
Institutional Emergence: Lessons From Two Computational Models

Greg Fisher
Institute For Complex Systems Simulation
University of Southampton
gjf1g13@soton.ac.uk

Seth Bullock
Department of Computer Science
University of Bristol
seth.bullock@bristol.ac.uk

1 Introduction

The aim of the work presented in this paper is to explore the emergence of informal institutional in economic systems; and also to consider whether formal institutions can substitute for informal institutions in such systems. Another aim was to add to the literature in computational social science and to further showcase the value of these models in the economics literature.

These aims were addressed by developing computational models of rudimentary economic systems and by employing concepts emphasised in the complexity sciences.

More specifically, we gave agents mental models (or ‘internal models’), which follows the work of Holland ([9]) and Denzau & North ([5]); we allowed agents to use inductive reasoning to make decisions (extending the work of Arthur ([3]) and Holland et al ([10])); and, in enabling the agents to interact, we allowed their mental models to co-evolve.

One of the motivators of this work was a quote from Field (in [8]) who stated that a lot of the work in New Institutional Economics has been concerned with ‘macro-level explorations of the consequences of institutional variation’ and less with ‘satisfactory micro-analytics ... that explain systematically how institutions are created, how they are sustained and why they vary.’ [8].

This quote from Field sits well with Epstein’s argument that social scientists should seek to ‘grow’ macro phenomena from the ground up (e.g., [6]).

Given the above aims and these inspirations, we sought to do the following:

First, we wanted to initially create a ‘space’ for informal institutions to emerge endogenously rather than imposing an institution exogenously, god-like, which agents would then use.

Second, we wanted the environment in which the agents existed and interacted to be a good, albeit abstract, proxy of a *whole* economic system. This emphasis on a whole system was in order to explore some of the feedback effects observed in such systems, which are often missing in partial analyses.

Third, ideally this economic system would contain parameters which could be adjusted in order to demonstrate the conditions under which informal institutions did and did not emerge.

Fourth, if informal institutions did emerge, the models should demonstrate (and, ideally, quantify) states of the world before, during, and after the emergence of such institutions.

Fifth, we had a strong preference to model two of the most important institutions in Western economies: property rights and markets.

Both of the models described in this paper include agents who spent their time foraging for resources from two ‘fountains’ in the environment (a form of work) and then spent some of their time trying to find each other to trade. The models were similar to the Sugarscape model in [7] except resources were not located geographically. In each round, agents first foraged for resources and then they attempted to trade on a grid, which represented a geographic region.

The first model was the simpler of the two: it assumed agents respected each other’s rights to the resources they either collected or received via exchange. The initial state of the model was one of uncoordinated activity. In the context of institutions, could the agents form a market? This challenge was similar to finding a Nash Equilibrium in a non-pure coordination game.

The second model was identical to the first except the agents did not necessarily respect property rights. Each agent was given a propensity to steal, which evolved with experience. The institutional challenge in this model was to see whether property rights emerged endogenously across the population.

This second model departed from conventional economics, especially rational choice theory, because the agents’ propensities to steal pre-determined (probabilistically) their strategies in interactions. This is related to a lot of empirical evidence and theoretical contributions concerning the use of heuristics in decision-making.

In discussing the second model, we will also highlight some preliminary work using neural networks as representations of agents’ internal models.

The sequencing of these models might look odd because we first look at the emergence of a market and then, in the second model, we change an assumption of the first model to focus on the emergence of property rights. This approach was taken because it will allow us to make a number of general points using a simpler model (the first), which will also be made in the context of a more complicated model (the second).

1.1 Summary of Findings

Perhaps the most important observation arising from both models was that *informal institutions emerged from the co-evolution of agents’ mental models*.

These mental models used inductive reasoning to enable the agents to make decisions but as agents interacted their mental models co-evolved in such a way that informal institutions emerged. In the first model, the institution was a market. In the second, it was respect for property rights.

The emergence of informal institutions in both models involve a reduction in uncertainty in some way and also an increase in efficiency.

In the first model, certainty was achieved when agents converged on a single market square: over time, agents’ inductive processes and the co-evolution of their mental models all led them to converge on the same grid square.

The efficiency achieved in the second model is best thought of as ‘coordination efficiency’, i.e., it is the gains which resulted from the agents being coordinated. The direct benefit of this was to increase the volume of transactions turnover to be closer to the market clearing level. The indirect benefit was division of labour among the agents, which resulted in productivity gains.

In the second model, uncertainty was reduced when the agents either fully respected property rights (or fully did not) and when each agent’s respect for property rights was consistent with its reputation within the mental models of other agents.

Efficiency was achieved in the second model when the agents came to respect property rights: direct gains included minimizing fight costs and maximizing the benefits of trade. The indirect benefits were due to the respect for property rights enabling the emergence of markets and all that entailed.

Indeed, in the context of Epstein’s emphasis on ‘growing’ macro phenomena, the second model demonstrated how the emergence of property rights enabled the emergence of markets and this in turn enabled the emergence of specialisation among the agents. This sequential emergence ultimately resulted in a productivity boom.

When we explored the parameter space and ran some experiments with the second model, we found that property rights did not emerge endogenously in some scenarios. We exploited these to consider whether different *formal* institutions might encourage agents to respect property rights. We found that in all 7 experiments there was some formal institution which achieved this; however, these results held only in the absence of corruption. If we ran the same experiments when corruption was possible, we found that in some of these scenarios it was impossible to achieve respect for property rights.

Finally, we observed that, after a market emerged, the mean price and net volume of transactions were almost identical to the market clearing price and quantity, i.e., market clearing in the general equilibrium model was approximated (albeit with only 2 resources). This result was contingent on the emergence of property rights and a market institution.

1.2 Overview of Paper

The first model is presented in Section 2 and the second model in Section 3.

In both of these sections, the model is described, null scenarios are presented and then the results from a default scenario (which contains all of the default parameters) are discussed. Both sections also contain a summary of the exploration of the parameter space and they summarise a number of experiments.

In the second model, we will explore in some more detail the work concerned with how formal institutions were used to encourage respect for property rights (both with and without corruption).

Section 4 concludes.

2 Market Emergence Model

2.1 Overview of the Model

There were two resources generated in the environment and agents had to consume *both* resources in order to survive. The objective of each agent was to stay alive, which they did by foraging from the fountains and, when it was useful, trading with other agents. In this model, agents respected property rights, i.e., agents did not steal the resources of other agents: they only traded.

2.1.1 The Environment and Time

The environment in which the agents were located was made up of 2 resource ‘fountains’ (denoted A and B), which follows Holland in [9], and a grid on which agents could trade.

At the beginning of each round both fountains were replenished with their specific resource, starting each round with L units ($L = 50$ in the default scenario¹: 2 units for each of 25 agents). All fountain resources perished in between rounds, i.e., no resources were carried through from one round to the next.

Each round was split in to two main phases: foraging and trading (Fig. 1).

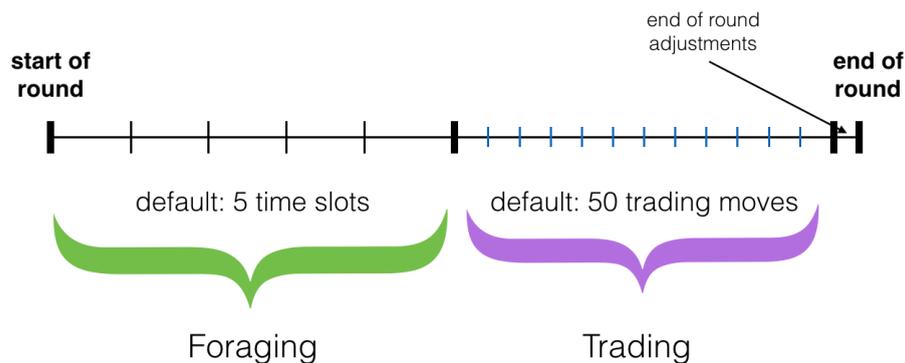


Figure 1: A Breakdown of Each Round

During the foraging phase, agents had 5 time slots in which they could visit either fountain (they could collect up to 1 resource unit per time slot).

After foraging, the agents were placed on a torus on which they could move, i.e., a grid which wrapped around top-to-bottom and side-to-side. This grid had dimensions of 50×50 .

Each agent had a home location on the grid from which it started its search for other agents during the trading phase. Each agent could see other agents on its current square and on adjacent squares, which here meant all of the eight Moore’s (or ‘King’s move’) squares. The trading phase was split in to 50 time periods: when traversing the grid during these periods, agents could only travel 1 grid square at a time.

¹For simplicity, we will specify a ‘default’ set of parameters and when we run simulations using these parameters, we will refer to this as the ‘default scenario’. The main parameters were varied when we explored the parameter space.

At the end of each round, after foraging and trading, the agents consumed the resources they held, updated their foraging strategies, and communicated with other agents. Once finished, a new round began until 1,000 rounds were completed.

2.1.2 The Agents

There were 25 agents in the default scenario. The following is a list of the agents' main state variables (individual agents are denoted by the subscript i):

- A personal resource array

$$\mathbf{r}_i = [r_i^A, r_i^B]$$

Each element of this array was a stock corresponding to each of the fountain resources (A and B). We can think of these resources as nutrients, both of which were essential for survival. The agents had to maintain a stock of each of these resources by consuming what they had collected or bought (metabolism depleted the agents' stocks). Once consumed, these resources became 'embodied', i.e., the agents could not un-consume the resources and trade them.

r_i^A and r_i^B were initialised with values drawn from a normal distribution with mean 50 and standard deviation 5.

Agent i remained alive if all of these resources exceeded 0, i.e.:

$$r_i^j > 0 \quad \text{for all } j = A, B$$

For the agents, therefore, the challenge in this model was to remain alive by sustaining a positive value in both cells in \mathbf{r}_i .

- A foraging strategy array

$$\mathbf{h}_i = [h_i^1, h_i^2, h_i^3, h_i^4, h_i^5]$$

The foraging phase of each round was divided in to 5 time slots, which meant agents had 5 discrete opportunities to forage from either fountain. This array determined which fountain each agent visited in each time slot. At inception this array was populated randomly, e.g., $\mathbf{h}_i = [A, B, B, A, B]$.

- A foraging skill levels array

$$\mathbf{p}_i = [p_i^A, p_i^B]$$

These were probabilities which corresponded to levels of skill in foraging. Specifically, each element reflected the probability that an agent would detect the particular resource associated with a fountain (this is explained further below).

At instantiation:

$$p_i^j = 0.5 \quad \text{for } j = A, B$$

These skill levels changed according to how much time the agents spent foraging for a particular resource: an agent's detection probability for j increased if it spent more time foraging for j , i.e., its skill improved, and vice versa. This process is at the heart of how agents became specialised and it will be explained further below.

- A basket array

$$\mathbf{b}_i = [b_i^A, b_i^B]$$

This array was used to keep track of the resources successfully collected during the foraging phase of each round, and which might have been subsequently traded. The resources remaining in the basket at the end of the round (post-foraging and post-trading) were consumed, i.e., they were added to \mathbf{r}_i .

- A memory array

$$\mathbf{m}_i = [m_i^1, m_i^2, m_i^3 \dots]$$

This array recorded locations where agent i had traded with other agents in previous rounds, and also locations where others had traded, which agent i had been informed about.

- A Home Location on the grid

Each agent was allocated a location on the grid from which it started the trading phase. In the Default Scenario, agents' homes were evenly spaced on the grid, as shown in Fig. 2 below.

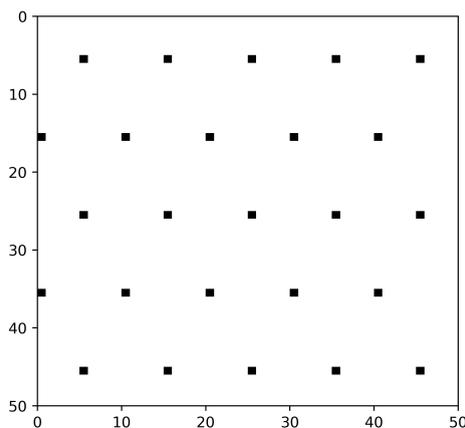


Figure 2: The location of agents' homes on the grid

An agent could sire children if both its resource holdings exceeded 125 units. If this was true of any two agents then they bore a child: 25 units of each resource were deducted from each parent's resource arrays, and a child was instantiated with a personal resource array containing 50 units of each resource. The child was given a home location in the grid square furthest away from any other agent (this maximised the sparsity of the population).

2.1.3 Foraging

The resource fountains were not given a geographic location: we assumed both fountains were within reach of all the agents. In a sense, the fountains represented work which the agents undertook in order to survive.

In the first time slot of the foraging stage, each agent visited the resource fountain specified by its foraging strategy array \mathbf{h}_i , e.g., $h_i^1 = A$. As a result, at each fountain there was typically a group of agents foraging from that fountain.

These agents formed a queue at the fountain (randomly generated), and they then took turns to forage from it. The likelihood of an agent detecting a resource when it was their turn is described below.

If an agent successfully collected one unit of a resource, this was deducted from the fountain's stock so it was not available to other agents, and it was added to the agent's basket array, \mathbf{b}_i .

After all the agents had foraged from both fountains, the 'clock' ticked forward to the second time slot. The agents moved to the fountain designated by their foraging strategy arrays, e.g., $h_i^2 = B$.

This process was repeated over the 5 time slots. After the foraging part of the round was completed, agents moved on to the trading phase.

Detecting Resources at Fountains

When foraging for resources, the success of an agent in detecting a resource was dependent on two factors: its skill in detecting that resource (determined by \mathbf{p}_i) and the resource's availability.

The simplest way to combine these two factors was to multiply Agent i 's detection probability by the remaining stock of a fountain relative to its replenished value. As a result, the (net) probability of an agent detecting resources at any fountain was:

$$n_i^j = p_i^j \frac{l^j}{L} \quad (1)$$

where:

- n_i^j was agent i 's net probability of detecting a resource unit at fountain f^j
- p_i^j was agent i 's gross probability of detecting a resource at fountain f^j , which was taken from the agent's foraging skills array, \mathbf{p}_i
- l^j was the stock of resources at fountain j at the time of foraging
- L was the stock of fountain j at the beginning of the round (50 by default)

For example, if agent i reached Fountain A when the Fountain's resource level was 40 and the agent had a gross detection skill of 0.5 for that resource then:

$$n_i^j = 0.5 \times \frac{40}{50} = 0.4$$

The agent had a 0.4 probability of detecting the resource. We can see that an agent was more likely to detect a resource unit if more resources were present at the fountain and if its detection skill was higher.

2.1.4 Trading

After the foraging was completed, agents were placed on the grid at their individual home locations, which were evenly distributed. See Fig. 2 (page 6) above.

After being placed in these home locations, each agent sought other agents to trade with.

In the default scenario, when agents had no memories, e.g., at the beginning of each simulation, they moved around the grid in a random walk. This meant they chose one of the 9 squares within range (the current square and any of the 8 adjacent squares) with equal probability. Eventually an agent interacted with others and memories of these locations were created and stored.

Alternatively, when agents had memories of previous transactions they selected a target location (from memory) to head towards from their home. An important assumption in this model was that agents used their memories of transactions in previous rounds inductively, as signals for where other agents were likely to be in the current round.

In the default scenario, agents used a ‘Roulette Wheel’ approach to choose between different locations in memory, i.e., the probability that an agent chose a specific location was proportional to its incidence in memory.

Once reached, agents remained at their target locations.

Here, the agents’ memories (along with the choice algorithm) constituted their mental models and the process by which grid targets were selected was a form of induction. Arthur ([3] referred to induction as:

... when we cannot fully reason or lack full definition of the problem, we use simple models to fill the gaps in our understanding. Such behavior is inductive. ([3], p 407)

Here, agents used what information they had to make their decision of where to move but there was no guarantee other agents would be there. In effect, the agents used ‘informed guesswork’.

If two agents did meet at the same target location, they interacted.

Transaction Prices

In order to determine the prices at which agents transacted, the model adopted the process used in [7]: prices were the geometric mean of agents’ marginal rates of substitution (MRSs). An MRS was simply the ratio of one resource holding (including the resources in an agent’s basket) relative to another, e.g.,

$$MRS_i^{AB} = \frac{(r_i^A + b_i^A)}{(r_i^B + b_i^B)}$$

When two agents, i and k , interacted their MRSs were compared and, if it was beneficial, the agents traded.

The price at which the agents would trade resources A and B was:

$$Price_{AB} = \left[\frac{(r_i^A + b_i^A)}{(r_i^B + b_i^B)} \times \frac{(r_k^A + b_k^A)}{(r_k^B + b_k^B)} \right]^{1/2} \quad (2)$$

The general idea with this approach to pricing was that if agent i had, say, significantly more of resource A than resource B, then MRS_i^{AB} would be relatively high. Suppose this agent then met agent k , with significantly more of resource B than A (MRS_k^{AB} was relatively low). The agreed price would be somewhere in between MRS_i^{AB} and MRS_k^{AB} : the agents would ‘meet’ in the (geometric) middle, sharing a ‘consumer surplus’. The trade would be beneficial because both agents would sell a resource they had more of and receive a resource they had less of.

2.1.5 End of Round Adjustments and Communications

Each agent added the resources in its foraging basket (\mathbf{b}_i), which remained after foraging and trading, to its personal resources (\mathbf{r}_i), and a metabolism cost of 1 unit for each resource was deducted:

$$\mathbf{r}_i^{end} = \mathbf{r}_i^{start} + \mathbf{b}_i - \mathbf{1}$$

Here, \mathbf{r}_i^{end} refers to the personal resource array at the end of the round, and \mathbf{r}_i^{start} at the beginning. The deduction of $\mathbf{1}$ represents the cost of metabolism.

Two other agent arrays were updated at the end of the round: the agents’ foraging skills (\mathbf{p}_i) and agents’ foraging strategies (\mathbf{h}_i).

Updating Foraging Skills

The agents’ resource detection arrays (\mathbf{p}_i) were updated to reflect the agents’ foraging in the round: if an agent spent a lot of time foraging for a specific resource, its skill (detection probability) would increase, and vice versa.

An adapted logistic equation was used, which allowed the 2 detection probabilities to vary between a minimum level of skill (\underline{p}) and a maximum level (\bar{p}):

$$\Delta p_i^j = \frac{t \cdot [w_i^j - d/x] \cdot (p_i^j - \underline{p})(\bar{p} - p_i^j)}{\bar{p} - \underline{p}} \quad (3)$$

There are three parts to this adjusted logistic equation:

1. t is a speed-of-adjustment variable
2. $w_i^j - d/x$ is a term which ensured the agent’s change in skill was positive if the agent spent more time than average foraging for a specific resource: w_i^j is the total number of time slots agent i spent foraging for resource j during the round; d is the total number of time slots; and x is the number of resources fountains. In the default scenario, $d = 5$ and $x = 2$. This meant that, in the default scenario, if an agent spent more than 2.5 of its 5 time slots foraging for resource j then its skill increased, and vice versa
3. $\frac{(p_i^j - \underline{p})(\bar{p} - p_i^j)}{\bar{p} - \underline{p}}$ ensured the adjusted logistic equation had a minimum value of (\underline{p}) and a maximum value of (\bar{p}). In the default scenario, agents’ skill probabilities had a floor of 0.2 and a ceiling of 1.0, i.e., the agents could attain skill ‘perfection’ and they always had some nominal level of skill even if they stopped foraging for a particular resource

Updating Foraging Strategies

In updating their strategies, agents were allowed to change one of the cells in their foraging strategy arrays (h_i). In doing this, agents sought to maximize the acquisition of the resource in which they were most lacking in their personal resource arrays, i.e., if they held less of Resource A than B then they would seek to increase their holdings of A in the next round.

In general, agents had two choices: (1) they could choose simply to forage for their lowest resource; or (2) they could choose to forage for the resource in which they were most skilled in foraging and then attempt to trade in to this lower resource (if different) in the next round.

Suppose, for example, an agent held fewer of Resource A and it was deciding whether to change h_i^2 from B to A. Suppose, also, that its detection probability for A was 0.3 and its detection of probability for B was 0.7. The agent has a conundrum: it could decide to forage from B (with a higher expected yield), hoping to trade from B to A; or it could forage from A.

What the agent decided depended on its detection probabilities and its expectation of being able to trade. If it believed the probability of trading was low then it would be less likely to forage for its specialised resource, and vice versa.

One addition made to this process was to add an error term to reflect the idea that agents did not have perfect information, nor perfect information processing power. Agents accurately estimated their expected yields for both resource choices and these were adjusted by a value drawn from a normal distribution with mean 0 and standard deviation of 0.1.

This updating of foraging strategies is central to whether agents became specialised: we found that once markets formed and the probability of trading increased toward 1 then agents specialised, and vice versa. We found that specialisation occurred when an agent entered a positive feedback loop: they foraged for the resource they were most skilled in foraging for, which typically increased the agent's skill for that resource, making it more likely the agent would forage for it when adjusting its foraging strategy.

Communications

At the end of each round, agents communicated with each other: there was a 1% probability that any pair of agents would communicate. If two agents did communicate, they shared transaction locations from the previous round with each other.

After these end-of-round processes were completed, the next round began until the 1,000th round was finished.

2.2 Results: Null Scenario - No Memories with Specialisation

Before showing the results from the default scenario, two null experiments are presented: one in which agents moved around the grid randomly, i.e., they were prevented from using any memories to help them find other agents (they could specialise); and a second in which agents used memories to decide on a target location but where agents were not allowed to specialise (their detection probabilities remained at 0.5). Here we look at the results from the first null scenario.

Fig. 3 below shows a heat-map of transactions in each grid square in the last 100 rounds of a typical simulation. It shows a random spread of transactions.

In this scenario, all transactions resulted from agents randomly bumping in to each other.

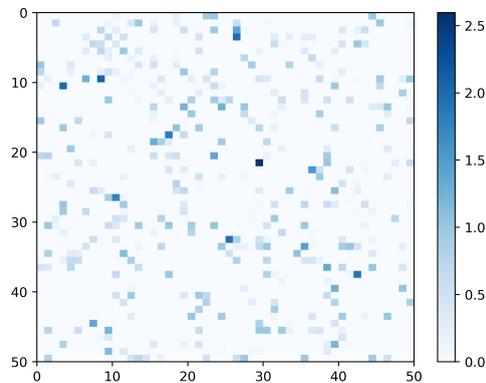


Figure 3: Heat-map of transactions during a typical simulation (last 100 rounds)

A useful representation of the inefficiency of this null scenario is to compare the actual volume of transactions on the grid with a representation of the market clearing volume of trade. Fig. 4 below shows the supply and demand curves for Resource A in a typical round in a typical simulation; and it shows the actual volume of transactions and the mean price (the red star).

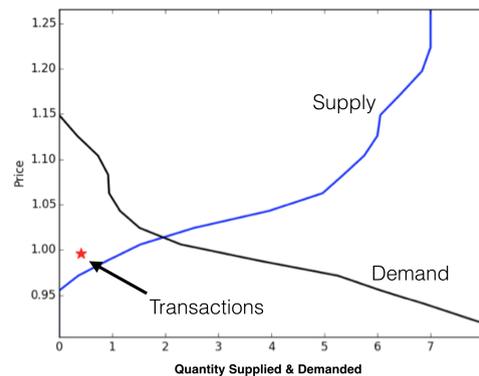


Figure 4: Supply and Demand Curves and Total Transactions (First Null Scenario)

Notes: The supply and demand curves were calculated from data in a typical round of one simulation, after foraging and before trading. The red star indicates the total volume (measured on the horizontal axis) and the mean price (measured on the vertical axis) of the transactions in the same round. Supply and demand curves were generated by calculating what each agent would have supplied or demanded for each resource over a range of prices; and then aggregating these over all agents.

On average we found that the volume of transactions was about 18% of the market clearing volume. This meant agents were left unsatiated: many wanted to trade but they did not because they did not interact.

In these simulations, none of the agents specialised: detection probabilities remained at about 0.5 for all agents. If we extended the simulations to 5,000 runs, we found the agent population fell to approx. 19-20 agents on average where it remained stable.

2.3 Results: Null Scenario - Memories without Specialisation

If agents used their memories of locations in order to choose a target location, we found that, over time, they converged on a single grid square where all transactions took place.

Fig. 5 below shows the heat-map of transactions during the last 100 rounds of a typical simulation. It is the equivalent of Fig. 3 above: here, 99.9% of transactions took place on a single grid square.

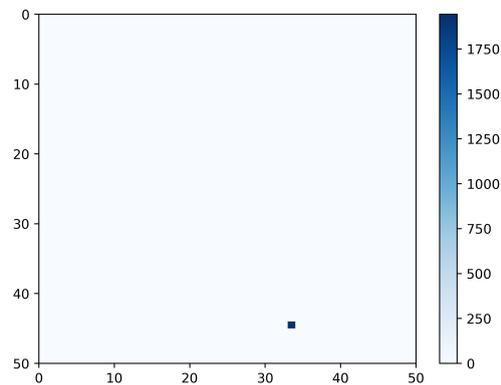


Figure 5: Heat-map of transactions during a typical simulation (last 100 rounds) - Second Null Scenario

A single market typically emerged over approx. 100 rounds: initially, several markets emerged as agents bumped in to each other and then (in some cases) reported these interactions to other agents. There was, however, ‘symmetry breaking’, which meant larger markets dominated smaller markets until only one market was left.

It is noteworthy, also, that markets emerged in different locations in different simulations: their exact locations were never known beforehand. It was an emergent property of the whole system.

Fig. 6 below is equivalent to Fig. 4 above: it shows how the existence of a market moved the actual volume of transactions much closer to the market clearing volume.

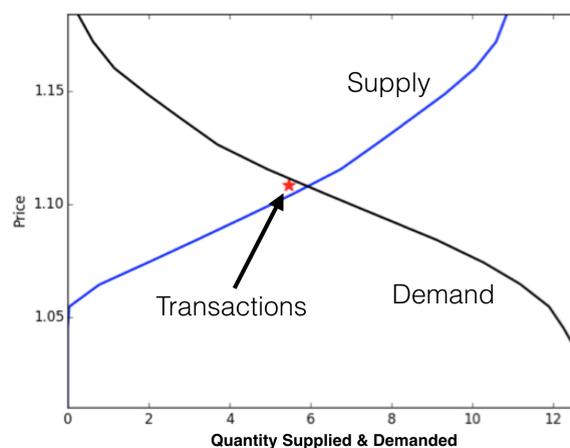


Figure 6: Supply and demand curves and total transactions (Second Null Scenario)

Fig. 7 below shows a time series of the actual volume divided by the market clearing volume of transactions (denoted here as the ‘turnover ratio’). We can see how this ratio rose from zero to approx. 1 over about 200 rounds.

Agents were not permitted to specialise in this scenario. This meant that detection probabilities were kept constant at 0.5. We found that, even though agents could transact efficiently, the overall foraging

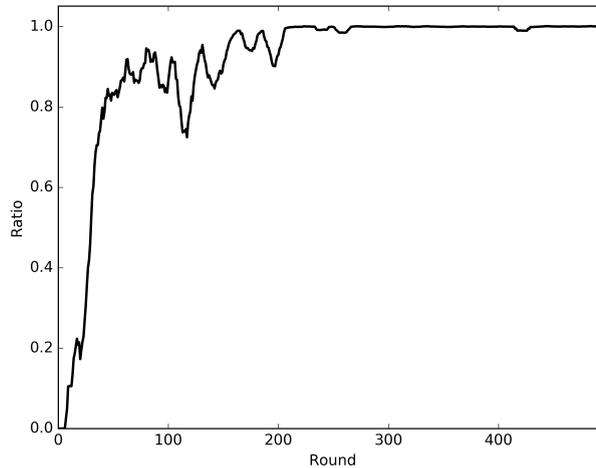


Figure 7: Turnover of Resources as a Proportion of Market Clearing Turnover

yield was not sufficient for all 25 agents to survive: if we extended the simulations to 5,000 rounds we found the population declined to approx. 18 agents on average where it remained stable.

We can appreciate that the market square represented a form of coordination: eventually, the agents only had 1 location in memory, which they kept returning to. This is equivalent to a market: it is a point in space where agents met to exchange resources.

2.4 Results: Default Scenario - Memories with Specialisation

Here, agents used their memories to decide on target locations and they were allowed to specialise.

The results were similar to previous scenario in that a market emerged but, this time, most agents became fully specialised.

We found the agent population increased to approx. 43 agents on average. Typically, approximately 37 agents of these reached full specialisation whereby they foraged from the same fountain in all 5 foraging time slots and their maximum gross detection probability reached 1.0. In a steady state there were also typically 5-6 agents that failed to reach full specialisation: these agents constantly died and they were eventually replaced with new, unskilled agents.

A measure which helps us visualise the move from generalists to specialists is shown in Fig. 8 below, which resulted from a typical simulation. The chart shows ‘the mean specialisation value’ over all agents during the first 500 rounds of one simulation (a value of 3 meant perfect generalist and 5 meant perfect specialists). In this run, most agents had specialised by approximately Round 200. The oscillations around 4.5 were because about 85% of agents achieved a mean specialisation value of 5 and approx. 15% of agents had mean specialisation values of 3 or 4.

This move from being generalists to specialists was due to the existence of a market, which increased the agents’ expected probabilities of transacting. When they updated their foraging strategy arrays, agents were generally encouraged to select the fountain associated with their highest detection probability. In a sense, when agents were confident of trading, they viewed the two resources as interchangeable, which meant they chose to forage for the resource with the highest expected yield.

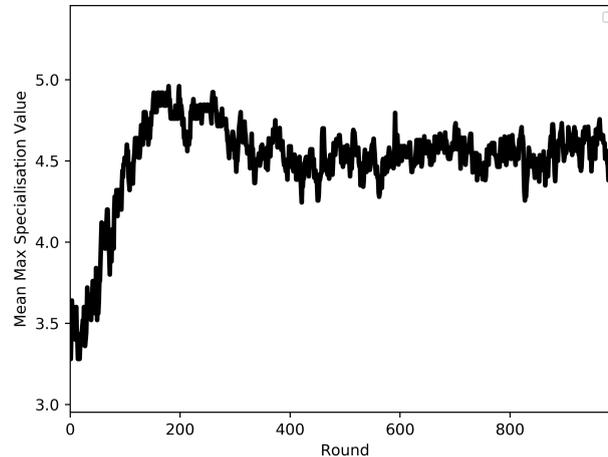


Figure 8: Mean specialisation value in a typical run.

Markets enabled specialisation in this model.

It is noteworthy that, prior to specialisation, agents met only to exchange marginal quantities of resources: as generalists, they foraged from both fountains and there was little need to transact (but not none). As agents became specialised, they foraged for a single resource but they needed both resources to survive, which meant exchange was necessary. As a result, the total average volume traded in the first 100 rounds of the simulations was 7.8 units whereas it was 20.5 in the last 100 rounds.

2.5 Exploration of the Parameter Space

Sixteen parameters were varied in order to explore: (1) their impact on the above results; and (2) the conditions under which markets emerged and when they did not.

In summary:

- The agents' memory length had to be at least 2 rounds for a market to emerge: if the agents had no memory or a short memory (of 1 round), a market did not emerge. Interestingly, a longer memory meant symmetry breaking was prevented: several 'local' markets emerged. Agents remained faithful to the first market they visited.
- If the population density was too low then no markets emerged. In the default simulations there were 10^2 grid squares per agent on average - if this measure was greater than approx. 60^2 grid squares per agent then the sparsity of the population typically prevented any agents from meeting.
- Communication was catalytic for market emergence: if agents were prevented from communicating then it took several thousand rounds for sustained markets to emerge.
- In the default scenario agents were prevented from transacting on the way to their target location: if we changed this assumption, markets still emerged, albeit with smaller 'satellite' markets close to a main market.

- In the logistic equation above (Equation 3, p 9), t represented the speed at which foraging skills changed. Markets emerged regardless of t : for very low values, the population would first decline to a lower, sustainable number of agents and, eventually, these agents would become fully specialised, resulting in new agents being born.
- If agents' home locations were randomised instead of being spread evenly across the grid as in Fig. 2 then markets tended to emerge more easily. When agents were located evenly across the grid, the distance between agents was maximized, which minimized the likelihood of them bumping in to each other. Randomizing home locations reduced the mean distance between agents, making it easier for them to find each other.
- In estimating future foraging yields when agents updated their foraging strategies, agents had imperfect intelligence and knowledge. As mentioned above, a value drawn from a normal distribution with mean 0 and standard deviation 0.1 was added to any estimate. If we increased the standard deviation of this error, markets emerged; however, when the standard deviation was above approx. 0.3, agents failed to specialise.
- In the default scenarios, agents could traverse the whole grid: they could travel for up to 25 grid squares and wait at their target grid square for whatever remained of the 50 time periods. When this travel distance was reduced to below 12 grid squares, agents struggled to find each other and to then specialise. Above (or equal to) 12 grid squares, typically 2-3 'local' markets emerged and most agents specialised.
- If we located the resource fountains on the grid and made the agents start the trading phase at the last fountain they visited, it was easier for agents to form markets: the agents started this phase among other agents they could trade with. Typically, two markets emerged (at the resource fountain locations) but, eventually, one market dominated. If all the agents had fully specialised, each started the trading phase among agents holding the same resources, so they could not trade. All the agents had to converge on the same grid square before trading could commence.

2.6 Experimentation

One of the useful characteristics of computational models is that we can ask what-if questions and develop experiments to answer these questions. For this first model, 6 experiments were conducted:

1. Famine Conditions: default simulations were run for 3,000 rounds but between rounds 1,000 and 2,000 the resources available to the agents were halved. In one experiment this was done for Resource A only and for the second experiment it was done for both resources. A market was sustained during both famines and it helped to mitigate the impact of the famines.
2. Attempting to Move the Market. What was required to move the market once it had emerged, i.e., to change an institution? An analysis of the data showed the markets acted like 'attractor basins' but they were unusual in that moving a market required a collective effort: no single agent could do this alone.
3. Engineering Local Markets. In this experiment the agents' travel distances were reduced and a policy maker attempted to ensure all agents had access to a market. A coordination failure existed when local markets emerged spontaneously: if a policy maker wanted all the agents to have access to a market then this required system-wide awareness and policy influence.

4. Fixing Prices. In this experiment the agents were forced to trade at a price of 1, i.e., the resources were valued equally. This quickly eliminated any trade and no markets emerged. Looking at the data, agents ‘wanted’ to trade in a narrow price range and this range varied between rounds with perturbations in aggregate foraging yields. Agents would not trade outside of this range.
5. Fixing Prices at the Market Clearing Price. This was a ‘Walrasian Auctioneer’ experiment. Markets formed and agents specialised in a way similar to the default simulations above. In general, agents were happy to trade at this price (which evolved between rounds as foraging yields changed).
6. Developing a Market Constitution. Agents were allowed to vote for 5 different constitutions, which determined the ‘opening time’ of the market and the prices at which agents traded. The first 4 constitutions are shown in Table 1 below. Opening times corresponded to the time period during the trading phase when agents could transact. Agents also had the choice of floating prices (bilaterally agreed between agents) or fixed prices (at the market clearing price). The fifth constitution was the full ‘Walrasian Auctioneer’ outcome where the agents traded at the market clearing price and there was full market clearing, i.e., agents bought and sold everything they wanted to. In this fifth constitution, there was no trading phase: the market automatically cleared after foraging.

Table 1: The First Four Constitutions

	Opening Time = 0	Opening Time = 20
Floating Prices	1	2
Market Clearing Prices	3	4

At first blush, one might expect the fifth constitution to be the overwhelming favourite because it should be the most efficient; however, this proved not to be the case.

In constructing the experiment the default model was run for 1,000 rounds and then each constitution was run for 1,000 rounds. In order to make an unbiased comparison, only the agents who were alive during all 6,000 rounds were allowed to vote: each voted based on their resource accumulation during each constitutional period.

This meant that only skilled, mature agents voted: any agent which was born or died after Round 1,000 could not vote. Mature, skilled agents generally preferred bilateral price negotiations and they were indifferent between opening times, i.e., they preferred the first two constitutions over all others.

An analysis of the data showed that bilateral pricing allowed skilled, wealthy agents to exploit unskilled, poor agents when they were close to death. At this time, these poorer agents would often accept extreme prices from other agents because their marginal rates of substitution were also at extremes. This benefited the wealthier agents. When the market clearing price was used, poor agents with very high / low marginal rates of substitution hardly influenced the price, so they were less exploited. As a result, the voting agents preferred bilateral pricing.

This experiment approximated the use of ‘power’ in agreeing formal institutions: here, wealthy agents voted for a constitution which suited them.

2.7 Discussion: Mental Models, Induction, Uncertainty and Efficiency

The emergence of a market in the default simulations was equivalent to finding solutions in (non-pure) coordination games so in this one sense the outcome of this first model was not new.

More important were the mechanisms by which the agents achieved this coordinated outcome. When we analysed the data generated by the model we found that while agents used the same inductive *process* to choose a target location on the grid, the *content* of their mental models (here, memories) evolved over time. Importantly, this evolution of individual agents' memories occurred in response to interacting with other agents both on the grid and at the end of each round. We can say that the agents' mental models *co-evolved* and this co-evolution gave rise to the emergence of markets.

Two important concepts in North's work (e.g., [13]) were uncertainty and efficiency. His view was that institutions mitigated uncertainty and they improved efficiency by reducing transactions costs.

In the model discussed above, it was clear that uncertainty was mitigated by the existence of a market: eventually, all the agents learned to visit one particular square on the grid. Moreover, bilateral interactions were relatively simple: agents either traded or they did not so there was no property rights institution required in interactions.

In terms of efficiency, it is perhaps more useful in the context of the model to consider the gains from coordination (direct and indirect) versus a state of non-coordination.

A useful quantification of this efficiency is to contrast the turnover ratio (recall this is the quantity of transactions / market clearing volume) when there was no market versus when there was a market. Fig. 4 above shows the supply and demand curves and the mean price / quantity traded in a typical simulation when agents walked around the grid randomly. The turnover ratio was approx. 18% in such conditions.

Fig. 6 shows the same data for one of the default simulations when a market existed. On average, the turnover ratio was very close to 100%.

In addition to the gains arising from coordination, markets also enabled the division of labour among the agents. There was no intent among the agents to create system-wide specialisation, rather it was an unintended consequence of the agents attempting to maximize their chance of living.

This division of labour led to productivity boom. In the model this led, ultimately, to a rise in the population (to a carrying capacity of approx. 43 agents). We can imagine that if we changed how agents bore children, this increased productivity could instead have given rise to agents who were wealthier on a per capita basis.

3 Property Rights Model

In the first model above we assumed agents respected property rights. Here we relax that assumption to explore whether respect for property rights might emerge endogenously among the agents.

The main change to the model concerned the agents' interaction on the grid: when two agents interacted, each agent had the choice of either attempting to steal the other's resources (what it held in its basket) or attempting to trade. Furthermore, if an agent wanted to trade but its counterpart wanted to steal then it had the choice of either acquiescing (giving its resources to the thief) or defending its resources.

Two new state variables were introduced in to the model, for each agent: a propensity to steal (denoted P_i^S) and a propensity to fight back (denoted P_i^{FB}). These were both probabilities which, respectfully, pre-determined the likelihood that any agent would attempt theft in an interaction, and defend its resources when it wanted to trade and the counterpart attempted theft.

The 'game' between two interacting agents is represented in Fig. 9 below.

Both agents first decided whether to trade or steal. If one agents wanted to steal and the other to trade, the latter then decided if it fought back (defend its resources) or acquiesced (gave all its resources without a fight).

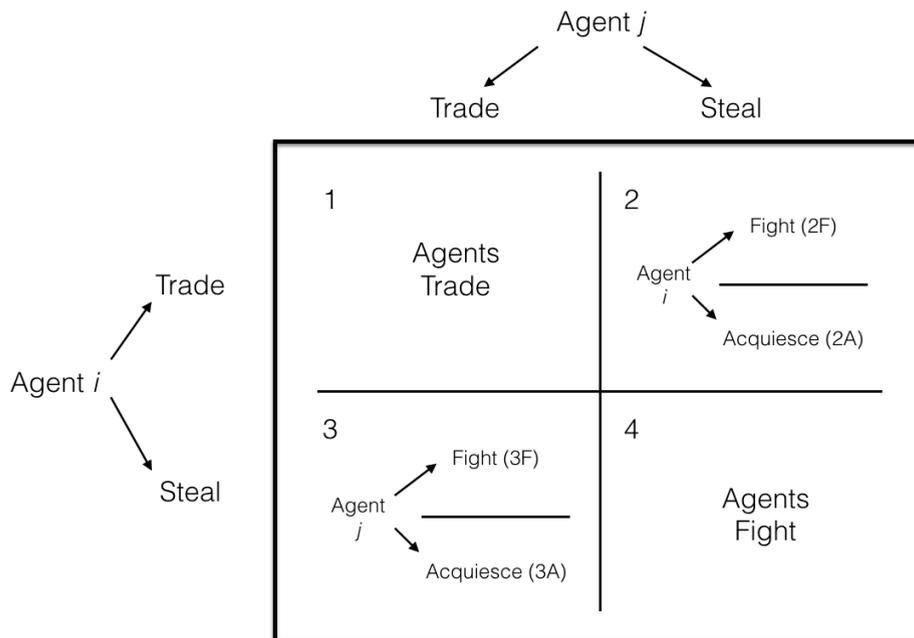


Figure 9: Overview of Agent Interaction

At first blush, this game looks similar to the stag-hunt game; however, this game structure was seen in only 0.1% of interactions over 20 simulations. In fact, there were 51 game types seen, determined by the resources held in the agents' baskets, which defined the pay-offs.

The idea that probabilities determined whether an agent stole or fought back stands in stark contrast to the rational agent model in conventional game theory, when agents would consider each interaction in isolation and determine their wealth-maximizing strategy. The approach here is akin to using heuristics (albeit probabilistic) to determine strategies. Moreover, for now we assume that agents

cannot differentiate between game types: they apply the same heuristics regardless of pay-offs. This links to Murray Gell-Mann's reference to coarse-graining in human cognition, i.e., agents cannot discern between interaction types.

We will see below some preliminary work using neural networks, which were employed by the agents to derive their propensity to steal and fight back in different game types. In bilateral interactions, the game's pay-offs were used as inputs in to the neural network: probabilistic heuristics were generated by these networks. Neural networks have the advantage of allowing agents to discern between game types. We will see that agents who used neural networks behaved almost identically to those using the simpler approach outlined in this section.

In the interaction architecture shown in Fig. 9, if both agents wished to trade then they did so in the same way as in the first model. If both agents attempted theft then they fought: one of the agents won the fight (determined by the flip of a coin) and received all of the other agent's resources. However, both agents incurred a fight cost (equivalent to an injury).

If one agent attempted theft and the other wanted to trade, the latter then decided whether to fight or acquiesce. If it fought back then both agents fought as described in the preceding paragraph. If the agent acquiesced, it gave all its resources to the other agent and neither agent incurred a fight cost.

The agents updated their propensities to steal and fight back according to their experiences with different strategies. Simply put, if a strategy proved advantageous to the agent (e.g., stealing), it was more likely to do this in the future; and vice versa. Logistic equations were again used to adjust these variables, with a minimum of 0 and maximum of 1.

This process in which agents updated their propensities was inductive in nature. Recall the quote from Arthur ([3]) on page 8. Here, adjusting probabilistic heuristics was the simple approach taken by the agents to inform decisions in a complex environment.

Another adjustment made to the model was the creation of a 'reputations architecture' whereby agents recorded reputation information about other agents. This information was used by the agents to estimate other agents' propensities to steal and fight back. Agents gained information through their own interactions and through communicating with other agents at the end of each round.

When agents moved around the grid in the second part of each round, they used the reputations information they had in addition to observations of others' resource holdings to judge the likely outcome of interacting with other agents. In general, agents moved away from others if they were likely to lose out, and moved toward or interacted with agents if they judged it would be beneficial. An error term was added to these estimates to reflect imperfect information and cognition.

A few other adjustments were made to the model: the agents' memories were increased from 4 rounds to 20 in the default scenario (this was in order to retain more reputation data); agents' resources at instantiation were increased to 300 units (the cost of fighting meant agents' resources were depleted much more quickly than in the first model); and the probability that any two agents communicated at the end of the round was increased to 0.25.

We will now turn to look at the results of three null scenarios. Following this we will examine the results of the new default scenario.

3.1 Null Scenario: No Theft

This model essentially replicated the original model because agents' propensities to steal were fixed at zero, i.e., the agents respected property rights. Propensities to fight back were irrelevant because agents only ever wanted to trade.

The results were essentially the same as the original default scenario. The most notable difference was the nature of the markets which emerged: this time, the markets covered an area on the grid rather than a single point; sometimes there were multiple markets; and these markets tended to migrate slowly over time. Fig. 10 below shows a heat-map of transactions over 100 rounds in a typical simulation.

Otherwise, the turnover ratio was close to 1 by the end of each simulation and all agents specialised. The carrying capacity of the environment was approx. 41-42 agents (see Fig. 11 below, which was from a simulation run over 5,000 rounds).

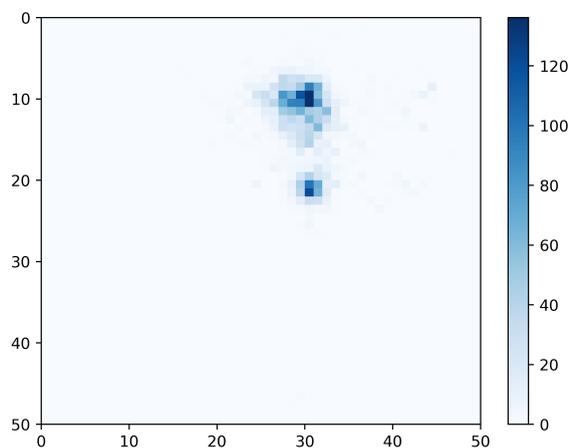


Figure 10: Heat-map of Transactions: First Null Model

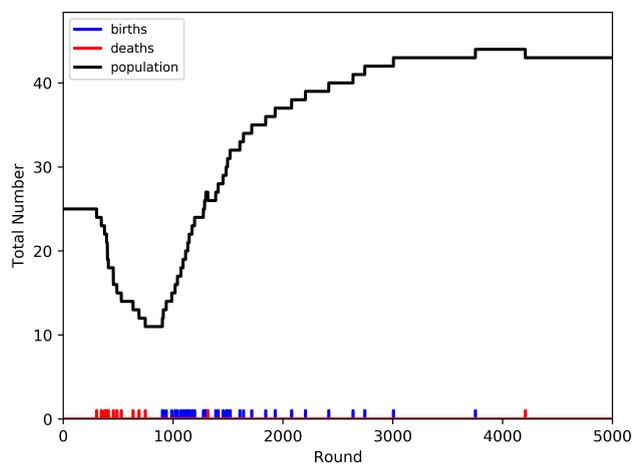


Figure 11: Agent Population: First Null Model

3.2 Null Scenario: Only Theft

In this simulation we fixed all the propensities to steal at 1, i.e., agents only every attempted theft, so they fought continuously.

The result was as one would expect: no markets formed (since there was no trading), none of the agents specialised, and most the agents died (under these conditions, the carrying capacity of the environment was about 2-3 agents). Fig. 12 below shows the agent population over 2,000 rounds in a typical simulation.

Similar results were seen when we fixed the agents' propensities to steal and fight back at 0.5, although the agents took longer to die and the carrying capacity of the environment was slightly higher at 3-4 agents.

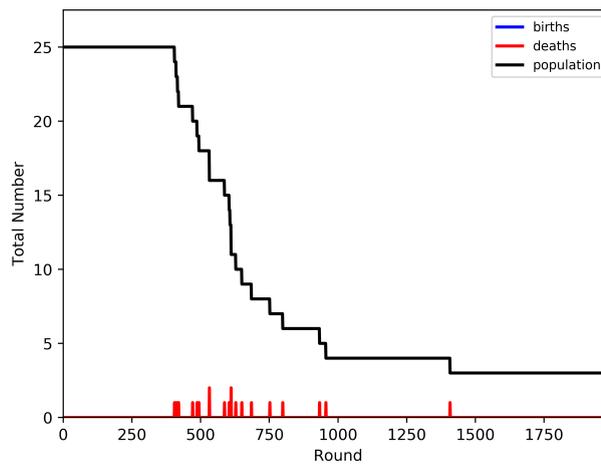


Figure 12: Agent Population: Second Null Model

3.3 Null Scenario: Classical Rational Agents

There were three ways in which we could apply rational choice theory to the model (instead of taking a 'probabilistic heuristics' approach). The first was to make all agents fight back in scenarios 2 and 3 in Fig. 9; the second was to make all agents acquiesce in these scenarios; and the third was to allow agents to choose between fighting back and acquiescing depending on expected gains / losses. In all three cases, the agent population collapsed, and the agents never specialised.

In the first case, when agents always fought back in scenarios 2 and 3, the agent population fell to 1-2 agents, none of whom specialised. Agents traded in only 0.04% of 385,000 interactions, otherwise they fought. The cost of fighting meant the carrying capacity of the environment was very low.

In the second case, the cost of fighting was lessened because agents acquiesced in scenarios 2 and 3. This meant the carrying capacity of the environment was higher: typically the agent population fell to approx. 7 agents, none of whom specialised.

In the third case, most agents in scenarios 2 and 3 chose to fight back unless they had no resources, in which case they acquiesced (there were few of these interactions, however, since agents tended not to interact with agents holding no resources). As a consequence, the results were similar to the first case: the population fell to approx. 2 agents and there was no specialisation.

3.4 Default Scenario: Allowing Agent Propensities to Vary

Here, agents' propensities to steal and fight back were instantiated with a mean of 0.5 and a standard deviation of 0.1. These were then allowed to evolve as described above.

Fig. 13 below shows the evolution of agents' propensities to steal over the first 500 rounds of a typical simulation (each line represents the propensity to steal of a single agent at the end of each round), Fig. 14 shows the agents' propensities to fight back, and Fig. 15 shows the mean propensities of agents who were alive throughout the simulation.

We can observe four patterns in the data:

1. Generally speaking, the agents' propensities to fight back increased toward 1 over time (Fig. 14).
2. After approximately Round 100, the agents' propensities to steal either increased toward 1 or fell to zero (Fig. 13). In the figure below, 13 agents saw their propensities to steal increase to 1 and 12 declined to zero.
3. For most of the agents whose propensities to steal declined toward zero, their propensities to steal initially increased (on average) before declining. This is consistent with Fig. 15.
4. All of the agents whose propensities to steal increased toward 1 died by Round 400. By contrast, all 12 of the agents whose propensities to steal declined to zero survived.

These patterns were observed in all 20 simulations of the set.

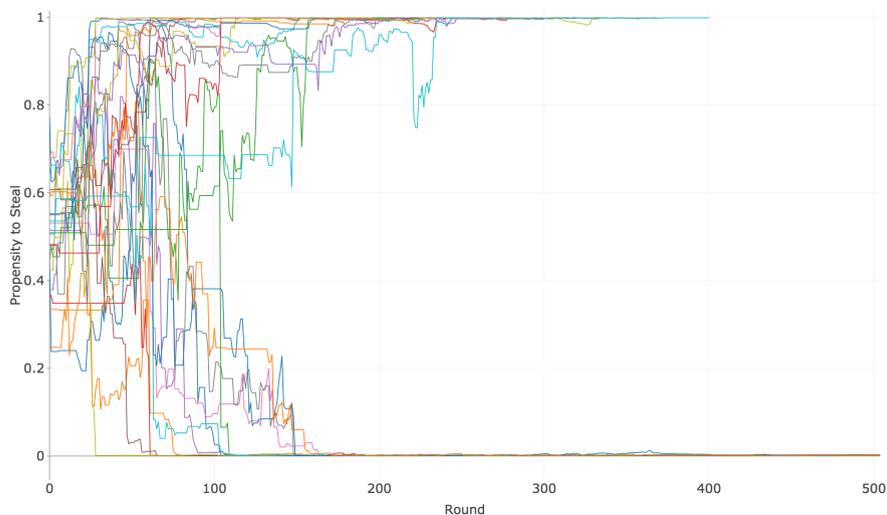


Figure 13: Agents' Propensities to Steal

The end result was a group of agents who, in general, did not try to steal from each other but who fought back when theft was attempted. Put another way, the agents came to respect property rights and, ultimately, they were likely to defend the resources they collected.

After the agents with high propensities to steal died and the surviving agents' propensities declined to approx. zero, we saw that markets emerged, the agents specialised, and the population increased to the carrying capacity of approx. 43 agents, i.e., the results of the first model were replicated.

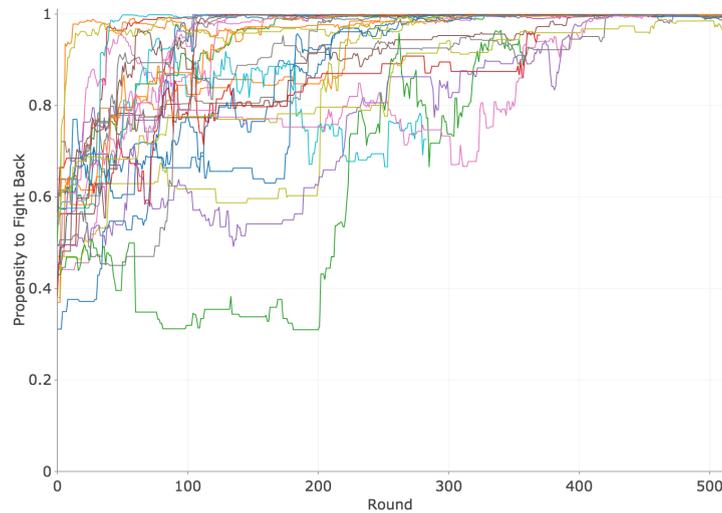


Figure 14: Agents' Propensities to Fight Back

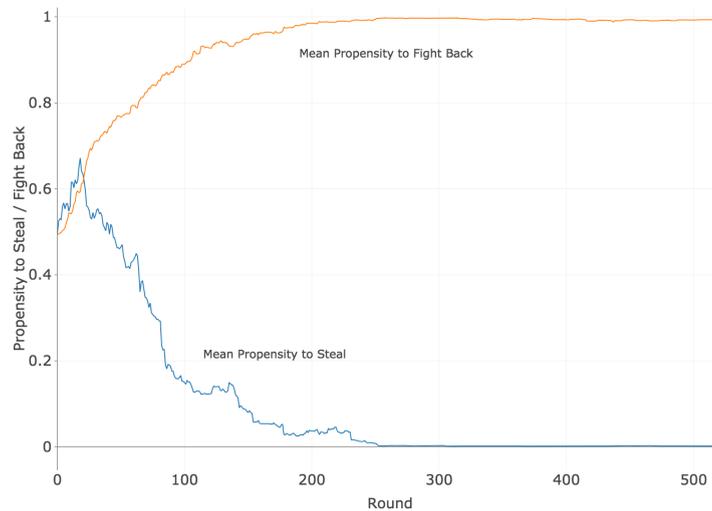


Figure 15: Mean Agents Propensities to Steal and Fight Back

When the default parameters were used, therefore, property rights emerged endogenously. In exploring the parameter space, we will see below the conditions under which this did and did not happen. Before doing that, let us look at why the agents' propensities evolved as they did. An analysis of the data showed there were four important phenomena at work.

First, it became clear to all the agents that it was better to defend their resources than to acquiesce (quadrants 2 and 3 in Fig. 9). If they acquiesced, they lost their whole basket whereas if they fought back, there was a 50% chance they would lose their basket and a 50% chance they would gain the other agent's basket. The default cost of fighting was relatively small (0.2 units of both resources) so, on average, agents learned it was better to defend their resources.

Second, whether it was better to steal or not depended on the likelihood of the other agents fighting back. We found that the *change* in agents' propensities to steal depended on the *level* of the agents' propensities to fight back. In general, we observed that when the agents' mean propensity to fight back was below approx. 0.75, the agents' propensities to steal increased toward 1, and when it was above approx. 0.75, these propensities to steal declined toward zero.

Third, the agents' propensities to steal and fight back tended to get locked in to extreme values. This was due to the use of a logistic equation to adjust these propensities: they changed less the closer they were to 0 and 1. In the early part of the simulation, a general increase in propensities to steal combined with 'noise' meant several of the agents' propensities reach 0.9 - 1.0 where they remained despite the third phenomenon noted above.

Fourth, after approx. 100 round, when the agents' propensities to steal diverged (to 0 or 1), there was a 'survival of the fittest' effect whereby agents with high propensities to steal died and those with low propensities lived. Agents with high propensities to steal incurred higher total fight costs and they benefited less from trade; whereas it was the opposite for agents with lower propensities to steal.

Fig. 16 above extends the agents' propensities to steal to cover 2,000 rounds (the data was taken from the same simulation as in Fig. 13). This figure is shown to highlight the propensities to steal of children, all of whom were born after Round 500.

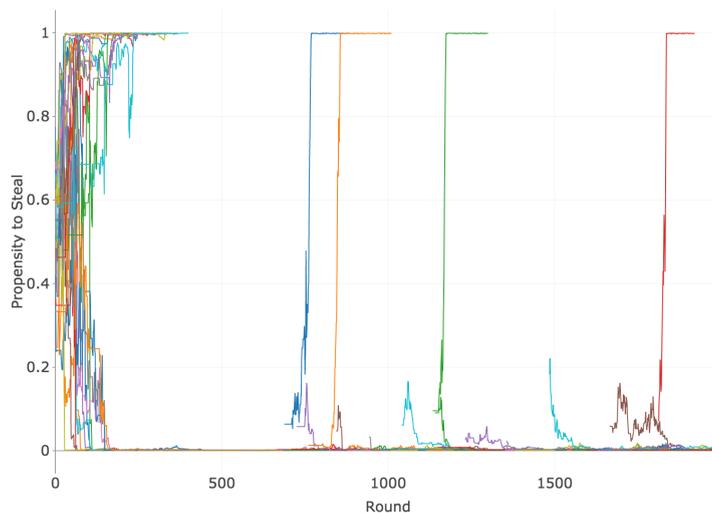


Figure 16: Agents' Propensities to Steal

In the model we assumed children were born with a propensity to steal equal to the mean of their parents' propensities to steal (+/- some variation). The same was true of their propensities to fight back.

The chart shows that agents were born with relatively low propensities to steal and that these generally declined to zero over time. Four of these children saw their propensities to steal rise toward 1 but they did not live for very long.

3.4.1 Game Types

We can differentiate between (i) games considered by agents when they evaluated the expected gains / losses of interacting with other agents; and (ii) games seen when agents actually interacted. In order to map the pay-offs from games with 6 scenarios (Fig. 9) to 4, we calculated the average pay-off of both agents in scenarios 2 and 3, weighted by their propensities to fight back.

There were on average 2.8 million interactions considered by agents in each simulation. The pay-off structure of each potential interaction was analysed and the model kept a tally of how many interactions of which type were considered. In total there 51 game types: 98.4% of these did not

fit the definition of any ‘classic games’ (Prisoners’ Dilemma, Stag Hunt, Strong Deadlock, etc.). Of the 1.6% of games with a pay-off structure identical to any classic game, 1.3% were Strong Deadlock games where both agents had strongly dominant strategies. The Prisoners’ Dilemma and Stag Hunt were both seen in 0.1% of considered interactions. Weak Deadlock (where one or both agents had a weakly dominant strategy) and Matching Pennies were also seen but only very infrequently.

There were 0.42 million interactions on average in each simulation, split among 22 game types. When we analysed these types we found that 99.8% of interactions did not match any classic game type and, again, both the Prisoner’s Dilemma and Stag Hunt was seen in approx. 0.1% of interactions.

3.5 Exploring the Parameter Space

Five parameters were varied and their impact on the results were analysed. We adjusted:

1. The agents’ starting propensities, which led to ‘initial conditions’ tests. Four tests were conducted, which used either low (0.1) or high (0.9) starting mean propensities for the agents. We found that the main results were seen in all four scenarios, that agents came to respect property rights and their propensities to fight back generally increased toward 1.

The emergence of property rights was most challenging when agents’ mean propensities to steal started at 0.9 and their mean propensities to fight back started at 0.1, i.e., approximately the reverse of what emerged in the default scenario. Typically the agent population fell to approx. 3-4 agents as there was more fighting but, eventually, propensities to fight back increased toward 1 and propensities to steal fell toward 0. Agents then traded, specialised, and bore children, leading to a recovery in the population.

2. The cost of fighting. In the default scenario, the cost of fighting was 0.2 units of each resource (deducted from the agents’ personal resource arrays). If the cost of fighting was zero, the agents’ propensities to fight back increased to 1 as in the default scenario but, here, their propensities to steal also increased to 1. In the default scenario, the agents learned that it was more costly to steal *on average*: if this cost was removed then they learned it was preferable to steal.

At higher costs of fighting, we found that agents were increasingly reluctant to interact with each other. At costs exceeding approx. 1.25 units for each resource, agents hardly ever interacted with other agents: there was a risk they would incur the cost of fighting². This lack of interaction meant agents did not learn to defend their resources nor to respect property rights.

When the cost of fighting was between approx. 0.75 and 1.25 resource units, agents interacted more but their propensities to fight back typically plateaued at around 0.75 and their propensities to steal drifted around 0.3. In this range, agents failed to specialise and the agent population stabilised at around 8-10 agents.

There was a ‘sweet range’ of approx. 0.1 to 0.75 resource units within which property rights emerged in the way seen in the default simulations.

3. The speed of adjustment of the propensities to steal and fight back (r). As mentioned above, a logistic equation was used to adjust agents’ propensities in light of experience. In extremis, if this speed of adjustment was zero, the agents’ propensities would not change so they would not reach a point of respecting property rights.

For very small values of r (e.g., 0.01), the agents’ propensities changed very slowly and the agent population typical fell to approx. 3-4 agents. However, the surviving agents went through the

²With a propensity to steal of, say, 0.5, agents would not trust themselves not to steal and incur a high fight cost.

same learning process as in the default scenario, albeit more slowly. Eventually, propensities to fight back increased to 1 and propensities to steal fell toward 0, enabling trading, specialisation, and a population boom.

For higher values of r (say, 1.0), agents' propensities to steal and fight back moved rapidly between extreme values of 0 and 1, which meant respect for property rights was never stable for individual agents. The population typically declined to approx. 3-4 agents, none of whom specialised nor bore children.

We found that property rights successfully emerged if, approximately, $0.01 \leq r < 0.45$.

4. The influence of the counterparts' experiences on the adjustments of the propensities. In the default simulations we assumed that agents learned from their own experiences in an interaction and they learned from their counterpart (with a weight (β) of 0.3 attributed to the counterpart's experience). If we reduced β to zero we found that agents' propensities changed more slowly but the same results were seen as in the default simulations: propensities to fight back increased toward 1 and propensities to steal fell to 0. Agents then specialised and a population boom followed. Higher values of β meant agents learned more rapidly but the same results were seen.
5. The standard deviation applied to children's' propensities to steal and fight back. In the default simulation, children were born with propensities to fight back and steal equal to the mean of their parents' propensities +/- a value drawn from a normal distribution with mean zero and a standard deviation of 0.1 (these propensities had a floor of zero and a ceiling of 1). If we increased this standard deviation we found that the children's propensities to steal remained at or declined to zero (eventually) or they died. The population of agents increased to approx. 43 agents regardless of this standard deviation.

3.6 Experimentation

Five experiments were conducted. Four of these are summarised below and the fifth, concerning formal institutions, is presented in the next sub-section.

1. 'Social Construction' Tests. Here we explored whether agents' propensity changes were independent of other agents' strategies. To do this we allowed only one agent to change its propensities to steal and fight back (denoted the 'change agent') and fixed all other agents' propensities at certain levels: (i) $P^S = P^{FB} = 0$; (ii) $P^S = 0, P^{FB} = 1$; and (iii) $P^S = P^{FB} = 1$.

The results were illuminating: in (i), the 'change agent' became a hawk ($P_i^S = 1$) among the doves who always acquiesced; in (ii), the agent became a dove ($P_i^S = 0$) because the other agents always fought back; and in (iii) the change agent became a dove ($P_i^S = 0$) in a world of hawks (and it typically died). These results emphasised a type of 'social construction' phenomenon: the change agent's learning process was contingent on the other agents' propensities.

2. 'Yellow Agents'. What happened if we allowed P^S to vary but we fixed $P^{FB} = 0$ for all agents? This was equivalent to eliminating scenarios 2F and 3F in Fig. 9: agents only acquiesced. We found that the agents' propensities to steal rose toward 1 and the population declined to approx. 2-4 agents.
3. 'Black Sheep'. What happened if the first child born only ever stole, i.e., we fixed its propensities to steal at 1? (we called this child the 'black sheep'). This is equivalent to an invading defector in iterated prisoners' dilemma games. We found that these children lived for about 200 rounds and then died. Crucially, the adults' propensities to fight back remained close to 1, which meant

the black sheep persistently incurred a fight cost. They were not bullies who exploited the adults. Furthermore, the adult population continued to respect property rights. Overall, the black sheep had little impact on the agent population.

4. ‘Power in Fighting’. What happened if the winner of a fight was determined by a measure of ‘power’ such as skill in fighting? In the default simulations, fights were decided by the flip of a coin. Here, we introduced the idea of power in a fight: a ‘fight skill’ variable was created for each agent, which was the deflated sum of past fights, i.e., the more an agent had fought in the past, the higher this skill was³. Here, the probability an agent won a fight was its fight skill divided by the total fight skill of the two fighting agents.

We found that the main change in the simulations was that agents with higher propensities to steal were in general more able to survive for longer: their fighting skills were higher on average than agents with lower propensities to steal, which allowed them to ‘bully’ these agents. This meant the ‘survival of the fittest’ phenomenon mentioned above was more beneficial to agents with higher propensities to steal / fighting skills.

In general, agents failed to specialise and the agent population declined to approx. 2-4 agents. In a few simulations, a number of agents with low propensities to steal survived long enough to have children. These children inherited low propensities to steal from their parents and a population of agents that respected property rights increased in size. However, such a population was never sustainable: crucially, these environments were ‘fertile ground’ for bullies to emerge.

We can see this pattern in Fig. 17 below, which shows the fight skills of living agents at the end of each round (each line represents one agent) in a typical simulation. In the first 1,000 rounds, agents were split in to two groups: bullies with fight skills of approx. 250-500 (and with high propensities to steal) and non-bullies with low propensities to steal and fight skills of approx. 100. By Round 1,500, only a few agents remained but they generally respected property rights and they had low fighting skills. Children were born but, eventually, a bully emerged and the agent population declined. Fig. 18 shows the mean propensities to steal and fight back of all living agents in the same simulation.

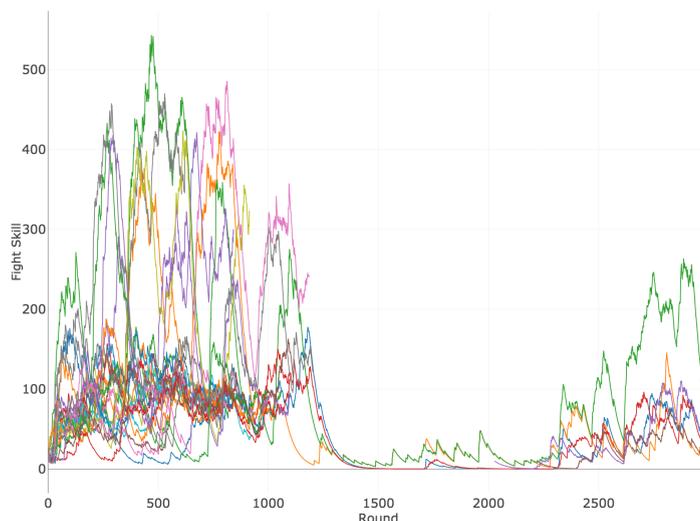


Figure 17: Agents’ Fight Skills over 3,000 Rounds

³At the end of each round we deflated the existing score by 2% and added the number of fights the agent had taken part in during that round.

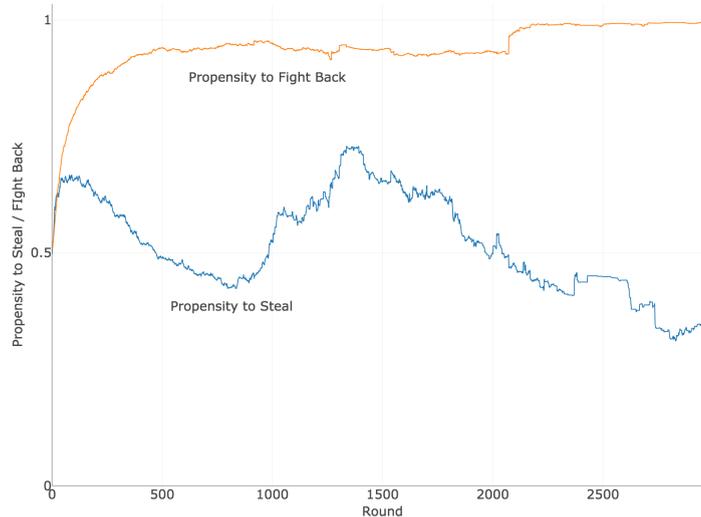


Figure 18: Agents' Propensities to Steal and Fight Back over 3,000 Rounds

3.7 Formal Institutions

In scenarios where property rights did not emerge endogenously, could we use formal institutions (here, laws) to encourage agents to observe property rights?

When we explored the parameter space and conducted various experiments, we observed 7 scenarios in which property rights did not emerge endogenously across the whole population. When we applied formal institutions to these scenarios, assuming no corruption, we found there were types of formal institutions which encouraged the population to respect property rights in all 7 scenarios. Below we look at two of these.

In the first sub-section below we discuss the results from applying formal institutions to the 'Yellow Agents' and 'Power in Fighting' experiments in the absence of corruption, i.e., when laws were applied perfectly by some policing authority. In the following sub-section we run the same experiments but in the presence of corruption.

3.7.1 Without Corruption

Two types of formal institution were applied: in the first, a fine (only) was applied to agents that attempted to steal from other agents; and, in the second, a fine was levied on an agent attempting theft but this fine was then given as compensation to the counterpart agent if this agent had attempted to trade. A fine represented a deduction from both resources in an agent's personal resource array and compensation was an addition to these arrays.

When estimating the gains / losses from interacting on the grid, agents took in to account any fines and compensation. In addition, when two agents interacted on the grid, the policing authorities applied fines and compensation to all interactions when theft was attempted.

Yellow Agents

Recall that in this experiment, agents never fought back when they wanted to trade and their counterpart wanted to steal: agents always acquiesced. Without any formal institution, the agents' propensities to steal increased to 1.

When we applied the two formal institutions to this scenario we found that fines alone did not discourage agents' propensities to steal from increasing toward 1. This was true for any value of fine.

However, when compensation was paid in addition to fines, agents' propensities to steal declined to zero if the fine / compensation was approx. above 0.35 resource units. Fig. 19 below shows the mean propensities to steal and fight back of all living agents over 20 simulations when the fine / compensation was 0.5 units.

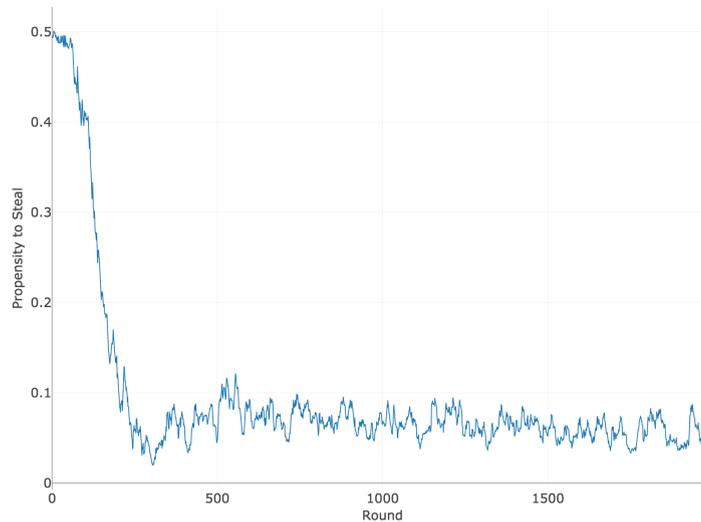


Figure 19: Agents' Propensity to Steal over 2,000 Rounds (with Fine & Compensation)

Power in Fighting

As mentioned above, when agents were given a 'fight skill', which gave them power in fights, the agent population as a whole failed to respect property rights in the long term. It was notable that if the propensities to steal of an agent population all declined toward zero, this proved to be fertile terrain for bullies to emerge.

When we applied both types of formal institution to this scenario we found that fines on their own were sufficient to encourage propensities to steal to decline to zero, provided the fine exceeded approx. 0.07 resource units and was less than approx. 1.6 units. Fines higher than 1.6 units had the effect of discouraging interaction between agents, which meant they failed to learn to defend their resources and respect property rights.

Fig. 20 below shows the fight skills of agents over 2,000 rounds in a typical simulation when agents were fined 0.1 resource units for theft (it should be compared with Fig. 17 above). Fig. 21 below shows the mean propensities to steal and fight back of all living agents over 20 simulations with the same fine (it should be compared with Fig. 18).

The figures show there was a decline in propensities to steal over approximately the first 1,000 rounds, and fight skills remained relatively low. However, we can see once again that a society of agents with very low fighting skills was fertile terrain for a 'bully' to emerge. The fine of 0.1 resource units was not sufficient to prevent this; however, the fine ensured bullies paid a significant cost in stealing, which meant they died or became non-bullies. In Fig. 20 below we can see that of the 3 bullies which emerged, 2 of these died before the end of the simulation.

Ultimately, despite a few bully agents emerging, most agents came to respect property rights, markets formed, the agents specialised and the agent population increased to approx. 40.

For higher fines (but less than 1.6 resource units), the discipline on bullies was greater, which meant fewer emerged and those that did died more quickly.

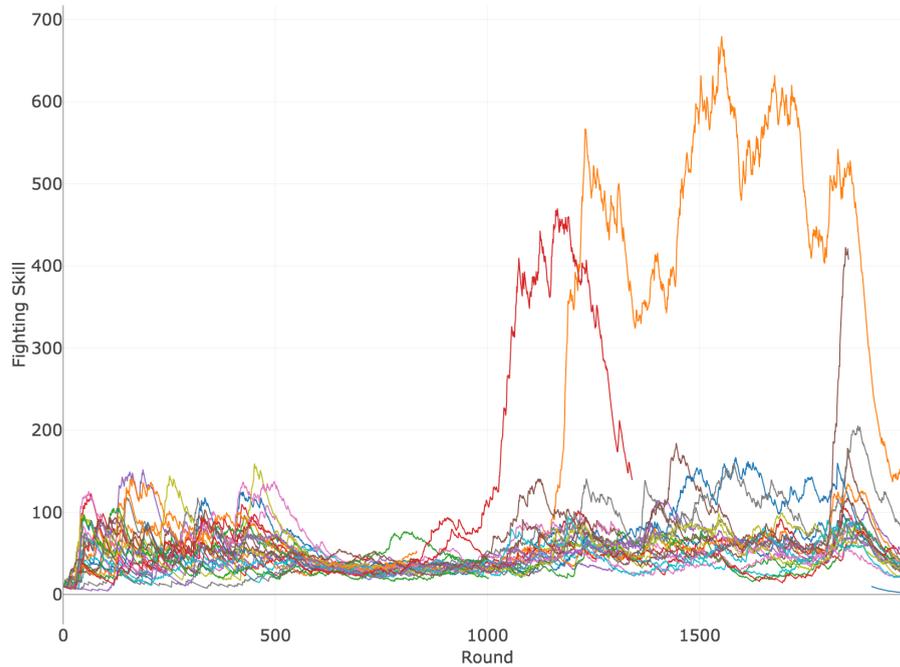


Figure 20: Agents' Fight Skills over 2,000 Rounds (with Fine)

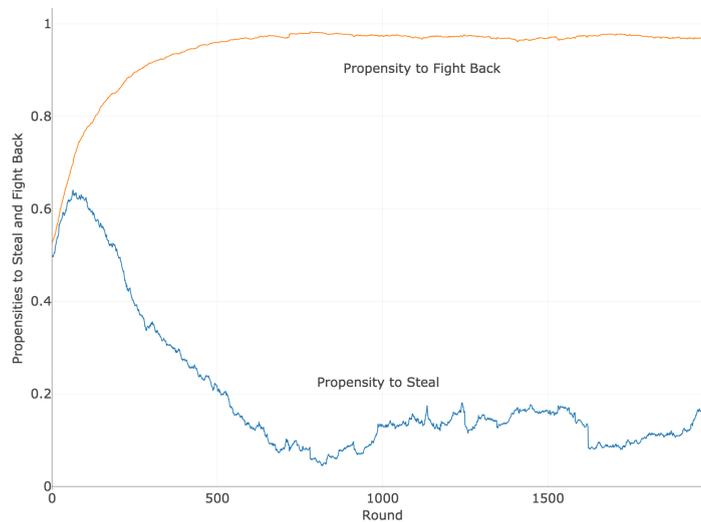


Figure 21: Agents' Mean Propensities over 2,000 Rounds (with Fine)

If compensation was paid to 'victims' then we found the minimum fine / compensation had to exceed 0.02 resource units for property rights to be respected, and there was no maximum.

3.7.2 Corruption

In reality, formal institutions are applied by people who live in societies and not by some selfless, abstract entity. Corruption is possible when the policing authorities might use the power they have owing to authority for personal gain.

In these experiments, when fines were applied without compensation, we assumed that stealing agents would pay the full fine if the policing agent was not corrupt but it would pay a bribe (less than or equal to the fine) if the policing agent was corrupt. Bribes were assumed to be a fixed proportion (denoted $0 \leq \rho \leq 1$) of fines, where ρ was a parameter in the simulations.

When the formal institution employed both fines and compensation, the fine and compensation were paid as above if there were no corruption but if there was corruption, the policing agent was paid a bribe and no compensation was paid to the ‘victim’.

A probability was used to determine whether a policing agent would be corrupt in each interaction. We decided to link this probability to the prevailing respect for property rights (propensities to steal) in the population. There is a good argument that corruption in these simulations should be associated with respect for property rights because corrupt policing agents would be violating any rights over resources dictated by a formal institution. For example, if a fine only were applied and a policing official took a bribe equal to 10% of the required fine then the official would take some of the resources for their own personal gain and the bribing agent would keep 90%.

Similarly, if corruption meant a victim did not receive compensation then they lost the right to resources as defined by the formal institution.

In these experiments, the probability that a policing agent was corrupt was equal to the mean propensity to steal of the living agents. This added complexity to the simulation because, now, formal institutions influenced the evolution of the agents’ propensities to steal *and* these propensities influenced the efficacy of the formal institutions, both at the same time.

Yellow Agents

There were two parameters in this experiment: the fine / compensation and ρ . For the sake of brevity, here we fixed the fine / compensation at 0.5 resource units and we looked at a range of values for ρ .

Fig. 22 below shows the mean propensities to steal of living agents over 20 simulations when $\rho = 0.5$. It should be compared with Fig. 21 above. We can see that agents’ propensities to steal increased to approx. 0.8 over 2,000 rounds: corruption meant the formal institution failed to encourage the agents to respect property rights.

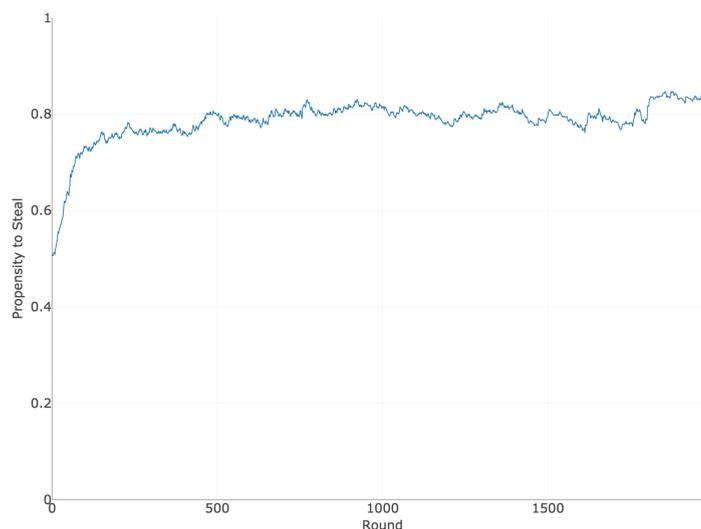


Figure 22: Agents’ Mean Propensity to Steal over 2,000 Rounds (with Fine and Corruption)

At the beginning of the simulation in Fig. 22, the mean propensity to steal of agents was approx. 0.5, which was also the probability that a policing agent was corrupt. Therefore, in half of interactions, stealing agents paid the full fine of 0.5 resource units and in the other half they paid a bribe of 0.25 resource units. Importantly, in half of interactions (when stealing agents bribed the policing agent) the ‘victim’ agents received no compensation.

The reduced efficacy of the formal institution meant agents’ propensities to steal increased (recall that without an effective formal institution, agents’ propensities to steal increased to 1). This created a positive feedback loop: higher propensities to steal led to more corruption and, eventually, the efficacy of the formal institution collapsed.

When we ran this experiment for values of ρ between 0 and 1 we found the same results for all values of ρ .

A question this feedback loop raised was about initial conditions. More specifically, what would happen if agents’ propensities to steal started lower, e.g., at 0.1? We ran these simulations and found that starting conditions did indeed matter. Fig. 23 below shows the average propensity to steal of all living agents over 20 simulations: notice how the mean propensity to steal increased from approx. 0.1 to approx. 0.25 and then declined back to approx. 0.1. In general, in these simulations, the agents came to respect property rights and the population increased.

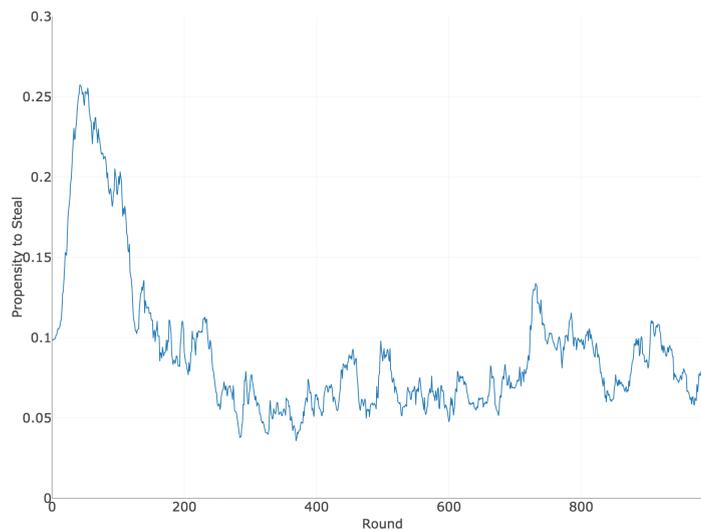


Figure 23: Agents’ Mean Propensity to Steal over the First 1,000 Rounds (with Fine and Corruption)

The initial rise of the mean propensity to steal in Fig. 23 was because approx. 1/3 of agents saw their propensities to steal rise to nearly 1 and about 2/3 of agents saw their propensities decline to approx. 0. This divergence was due to the imperfect application of the formal institution: some agents received windfalls when they stole and successfully bribed. Once again a positive feedback loop meant the effectiveness of the formal institution worsened during the first 50 rounds.

Despite the increase in mean propensities to steal and the counterpart increase in corruption, the formal institution was still effective enough to ensure that agents with high propensities to steal incurred both fight costs and fines / bribes. In addition, agents with low propensities to steal were net recipients of resources via compensation. This meant the former agents typically died after Round 100 (approximately) while the latter agents survived. This resulted in a decline in propensities to steal and a more effective formal institution.

Overall, despite the spike in propensities to steal in the first 100 rounds, when agents started with lower propensities to steal the whole system became locked in to lower corruption and a more effective formal institution, which maintained low propensities to steal (and low corruption). When we ran these experiments with higher mean starting propensities, we found a threshold of approx. 0.2 for agents' starting propensities to steal, above which these propensities tended to approx. 0.8 (the spike in mean propensities was higher with higher starting propensities) and below which they tended to 0.1.

Power in Fighting

Understanding the impact of corruption on the agents was simpler when fines were levied without compensation: for values of $\rho = 1$, bribes were identical to fines so the results were the same as those seen without corruption. For values of $\rho < 1$, the amount paid by stealing agents (in fines and bribes) declined, depending on the probability of corruption.

The experiments here were implemented using a fine of 0.1 resource units.

In general we found that for higher values of ρ (say, $\rho > 0.7$), the results seen without corruption were approximately replicated: the fines / bribes paid by stealing agents were relatively high, which meant most agents came to respect property rights - there were fewer bullies and they lived relatively shorter lives. All of this resulted in a probability of corruption of approx. 0.15, which meant the formal institution remained effective and the cost of stealing remained high.

For lower values of ρ (say, $\rho < 0.3$), the results were closer to scenarios when there was no formal institution (Section 3.6). The cost of stealing was lower, more bullies emerged on average and they lived longer, enabling a higher probability of corruption. This in turn reduced the efficacy of the formal institution, which limited the cost of stealing.

Note, however, that fighting agents still incurred a fight cost in addition to any fine or bribe: this limited the number of living bullies to a maximum of approx. 5 among approx. 15 non-bully agents. As a result, the probability of corruption was typically capped at about 0.25 in these simulations.

Finally, the simulations in these experiments showed considerable variation because the emergence of bullies was unpredictable. In some simulations a population with agents who all respected property rights lasted several hundred rounds before a bully emerged and in others, it lasted only a few rounds.

3.7.3 Formal Institutions: Conclusion

We found it was possible to use formal institutions to encourage agents to respect property rights in the two scenarios we looked at, when there was no corruption.

In the 'Yellow Agents' scenario, a fine on its own was not sufficient to bring about respect for property rights: compensation was also required. In the 'Power in Fighting' scenario, a fine on its own was sufficient.

It is worth noting also that the fine / compensation in 'Yellow Agents' had to be much larger (above 0.35) than the fine in 'Power in Fighting' (above 0.07). This was also true of the fine / compensation in the latter (above 0.02).

The impact of corruption was also significantly different in the two scenarios: in 'Yellow Agents', the whole system could get locked in to different states. In one such state, there was little corruption and formal institutions remained effective, which helped to maintain a respect for property rights

and low levels of corruption. In another state, corruption was rife, which reduced the effectiveness of the institution and agents failed to respect property rights. Which state transpired depended on the initial conditions.

In the ‘Power in Fighting’ scenario, the impact of corruption was less clear: the results depended on ρ . In general, and paradoxically, high values of ρ typically led to lower levels of corruption (and fewer bullies) and vice versa.

3.8 Neural Networks

In the simple ‘probabilistic heuristics’ approach described above, game pay-offs were ignored by the agents: they exhibited tendencies toward stealing and defending their resources irrespective of potential gains or losses. As was also mentioned above, we can view this approach as the opposite of the classic rational choice approach, which is prevalent in conventional game theory, whereby choices are determined only by the agents’ perceptions of pay-offs

Using neural networks to mimic internal models appeals for two reasons. First, neural networks are inspired by human cognition, so there might be some merit in using them to model agent decision making and learning⁴. Second, neural networks allow for inputs to be used, which means agents can account for game pay-offs, unlike the simpler method above.

The question which guided our work here was whether agents’ internal models came to replicate standard rational choice theory or the simple heuristics approach used above.

For readers unfamiliar with neural networks, Fig. 24 below illustrates a typical network.

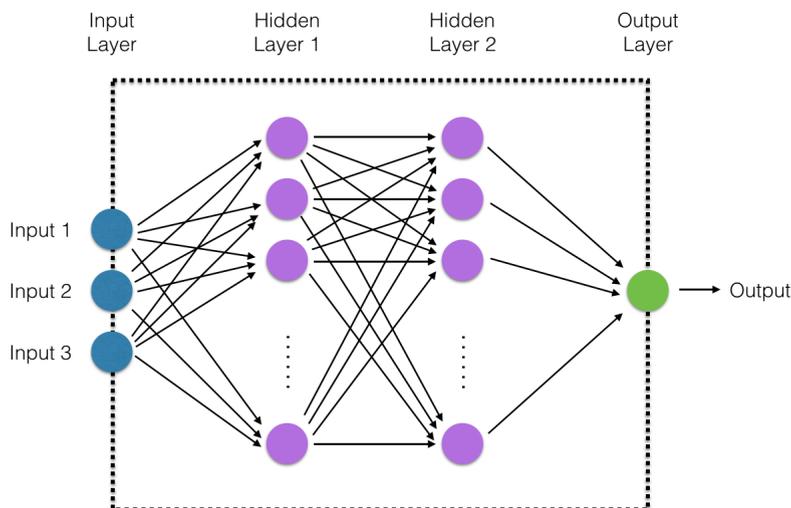


Figure 24: A Neural Network

These networks typically have an input layer consisting of multiple inputs; hidden layers (the network’s inputs are ‘received’ by the first hidden layer and the outputs of each hidden layer are inputs in to the next layer, until the last); and outputs. The mathematics of neural networks will not be discussed here.

⁴A comprehensive discussion of this is beyond the scope of this paper.

In these experiments, each agent was given two neural networks, one corresponding to each of its propensities to steal and fight back. There were 14 inputs in each of these networks: the 12 pay-offs in each interaction (one for each scenario in Fig. 9, for each agent) plus the agent's prediction of its counterpart's propensities to steal and fight back (making use of the reputations architecture mentioned above).

The agents' neural networks were given 3 hidden layers of 14 nodes each, and each network generated an output which was a probability: the agent's propensity to steal or fight back.

In each interaction the inputs were processed by the agents' neural networks, resulting in outputs. Typically, in the initial interactions of a simulation, the neural networks generated propensities to steal and fight back of approx. 0.5.

Neural networks learn through a process known as backward propagation, which is when the results of outputs are fed back to the coefficients in the network. In these experiments, suppose, for example, an agent's propensity to steal was high and this resulted in a decision to steal, and suppose also that its counterpart acquiesced, resulting in a gain to the agent. This gain was fed back to the neural network, resulting in an adjustment of the coefficients in the hidden and output layers. Simply put, if the same situation (inputs) were observed again, this adjustment would make it slightly more likely the agent would steal.

Over time, if an agent were exposed to different pay-offs (game types), this learning process might mean its outputs will vary according to the nature of the game, e.g., it might generate a high propensity to steal for Prisoners' Dilemma games and a low propensity for Stag-Hunt games.

It is perhaps worth noting that in each of the default simulations described above, agents born at instantiation interacted approx. 15,000 times on average, and agents learned after every interaction.

3.8.1 Preliminary Results

This is work in progress so only preliminary results are presented here.

Under certain parameter settings (for the neural networks), the results of the default simulations above were replicated: agents came to respect property rights, markets emerged, the agents specialised and the population increased to its carrying capacity of approx. 43 agents.

An analysis of the data showed that despite the inputs providing information to the neural networks, the outputs of the networks were almost completely insensitive to the inputs. This meant the neural network approach closely approximated the simpler heuristics approach on which most of the results above were based.

3.9 Discussion of the Property Rights Model

In considering the results of the default simulations in the second model, it was not immediately clear to us whether or not each agent's propensities to steal and fight back evolved independently of other agents' propensities. An analysis of the data and the results from the 'Social Construction' experiments indicated that these propensities did not evolve in isolation: once again we saw that agents' internal models *co-evolved* and, from this co-evolution, property rights emerged.

As mentioned in Section 2.7 above, the first model led to a reduction in uncertainty and there were efficiency gains which arose from coordination.

In the second model, there was a reduction in uncertainty regarding property rights in the default simulations. We would argue that this reduction in uncertainty had two aspects to it. First, and most importantly, the expected propensity to steal of an agent had to be either 0 or 1 for another agent to be certain of its behaviour: if this propensity were anything else, the agent could attempt to steal or trade.

Second, for certainty to be maintained over time, an agent's actual propensity to steal had to be consistent with its perceived propensity. If there was any difference then its reputation would change.

We saw that such a state of reduced uncertainty was achieved in the default simulations of the second model.

From an efficiency perspective, we can once again identify direct and indirect gains from an institution (here, respect for property rights). Direct gains included a decline in fight costs owing to agents attempting to steal, and the benefits of agents trading.

The indirect benefits included the emergence of markets and the resulting productivity boom owing to specialisation, which were seen in the first model.

The second model also helps us appreciate that it was possible for an *inefficient* institution to exist: if the agents' propensities to steal increased to 1 (as they did in the 'Yellow Agents' experiment) then this also reduced uncertainty (since all agents expected theft) but the result was highly inefficient.

It is also worth emphasising that the right to defend one's resources (a high P^{FB}) was also an institution and there was inter-dependence between this and respect for property rights. On the one hand, as we have seen, the evolution of P^S depended on the prevailing propensities to fight back. On the other hand, agents' propensities to steal determined the number of interactions in quadrants 2 & 3 in Fig. 9 (page 18), which is when the agents learned to defend their resources (or not).

The second modal also demonstrated that it was possible for formal institutions to encourage agents to respect property rights but different institutions were required in different scenarios. In addition, success depended on the existence of corruption.

4 Conclusion

The first aim of the work presented in this paper was to explore the emergence of informal institutions in economic systems. One of the motivations mentioned in the introduction was a quote from Field who emphasised the need to ‘explain systematically how institutions are created’ [8]. Another motivation was Epstein’s argument, consistent with Field, that social scientists should seek to ‘grow’ macro phenomena from the ground up ([6]).

We pursued this aim by developing computational models and by employing various concepts which are emphasised in the complexity sciences, including emergence, co-evolution, mental models, and inductive reasoning.

Looking at the work as a whole, we can see that the default simulations of the second model showed the emergence of property rights enabled the emergence of efficient markets, and this in turn enabled the division of labour and a productivity boom.

However, while informal institutions emerged in the default simulations of the first and second models, we showed that in various parts of the parameter spaces, these institutions did not emerge. This was also true in certain experiments.

In the second model we explored these scenarios - when property rights did not emerge - to consider whether formal institutions could encourage or enable such rights. This was a second aim of the work presented here. As mentioned in Section 3.9 above, it was possible to develop formal institutions which led to agents respecting property rights but this result was contingent on there being no corruption.

When we introduced corruption, which we linked to the agents’ propensities to steal, we found that this phenomena could prevent the emergence of property rights and, by extension, efficient markets. In one experiment (‘Yellow Agents’) and under certain initial conditions, the possibility of corruption led to system-wide outcomes where agents failed to respect property rights and corruption was rife.

We believe it likely that if formal institutions were used in an attempt to prevent corruption then corruption itself would also mitigate the efficacy of such institutions. This helps us emphasise the leading role of informal institutions in economic development: formal institutions are not necessarily the answer if corruption undermines their efficacy. This work also support North’s emphasis in [14] on the need for formal institutions, e.g., constitutions, to be supported by social norms and values.

Uncertainty and efficiency were discussed in the context of both models. The different emergent informal institutions helped to mitigate uncertainty, consistent with North ([13]): in the first model this was about geographic coordination and in the second model it was about agents respecting property rights.

Efficiency in these models was measured as the gains from coordination rather than the minimization of transaction costs. The direct gains from the emergent institutions could be interpreted as a reduction in transaction costs if we took a broad enough definition of this term but in these ground-up model we felt it more appropriate to articulate efficiency as the gains from coordination.

Furthermore, the coordination of behaviour which the informal institutions achieved also gave rise to indirect gains. In the second model, this was the enabling of markets; and in the first model it was the division of labour. It is noteworthy that the efficiency gains arising from the division of labour had nothing to do with transaction costs.

Finally, it is interesting to note that the some of the results from the default simulations of the larger, second model replicated the general equilibrium model ([2]): after the emergence of markets, the mean

price of transactions was very close to the market clearing price in each round, and the total volume traded was approximately the same as the market clearing quantity.

These results were, however, contingent on the emergence of property rights and markets.

References

- [1] Archer, M.: *Structure, Agency and the Internal Conversation*. Cambridge University Press (2003)
- [2] Arrow, K., Debreu, G.: Existence of an equilibrium for a competitive economy. *Econometrica* 22 (3), 265–290 (1954)
- [3] Arthur, W. B.: Inductive Reasoning and Bounded Rationality (the El Farol Problem). *American Economic Review (Papers and Proceedings)*, 84, pp 406-411 (1994)
- [4] Bronk, R.: Reflexivity unpacked: performativity, uncertainty and analytical monocultures. *Journal of Economic Methodology*, Vol. 20, Issue 4 (2013)
- [5] Denzau, A., North, D.: Shared Mental Models: Ideologies and Institutions. *Kyklos*, Vol. 47 (1994)
- [6] Epstein, J.: *Generative Social Science: Studies in Agent-Based Computational Modeling (Princeton Studies in Complexity)*. Princeton University Press (2007)
- [7] Epstein, J., Axtell, R.: *Growing Artificial Societies : Social Science from the Bottom Up*. MIT Press (1996)
- [8] Field, A.: Beyond Foraging. *Journal of Institutional Economics* Vol. 3 No. 3 (2007)
Alexander Field in a paper entitle *Beyond Foraging* (2007, *Journal of Institutional Economics* Vol. 3 No. 3)
- [9] Holland, J.: *Signals and Boundaries: Building Blocks for Complex Adaptive Systems*. MIT Press (2014)
- [10] Holland, J., Holyoak, K., Nisbett, R., Thagard, P.: *Induction (Computational Models of Cognition and Perception): Processes of Inference, Learning, and Discovery*. A Bradford Book (1989)
- [11] Knight, J.: *Institutions and Social Conflict (Political Economy of Institutions and Decisions)*. Cambridge University Press (2008)
- [12] North, D. C.: *Structure and Change in Economic History*. New York: Norton (1981)
- [13] North, D. C.: *Institutions, Institutional Change and Economic Performance*. Cambridge University Press (1991)
- [14] North, D. C.: *Understanding the Process of Economic Change*. Princeton University Press (2005)
- [15] Schofield, N.: Anarchy, Altruism and Cooperation: A Review. *Social Choice and Welfare*, Vol. 2, pp 207-219