevention of future crises by, for example, encouraging market stability or reducing risk [3]. Overall, this move has highlighted the importance of understanding how complex economic systems, such as the financial markets, form and evolve over time.

Agent-based computational economics can be used to identify conditions that promote positive economic behaviour and provide a testing platform for potential future policies [4].

[1]This work was supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1).

Within the last ten years, research in economics has embraced networks as an important explanatory tool for the highly complex nature of economic systems [6]. For instance, networks have been used to study social networking, trade routes and the spread of information [5]. Networks are widely used across disciplines as a way of representing and exploring interaction patterns between entities, and are particularly useful for studying the effect of low-level component changes at an aggregate level. The structural representation of networks also enables the detection of relationship patterns that might otherwise be indiscernible, yet may still play an integral role to the development of the system as a whole.

The formation and evolution of network structures is of key interest to economists. In [1], Johnson and Gilles present a model of endogenous network formation in which there is a *spatial cost topology* (i.e. agents that are spaced further from each other incur a larger cost of connection). This extension to existing network formation models enables the interaction between agents to be considered, in which the distance between agents can represent geographical, social or other differences.

However, agents were fixed in position throughout the formation process. When considering dynamic complex systems, such as the economy, such static equilibrium models can have limitations.

To address this, first the spatially embedded network formation model is made dynamic and simulated so that the formation process itself can be examined in more detail. Second, the dynamic model is further adapted to enable agents to be spatially mobile. This mobility further extends the capability of the model in expressing real-world systems, in which such dynamics are often present. A cost of travel in the spatial domain can also be applied to agents, acting as a limiting factor.

This paper identifies the conditions for stable and efficient network formation, and explores the dynamic behaviour of agents over time.

The rest of this paper is organised as follows: Section 2 provides a brief discussion on the relevant background literature; Section 3 presents a model of spatial mobility in which travel may be costly; Section 4 examines the behaviour of the model;

Section 5 concludes.

II. Related Literature

The foundations of this model lie in Jackson and Wolinsky's model of endogenous network formation [7] and Johnson and Gilles' spatially-embedded extension [1].

In the Connections Model [7], agents endogenously form mutually-desirable links with other agents. An agent may benefit from both direct and indirect connections, but there is an associated cost of maintaining a link. This is captured in the agent's utility function, which is used to determine whether or not a link is worthwhile. Agents, therefore, have the ability to add, maintain or remove links.

Using the Connections Model, Jackson and Wolinsky introduce the concept of pairwise stable networks as those in which no two unconnected agents would choose to form a link and no single agent would choose to remove one. Additionally, they consider the efficiency of networks. These two concepts of stability and efficiency are fundamental to the exploration of economic networks, and feature heavily in our analysis.

It was noted that though pairwise stable states were always reachable, efficient network configurations did not necessarily form. For instance, under certain conditions star topologies, though efficient, were not pairwise stable.

A specialisation of this is the Symmetric Connections Model in which all agents and links are identical. In this model, there is a common, fixed cost of linking and each agent has the same 'intrinsic value' to offer other agents. Therefore, agents only differ in their connections and position in the network.

The abstractness and adaptability of the model enables it to be applied in many areas. However, as an equilibrium model, the dynamics of behaviour within the network or its evolution over time can only really be considered through simulation.

In [8], a dynamic implementation of the Symmetric Connections Model was explored and robustness tested against errors in agent decision-making. Dependent on the likelihood of an error being made, the theoretical predictions of [7] were found to be robust. It was also found that the time taken for a network to form can increase exponentially after the introduction of error.

The dynamic implementation highlighted the importance of simulating static models: a more fine-grained cross-section of the original model's behaviour was able to be identified, and the effect of what were previously considered design choices could be studied in more detail.

In [1], Johnson and Gilles' extend the Connections Model by introducing a spatial cost topology designed to incorporate the concept that during the formation of most real-world networks the distance between agents can also affect how worthwhile a connection is. This distance does not necessarily have to be a physical space, but can be used, for example, to represent the differences between agents.

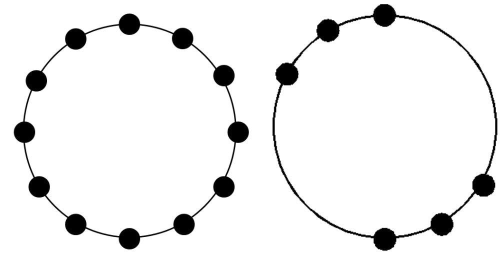

(a) 1 Neighbourhood, 100% Occupancy (b) 2 Neighbourhoods, 50% Occupancy

Fig. 1: Examples of Occupancy and Neighbourhoods

Johnson and Gilles' Spatial Connections Model situates agents on the real line, which is used to represent the social distance between them. Agents again choose which connections to add, maintain and remove dependent on a utility function, which now includes the imposed linear cost topology. They find that when the cost of a connection is high, there is a unique stable network, which is not the case when costs are low.

Spatial mobility in models with both spatial and social capital is a familiar concept in demographic models [9]. Of particular interest to simulation modellers is the order in which the processes take place. For instance, should agents be allowed to move before they connect? Should they be allowed to move as often? This combination of connection, spatial embeddedness and spatial mobility remains an interesting modelling challenge, and is a key theme of this paper.

III. The Model

A finite set of *agents* $N = \{1, 2, ..., n\}$ are situated at discrete *locations* $L = \{1, 2, ..., l\}$. No more than one agent can occupy a single location at a given time[2], and not every location is occupied.

Let *occupancy* refer to the proportion of unique locations containing an agent, and a *neighbourhood* refer to a set of agents such that for any two agents there is a path between those agents that has full occupancy (see Figure 1).

Unlike [1] in which agents are situated along the real line, a finite number of locations exist in the agent world. This is because on the real line, unless agents may occupy the same location, it is always possible to lessen the distance between two agents. This poses implementation problems when spatial mobility is introduced, since agents that wish to minimise the distance between themselves would end up stretching the computational limits of numerical accuracy. Therefore, we finitely discretise the space and introduce occupiable gaps

[2]This is a design choice and has little effect on the overall dynamics of the system.

between agents. Additionally, the measure used by Johnson and Gilles to determine the distance between two agents is the cardinality of the set of agents between them. Here, distance refers to the number of locations between two agents.

It must be noted that the spatial boundary conditions that exist in [1] result in an implicit asymmetry amongst agents. This asymmetry becomes more pronounced as the fixed cost of connection increases, causing the position of an agent to have a greater effect on its utility. Therefore, in order to better explore the effect of introducing spatial mobility, spatial boundaries are removed (i.e. locations 1 and l are neighbours, as are 1 and 2). Though this is perhaps less realistic when locations are used to represent geographical placements, it enables agent dynamics to be considered more clearly.

A *network* $g \in G$ represents the relationships between agents: $g \subseteq (N \times N)$. The binary relationship between two agents i and j is given by $ij \in \{0, 1\}$, where 1 indicates a connection and 0 indicates no connection. Note that $ii = 0 \ \forall i$. Since relationships are symmetric: $ij = ji$ and g can be shown by an undirected graph.

It is often useful to refer to the topology of the network during formation.

> A network g is **fully connected** or **complete** when forall $i\ j$, $ij \in g$.
>
> A network g is a **line** when forall $ij \in g$ there is always a path between i and j, exactly two agents have one connection, while remaining agents have exactly two connections.
>
> A network g is a **ring** when forall $ij \in g$ there is always a path between i and j and all agents have exactly two connections.

A *situation* $s \in S$ refers to which agents are at which locations: $s \subseteq (N \times L)$.

The agent *world* $w \in W$ is a collection of agents, their locations and their connections: $w \in (S \times G)$.

At initialisation in each simulation, each agent is assigned a unique random location and has no connections.

An agent may perform two types of action: spatial (1) and social (2). Therefore, we define two action sets as follows:

$$A^i_{sp} = \{maintain_s, move_l\} \quad (1)$$
$$A^i_{soc} = \{maintain_g, add_{ij}, remove_{ij}\} \quad (2)$$

where $maintain_s$ is to maintain its current situation s;
$move_l$ is to move to location l;
$maintain_g$ is to maintain its current connections g;
add_{ij} is to add the connection ij;
$remove_{ij}$ is to remove the connection ij.

When selected to act, an agent i performs a spatial action $a^i_{sp} \in A^i_{sp}$ and a series of social actions $a^i_{soc} \in \mathcal{P}(A^i_{soc})$. That is, they choose a location (spatial) and connection set (social) to maximise their utility, and then perform the actions to change the state of the world[3]. The agent's current or potential social connections incentivises its spatial decision of whether or not to move. This presents the model with two timescales: one for spatial and one for social actions.

The presence of two such timescales is often found in demographic literature [9]. It is often seen to be the case that spatial actions occur less often than social actions. In other words, a change in location is a longer process than the formation of connections. This is the approach adopted here. Hence, in our model, once selected an agent may perform one spatial action followed by a series of social actions.

Despite this difference in timescales, unselected agents cannot perform even the shorter-timescale social actions until both the spatial and social actions of the selected agent are executed. This is done to simplify the model.

An individual agent's turn terminates when it chooses to maintain its current network, rather than add or remove any connections: $maintain_g$. This assumption of inertia (that is, that satisfied agents stay put) is common in simulation literature [10]. Note that it is still possible for an agent to prefer to add a connection rather than maintain its current network, but it cannot choose this action unless the agent it wishes to connect to consents. However, an agent does not need consent to remove one of its existing connections.

Until stability is reached, an agent is selected at random to act. Note, the existing definition of stability is here extended to include spatial actions:

> A world w is **stable** if each agent would choose to both maintain its current location ($maintain_l$) and its current connections ($maintain_g$).

During simulation, only agents that have not yet reached stability may be randomly selected to act. These agents are termed textitunstable. This does not effect the outcome of the simulation, only its duration since.

The concept of efficiency is similarly extended:

> A world w is in an **efficient** state if both the spatial and social configurations maximise the total utility of all agents over the set of all possible worlds.

In reality, there may be some cost of travel c_t experienced by an agent when moving between spatial locations in the same way that there is a cost of connection when forming a connection with another agent. Similarly, this cost of travel should be proportional to the distance between locations. For example, it should be more costly to travel longer distances. This is true of many real-world scenarios involving travelling between locations; consider the difference in the cost of petrol

[3]Note that even though they 'plan' a series of actions, agents still do not use foresight: each subsequent action selected is chosen because it is a utility maximising action given that the previous actions have taken place.

when driving somewhere relatively local and somewhere far away.

In evaluating its action choices, an agent considers its utility function, which considers two states of the agent world: the current state (g', s') and the state that the agent's proposed action set (a^i_{sp} and a^i_{soc}) would bring about (g, s).

The utility to player i of social connections network g and spatial configuration s given the current network and configuration g' and s' is given by:

$$u_i(g, s|s') = \sum_{j \neq i} \delta^{t_{ij}} - c \sum_{j:ij \in g} d_{ij} - c_t \cdot d_i(s, s') \quad (3)$$

where δ is the intrinsic value of an agent;
t_{ij} is the number of links in the shortest path between i and j;
c is the base cost of maintaining a direct link;
d_{ij} is the geodesic distance between i and j;
c_t is the cost of travel;
$d_i(s, s')$ is the distance between agent i's position in s and s'.

Note: $0 \leq \delta \leq 1$ and $0 \leq c \leq 1$. This maintains the idea that the more indirect the graphical connection between agents, the less the benefit an agent receives from it.

When selected, an agent uses its utility function to evaluate the payoff it would receive by either moving to any free location on the grid or maintaining its current location, and adding, removing or maintaining any links. The agent here operates with complete information when evaluating potential connections and positionings. Since agents are self-interested and utility-maximising, the actions it chooses to execute are those that provide the greatest immediate payoff.

Simulations were run in Java using Eclipse IDE and the JUNG library.

IV. Results

An agent always seeks to minimise its cost and maximise its profit during decision making. Since cost is dependent on distance, an agent always chooses the shortest path when evaluating a connection to another agent.

We introduce the following notation:

Let H be the number of neighbourhoods;
O be the occupancy;
$d_{\min}$ be the smallest distance between two agents;
$d_{\max}$ be the largest distance between two agents.

Note: initially, the cost of travel is removed ($c_t = 0$). This is true for all cases unless otherwise stated. This greatly simplifies the model and enables a more detailed investigation into its behaviour.

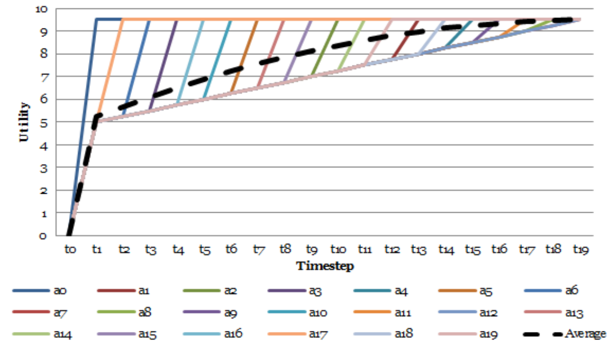

Fig. 2: Graph Showing Individual and Average Utility over time for $c < \frac{\delta - \delta^2}{d_{\max}}$ ($n = 20$, $O = 100\%$, $\delta = 0.5$, $c = 0$)

A. Consider $O = 100\%$

Under full occupancy, no agent has the opportunity to move. This is similar to a simulated Spatial Connections Model with no spatial boundaries and agents that can perform multiple social actions during a single turn.

Here, $d_{\min} = 1$ and $d_{\max} = \lfloor \frac{|L|}{2} \rfloor$. This presents five cases for investigation. However, dependent on the input parameters, some of these may not always be applicable[4].

When $c < \frac{\delta - \delta^2}{d_{\max}}$**:**

The payoff from a direct connection is always positive and preferable to the payoff from an indirect connection. Once selected to act, an agent will therefore choose to connect to all other agents.

There is only one stable network: the fully connected network. Since only agents that do not want to maintain their current network are chosen to act, stability is always reached, and takes $n - 1$ timesteps.

This unique stable network is also the unique efficient configuration.

Figure 2 shows the utility of each of the agents at each time step during a simulation. The average utility at each timestep is also shown. As can be seen, once selected, an agent's actions cause it to reach its maximum utility, simultaneously increasing the utilities of unstable agents and raising the population average. Due to the symmetry in the network and model, all agents receive the same payoff at equilibrium.

This result is fully robust to random perturbations such as the removal of links or agents, since any unstable agents will always act to restore the unique stable state, as well as changes in parameter values that are within the pre-defined limits.

When $\frac{\delta - \delta^2}{d_{\max}} < c < \frac{\delta}{d_{\max}}$**:**

[4]For example, the third case does not exist when $\delta = 0.5$ and $N = 5$.

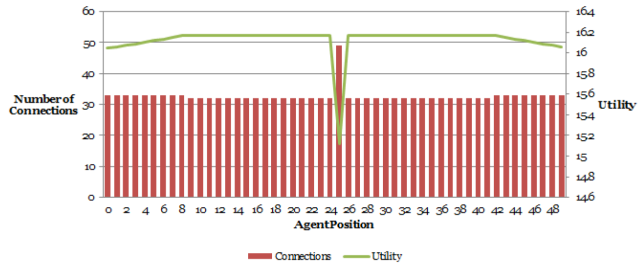

Fig. 3: Histogram showing the number of connections and final utilities of each agent. Note: the order of selection began with agents 25, 16 and 8.

The payoff from a direct connection is always positive but may not always be preferable to an indirect connection. An agent is able to afford (that is, receive a positive payoff from) directly connecting to all other agents in the world, but would prefer some indirect connections.

The first agent selected will always connect to all other agents. This is because there are no existing indirect connections it could utilise in order to minimise cost. However, this is not the case for subsequently chosen agents who, dependent on their proximity to other agents, may not need to connect to all of the agents they are not currently connected to. The result of this is that, once stability has been reached, not all agents have the same utility or the same number of connections (see Figure 3), and the first agent is always significantly worse off.

There is no longer a unique efficient configuration, since both the order of agent selection and the relative spatial positionings of the agents can be seen to affect the network formation. However, it is interesting to note that, due to the symmetry in the model, the overall value of the stable network that is formed is consistent throughout simulations.

Dependent on the order of selection, the time to stability can vary.

As would be expected, the utility of each agent over time behaves in a similar way to the previous case with the exception of the first selected agent. This agent is significantly penalised *for the duration of the simulation* (see Figure 4). The larger the number of agents, the larger this divide becomes.

Though the actions of the remaining agents do still increase the utility of each other, not all of them reach the highest obtainable utility values: a permanent imbalance in utilities has be introduced.

When $\frac{\delta}{d_{\max}} < c < \frac{\delta-\delta^2}{d_{\min}}$:

No agent can afford to directly connect to all other agents but some direct connections may still be preferable to indirect connections.

Just as before, the first agent chosen connects to all agents within an affordable range. Subsequent agents attempt to utilise existing groups of already connected agents to minimise their costs. This causes a wider spread in agent utility than before. However, as the cost of a direct connection is proportionally increasing, this becomes more pronounced (see Figures 5 and 6).

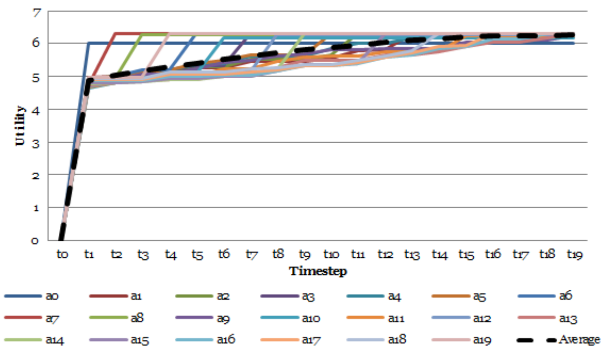

Fig. 4: Graph Showing Individual and Average Utility over time for $\frac{\delta-\delta^2}{d_{\max}} < c < \frac{\delta}{d_{\max}}$ ($n = 50$, $O = 100\%$, $\delta = 0.5$, $c = 0.035$)

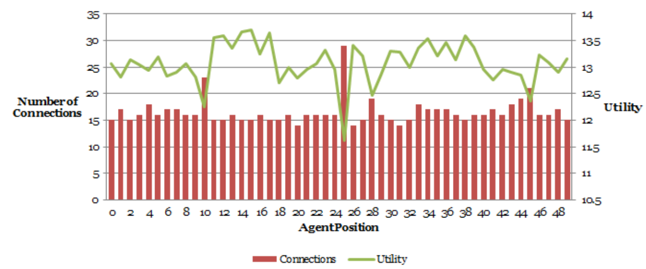

Fig. 5: Histogram showing the number of connections and final utilities of each agent. Note: the order of selection began with agents 25, 10 and 24.

It is interesting to note that, beyond the first few agents chosen to act (i.e. those that cannot utilise existing indirect connections), the order of selection does not appear to play a significant role in affecting the agents' final utilities. It is no longer the case that the first agent to act receives a lower payoff than all other agents in the final state.

As costs increase, the number of connections that an agent can maintain is reduced. Stable networks therefore begin to contain fewer links. Distant links are less desirable, so the network begins to look more ring-like in its topology. Evenly distributed ring-like structures (i.e. those in which all agents have approximately the same number of connections) are more efficient than other topologies, but the order of selection can result in stable configurations with less efficient distributions.

When $\frac{\delta-\delta^2}{d_{\min}} < c < \frac{\delta}{d_{\min}}$:

The payoff from a direct connection may be positive but not always preferable to an indirect connection. An agent always prefers indirect connections to direct connections.

This results in a ring-like network forming. This is the unique

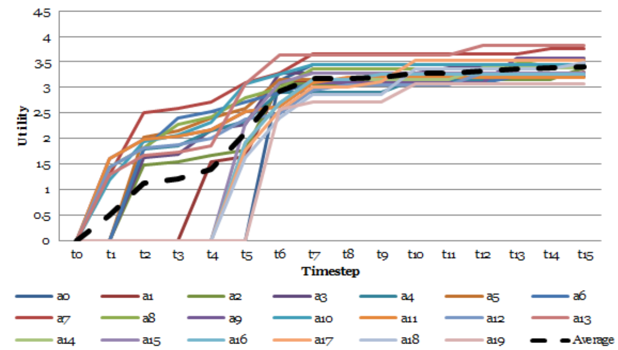

Fig. 6: Graph Showing Individual and Average Utility over time for $\frac{\delta}{d_{\max}} < c < \frac{\delta - \delta^2}{d_{\min}}$ ($n = 20$, $O = 100\%$, $\delta = 0.5$, $c = 0.15$)

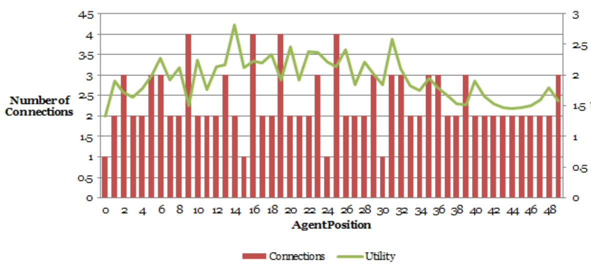

Fig. 7: Histogram showing the number of connections and final utilities of each agent. Note: the order of selection began with agents 10, 12 and 20.

efficient network configuration and also stable. However, there are additional stable configurations in which the network branches. This is because, dependent on the order of selection, it may be more profitable for an agent to make more than the two connections to its nearest neighbours.

The order of selection and relative position of agents to agents that have already acted, becomes more significant in determining an agent's utility over time. There is much more variance in an agent's utility transitions over time.

When $\frac{\delta}{d_{\min}} < c$:

Here, no agent can afford the cost of a direct connection (even to their neighbour). Therefore, the only stable state is one in which all agents are disconnected from each other. This is also an efficient configuration since all other configurations result in negative payoffs.

B. Consider $0 < O < 100\%$ and $H = 1$

As an agent always seeks to minimise its cost during decision making, and since cost is dependent on distance, an agent always chooses the shortest path when evaluating a connection to another agent: the geodesic distance. The behaviour of the systems can now be divided into two cases.

Let d_{gap} be the length of the path between two peripheral agents that passes through no other agent.

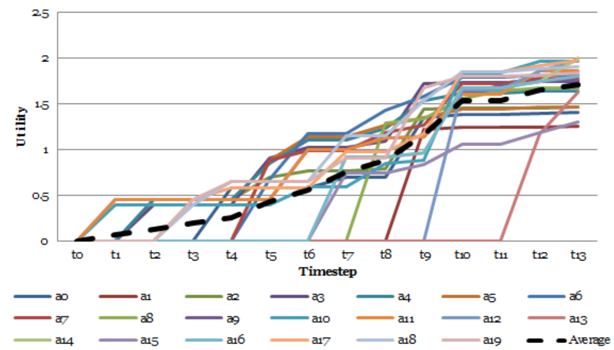

Fig. 8: Graph Showing Individual and Average Utility over time for $\frac{\delta - \delta^2}{d_{\min}} < c < \frac{\delta}{d_{\min}}$ ($n = 20$, $O = 100\%$, $\delta = 0.5$, $c = 0.3$)

When $\frac{\delta}{c} > d_{gap}$:

The largest geodesic distance between two agents is $d_{\max}$. This results in similar behaviour to the previous model, despite the potential for agents to move location.

When $\frac{\delta}{c} < d_{gap}$:

The largest geodesic distance between two agents is $n-1$ (i.e. dependent on the number of agents). This is equivalent to a bounded, linear model similar to [1].

Network Topologies

As c increases, stable networks form with fewer links. However, agents utilise existing paths within the neighbourhood by forming line-like connections along it. For smaller c, agents also create small loops along the line to reduce the number of indirect links they have. As would be expected, when $c < \frac{\delta - \delta^2}{d_{\max}}$, all links are made, and when $c > \frac{\delta}{d_{\min}}$, no links are made. Therefore, topologies favour line formations.

Agent Mobility

Since agents are already in a neighbourhood, and therefore have at least one neighbour with whom direct connection is available at the minimum cost, agent movement is unlikely, but possible. For higher occupancy values, agent movement is more likely.

Agents that move do so in order to create the best trade-off between their connections and location. This can result in them forming their own neighbourhood, which might attract other agents.

Once local connections have been made, an agent can become locked into its position: in order to maintain a larger number of cheaper, closer connections, they do not move. This effect can be detrimental to agents at the edges of neighbourhoods, since for any positive cost of connection, they are much more limited in the payoff they receive. This results in a arch shapes in utilities across the neighbourhood with perifery agents performing significantly worse than central agents. This effect becomes much more pronounced for larger n and c and

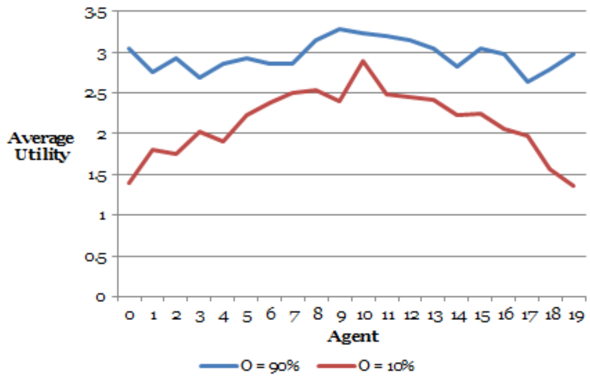

Fig. 9: Graph Showing Average Agent Utility for High and Low Occupancy Values Across 30 Runs ($n = 20$, $H = 1$, $\delta = 0.5$, $c = 0.3$)

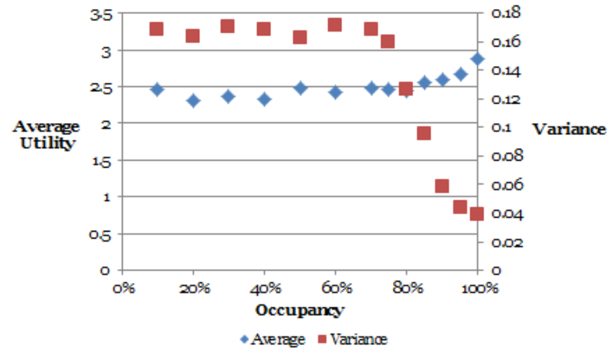

Fig. 10: Graph Showing Average and Variance in Agent Utility across Occupancy Values Across 30 Runs ($n = 20$, $H = 1$, $\delta = 0.5$, $c = 0.3$)

also makes it more likely that agents perform the majority of their spatial actions earlier in the simulation.

Efficiency

For positive c, as occupancy decreases, the efficiency of the final network decreases as well: this shows a strong positive correlation (see Figure 10). However, the spread of utility across agents increases. This pattern is not surprising:

Since there is only one neighbourhood, and due to the spatial cost topology, agents are likely to first prefer connecting to local agents. For larger occupancy values, due to the removal of spatial boundaries, agents at the peripheries of neighbourhoods may still be able to form connections with each other, even if they are unable to connect to more central agents. However, as occupancy decreases, the gap between these peripheral agents increases and so connections become less likely. Therefore, the efficiency of these agents reduces, as does the overal network efficiency. However, central agents are less affected. This pattern can be seen in Figure 9, which shows the distribution of utilities across agents for high a low occupancy values.

As soon as $\frac{\delta}{c}$ becomes less than d_{gap}, peripheral agents become unable to afford to connect to agents on the other side of the gap. This results in the network behaving like a bounded model, rather than an unbounded one. This can be seen in the variance of agents in Figure 10. In the graph, when occupancy is becomes less than 87%[5], there is a sharp increase in the variance of agent utilities, as would be expected. This eventually levels out.

C. Consider $0 < O < 100\%$ and $H > 1$:

As soon as multiple neighbourhoods are permitted, the complexity of the system increases. The time for a stable network to form can vary greatly, as can agent efficiency, the frequency of spatial actions etc. Dependent on the occupancy, the initial distribution of agents and the cost of connection, different types of network can form: the order of agent selection plays an increasingly significant role in the network formation.

However, the patterns identified in the simpler cases above are still applicable. For instance, for valid c and δ, agents will always seek the most connections they can afford, whilst minimising the cost. Therefore, agents will move to areas that are more densely populated: larger neighbourhoods can act as attractors, as do small gaps between groups of agents. An example of this can be seen in Figure 11, in which the agent world starts with six neighbourhoods and ends with two. Despite no cost of travel yet imposed, agents do not form a single neighbourhood. This is because of the lock-in effect of connections described earlier. The behavioural boundaries are still defined by d_{gap}, which is now a peripheral agent's distance to another peripheral agent from the nearest neighbourhood.

It is interesting to note that, though spatial mobility improves the overall efficiency of an agent world, there is no evidence to suggest that agents that move perform better (i.e. receive a higher payoff) than those that maintain their location. It is thought that this is due to the bidirectionality of connections.

D. Consider $c_t > 0$:

The introduction of a cost of travel acts as a limiting factor in agent mobility.

Trivially, if the cost of travel is too large, regardless of the benefit of potential connection, an agent will never change their position. This happens when, despite more beneficial locations existing, the cost of moving to them is less than the benefit. The precise threshold value of this is dependent on the specific setup (i.e. occupancy, distribution of agents, c, δ, ...).

Similarly, when the cost of travel is always less than the benefit

[5]This is the value at which $\frac{\delta}{c} < d_{gap}$.

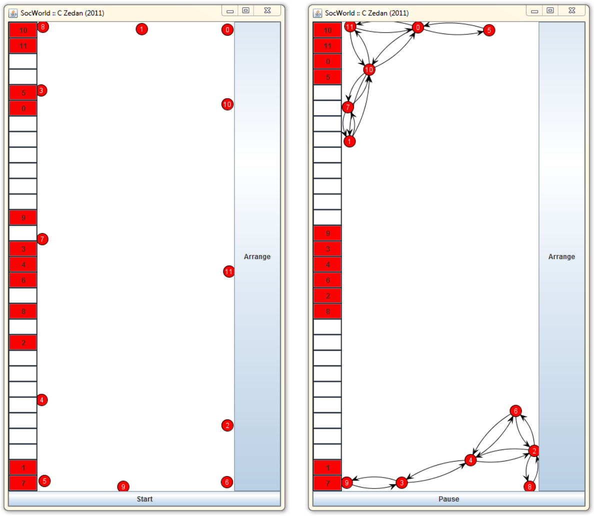

Fig. 11: Screenshots of the First (left) and Final (right) timesteps of a simulation ($n = 12$, $H = 6$, $O = 40\%$, $\delta = 0.5$, $c = 0.2$). In each window, the left panel represents the agent grid world, and the central panel represents the connections between agents.

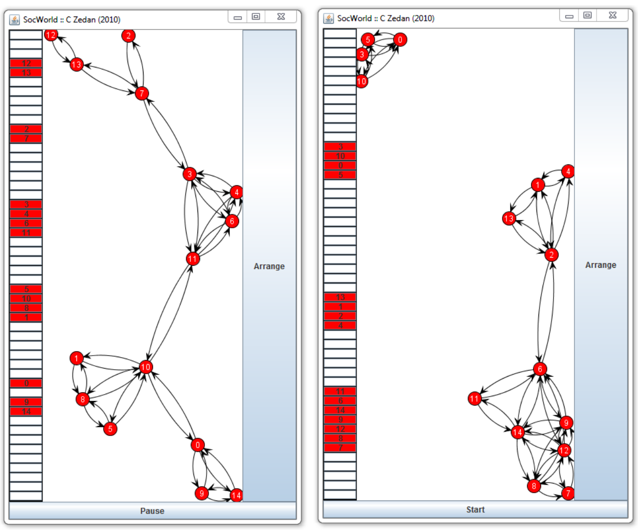

Fig. 12: Two sample simulation runs for $N = 15$, $O = 30\%$, $\delta = 0.5$, $c = 0.1$, $c_t = \frac{1}{10}$

to be gained by moving, agents are not limited in how far they can travel. The simplest cases of this is when $c_t = 0$.

For intermediate cases, dependent on the system setup some agents may be able to move to shrink the gap between them and other agents, but not necessarily remove it entirely. Similarly, it is no longer the case that agents form a single connected unit in a single neighbourhood; both c and c_t limit the networks that form (see Figure 12). Since the order of agent selection has an even greater effect that before on the network formation, the likelihood of efficient networks forming are again reduced.

E. Robustness

The results derived from all models in this paper were tested for robustness by varying initial parameter settings, such as N. Additionally, noise was injected to the system by removing existing links, adding new links, and repositioning agents at random.

It was found that the trends identified were robust to these perturbations, further supporting the findings.

V. Conclusion

We have presented a model of endogenous network formation that incorporates the concept of agents that are spatially mobile. We have also identified the states that promote efficient network formation and those that encourage rich-poor divides across agent populations.

In particular, it was found that though agent mobility improved the overall efficiency of an agent world, there was no evidence to suggest that those who moved performed better than those who maintained their location. Additionally, it was found that the order in which agents were selected to act, along with their spatial position, had the most significant effect on their individual performance.

However, there are several limitations to our model, such as the division of timescales in agent decision-making and the linearity of the spatial world. These were largely design choices and provide possible starting points for extension. Other extensions include agent heterogeneity, games on networks, games of imperfect information, signalling and indirect communication. Additionally, it would be interesting to consider the effect that agents with bounded foresight would have on the model.

References

[1] Johnson, C. and Gilles, R P. (2000), *Spatial Social Networks*, Review of Economic Design, vol. 5, pp. 273-299.

[2] HM Treasury. (2010), *A New Approach to Financial Regulation: Judgement, Focus and Stability*, The Stationery Office, ISBN 9780101787420.

[3] Financial Services Authority. (2009), *A Regulatory Response to the Global Banking Crisis*, FSA Discussion Paper, ref. 003299.

[4] Tesfatsion, L. (2006), *Agent-Based Computational Economics: A Constructive Approach to Economic Theory*, Handbook of Computational Economics, vol. 2, pp. 831-880.

[5] Goyal, S. (2009), *Connections: An Introduction to the Economics of Networks*, Princeton University Press, ISBN 9780691141183.

[6] Schweitzer, F., Fagiolo, G., Sornette, D., Vega-Redondo, F., Vespignani, A. and White, D R. (2009), *Economic Networks: The New Challenges*, Science, vol. 325, pp. 422-425.

[7] Jackson, M O. and Wolinsky, A. (1996), *A Strategic Model of Social and Economic Networks*, Journal of Economic Theory, vol. 71, pp. 44-74.

[8] Zedan, C. (2010), *The Endogenous Formation of Economic Networks*, MSc Thesis, University of Southampton.

[9] Metcalf, S. and Paich, M. (2005), *Spatial Dynamics of Social Network Evolution*, 23rd International Conference of the System Dynamics Society.

[10] Schelling, T. (1969), *Models of Segregation*, The American Economic Review, vol. 59, pp. 488-493.