



ISSN 1178-2293 (Online)

**University of Otago**  
**Economics Discussion Papers**  
**No. 1015**

**School of Business**  
Te Kura Pakihi

October 2010

---

## **Virtually Science: An Agent-based Model of the Rise and Fall of Scientific Research Programs**

**Daniel Farhat <sup>1</sup>**

*Address for correspondence:*

Daniel Farhat  
Department of Economics  
University of Otago  
PO Box 56  
Dunedin  
NEW ZEALAND  
Email: [dan.farhat@otago.ac.nz](mailto:dan.farhat@otago.ac.nz)  
Telephone: 64 3 479 8645  
Fax: 64 3 479 8174

## Abstract

Is there more to ‘good science’ than explaining novel facts? Social interaction within the scientific community plays a pivotal role in defining acceptable research. This paper develops an agent-based computational model of scientific communities to explore the connection between research outcomes and the socio-cultural environment they are constructed in. Agent-to-agent interaction is added to the notion of scientific research programs (SRPs) developed by Lakatos (1969, 1970) as an important factor guiding the practice of researchers. Early in scientific inquiry, when there are many novel facts to explain, the research community is fragmented and researchers prefer to use their own talents (or *innovation*) when conducting research. Over time, consensus emerges in the form of a dominate SRP which utilizes intra-disciplinary approaches (or *adaptation*). This result illustrates the rise and fall of scientific paradigms as described by Kuhn (1970). The history of Keynesian macroeconomic theory provides colloquial support for the simulation model.

**Keywords:** Methodology, Agent-based models, Sociology of Social Science, Scientific Research Programs.

# 1 Introduction

What is ‘good’ science? Many active researchers in both the physical and the social sciences would agree that some semblance of the hypothetico-deductive<sup>2</sup> model sets the standard of practice when constructing knowledge about the world around us. For the social sciences, however, following the steps laid out by hypothetico-deductivism is not always practical. In many cases, social scientists derive understanding about human behaviour from a single, non-repeatable observation (*the way things are*). Because testing hypotheses by repeated experimentation is not possible in these cases, validation of one’s research results becomes difficult. In its place, notions of peer assessment (in particular, peer *acceptance*) have been interwoven into the definition of ‘good research’ in the ‘soft-sciences’ as a means to reduce the need for corroboration via experimentation. Peer-to-peer interaction in these fields results in a ‘culture of science’ which guides the work performed by social scientists and validates the knowledge they produce.

Cultures of science may make it difficult for methodologists to evaluate research practices. Within the scientific community, rules and norms are established that determine what sorts of questions a researcher might ask and what techniques she may use. The socio-cultural environment determines how results are disseminated and how they are appraised by others, also affecting the process of research. Further, cultures of science, like other cultures, change over time and adapt to the needs and wants of those within them. The complexity and dynamic nature of peer interaction within scientific communities make it difficult to clearly pin down the structure of scientific inquiry and the evolution of research methods. Without accounting for the influence of cultures of science on the definition of acceptable practice, many methodological debates can (and do) remain unresolved.

To understand the role of social interaction plays in the creation of scientific knowledge, I build a scientific community within a virtual environment. Using a computerized framework, I create individual scientists, each with a set of knowledge, a set of available actions and a set of rules for guiding their behaviour. I impose a group structure on the researchers using a popular idea from methodological theory which captures a notion of ‘community’. Agent-to-agent interaction influences how individual scientists make decisions about practice and how the scientific community develops over time. I evaluate the performance of the virtual community to extract the relationship between social behaviour and scientific practice. I focus specifically on the interactions between scientists instead of the actual creation of science to pin down this relationship.

Results from the simulated scientific communities suggest that the value society places on successfully replicating observed facts influences and is influenced by the practices performed by researchers. Even though it is possible for the virtual scientific community to prefer inferior approaches, the society assigns large amounts of prestige to the most successful types of research on average. In addition, when science is young and little is known about the perceived world, scientists rely on their own abilities to perform research and the scientific community is fragmented. The tendency to ‘go your own way’ when conducting research distorts the probability a research project is successful. As more research is undertaken and disseminated between scientists, a unique approach to inquiry often emerges (*consensus*). This method involves mixing the successful approaches used by others in the past (an *intra-disciplinary* approach), which is shown to enhance the probability research is successful. The scientific community allocates the most prestige to this method, which reinforces the state of consensus. We can infer from this pattern that the scientific community becomes fragmented when a large number of new stylized facts appear, as they often do in the social sciences, but then coalesces over time.

The simulation results link together sociological approaches to understanding knowledge construction (Hands (2004)), the notion of scientific research programs (Lakatos (1969, 1970)), and the rise and fall of paradigms of science (Kuhn (1970)) within a computerized framework. To substantiate the model’s intuition, I refer to recent developments in the field of macroeconomics: the rise and fall of Keynesian macroeconomic theory. In the sections that follow, I describe the motivation behind the project and the structure of the virtual model. Next, I perform and discuss a series of simulations. After connecting the simulation results to methodological theory, I end with a short conclusion.

## 2 Scientific Communities

Recently, the sociology of scientific knowledge has received a lot of attention as a step towards dissecting cultures of science in the hope of resolving several methodological debates (Hands (1994, 2004)). From the sociological perspective, we can enrich our understanding of proper research practices by studying the way scientists interact (through seminars, journal articles, pedagogy, etc.) and the institutions they create (such as professional associations, journals, etc.). If the definition of ‘good science’ is interrelated with successful interactions within these institutions<sup>3</sup> *in addition* to explaining observed facts, then the social exchanges associated with active research is an essential part of defining methods of science and the meaning of acceptable research. We see these ideas starting to emerge in the 1970s with the early work by Kuhn

(1970) on paradigms of science and Lakatos (1969, 1970) on scientific research programs. Although explaining novel facts is paramount in their models of scientific practice, it is done so within some notion of the ‘scientific community’. Further, concepts such as *rhetoric* (McCloskey (1983)) and *storytelling* (Morgan (2009)) have highlighted the importance of interactive behaviour between peers (through persuasion or argument construction) in the analysis of scientific methodology.

Conveniently, sociological and institutional approaches to scientific inquiry accommodate a specific computational technique: agent-based modelling (ABM). Agent-based modelling has become a novel way to add realism and complexity to models which require interaction between individuals. In ABM, artificially intelligent agents are provided with private, local information and a set of objectives to achieve. They are then allowed to communicate with each other within the virtual environment. As agents react to a mixture of private information and interpersonal contact, they affect aggregate outcomes and build group behaviour from the “bottom up”.<sup>4</sup> Because both the agent’s behaviour and the way they interact is set by the programmer using a series of rules and meta-rules<sup>5</sup>, ‘experiments’ can be performed by augmenting initial conditions or by turning on and off features of the computerized model. With this approach, methodological studies can become more proactive. The methodologist can create scientific communities according to their theories of scientific practice, design the behaviours of the researchers, generate alternate realities and then compare the results to what she observes in her field.

There are relatively few such simulation models focusing on the ‘sociology of social science’. Gilbert (1997) models the production of *academic* science as a structural system in which researchers write papers based on past work they have completed and the work done by randomly selected peers. The simulated model is able to replicate citation behaviour seen in academic journals and exhibits the endogenous formation of specialities amongst scientists. Sun and Naveh (2009) build on Gilbert (1997) by replacing the structural system with a regime featuring agent cognition. Agents actively identify lucrative strands of research non-randomly, establish private valuations for different research outputs, and either become successful or die out. This approach produces similar results as Gilbert’s original model within a more realistic (and flexible) framework. Edmonds (2007) also builds on Gilbert (1997), but uses a system of logic rules to create scientific knowledge. Agents pull items of knowledge in the form of theorems from a public repository, apply a rule of logic for creating new theorems, and then publish a set of their results via a public repository (or ‘journal’). All these studies generally focus on the social processes that *produce* science and understate the impact of the construction of knowledge on social outcomes.

I add to this literature by focusing specifically on the development of the scientific community instead of the actual creation of science to extract the relationship between social behaviour and scientific practice. To do this, I impose a group structure on agent behaviour in the virtual environment inspired by the work of Lakatos (1969, 1970) on scientific research programs (SRPs).<sup>6</sup> I add agent-to-agent interaction to disseminate research results and to determine proper research practices, which forms the socio-cultural (institutional) foundation of science. In the model, virtual agents receive private, incomplete information from one another about the success potential of each SRP. Based on this information, they choose which SRP to join. They perform science and contribute to the success of their chosen SRP. Research results are sent out to a subset of other scientists, who then update their private perception of each SRP's success potential. This process then repeats. In the model, I allow agents to choose from different 'types' of research practices: imitation, innovation and adaptation (or mutation). I allow the scientific community to assign prestige to each practice, which affects the type chosen by the agent. During the simulation, I record the number of agents within each SRP, the social value of each agent type, and how successful science becomes as SRPs grow, develop and die.

### **3 The Agent-based Model**

I populate a virtual research environment with scientists. These scientists engage in research activity within scientific research programs in order to match stylized facts. There are 5 significant pieces of the simulated model:

1. How the virtual world is designed.
2. How science is performed.
3. How information is spread amongst agents.
4. How agents choose research approaches.
5. How agents select and join SRPs.

#### **3.1 The Virtual World**

The virtual world is populated by  $I > 1$  agents indexed by  $i \in \{1, \dots, I\}$ . Agents perform science in an attempt to replicate a stylized fact. Denote the number of existing stylized facts as  $D > 1$ , indexed by  $d \in \{1, \dots, D\}$ . Science is performed within  $J \geq 1$  scientific research programs, indexed by  $j \in \{1, \dots, J\}$ . Each agent is born with the following:

- **Innate Ability.** Agent  $i$  can reproduce a stylized fact,  $d$ , with probability  $p_{i,d}$  independent of their SRP. This ability is a random draw from a  $U[0,1]$  distribution which is given to the agent at the start of the simulation.
- **Mailboxes.** Agent  $i$  can access a mailbox,  $M_i$  which contain scientific reports. The mailbox is divided into pigeonholes, one for each SRP ( $M_{i,j}$ ), and facilitates record-keeping by individual agents. The reports are used by the agent privately to assess the success potential of each SRP in matching any given stylized fact.

Agents maintain their perceptions about the research environment privately through knowledge of their own capabilities and knowledge of a *subset*<sup>7</sup> of research done in the field.

Some unobserved information about SRP performance must be recorded and utilized during the simulations. To do this, the following structures are used:

- **SRP Research Attempts.**  $Attempts_{j,d,t}$  = the total number of attempts that members of some SRP,  $j$ , made to match stylized fact,  $d$ , in period  $t$ .
- **SRP Research Successes.**  $Success_{j,d,t}$  = the number of times members of some SRP,  $j$ , successfully matched a stylized fact,  $d$ , in period  $t$ .
- **SRP Membership.**  $Membership_{j,t}$  = the total number of researchers that are members of SRP,  $j$ , in period  $t$ .  $SRP_{i,t}$  = the SRP that agent  $i$  is a member of during period  $t$ .

## 3.2 The Structure of Science

Lakatos' original description of SRPs (the core and the heuristic) is generalized using a framework of 'random science'. This approach is similar to that of Peirce (1967) and provides a convenient way to model the actual practice of science within an SRP.<sup>8</sup> Assume that when a testable model is generated from core axioms, it succeeds in matching a given novel fact with some probability. Denote the probability that practice within and SRP,  $j$ , matches stylized fact,  $d$  in period  $t$  as  $p_{j,d,t} \in [0,1]$ . Valuing a scientific success at unity and a failure at 0, the expected *value* of the SRP at any period is then  $\sum_d p_{j,d,t}$ . Progress of the SRP is then measured by both an increase in the number of stylized facts they address and an increase in the probability of matching the data. Note that one could also interpret  $p_{j,d,t}$  as an abstract measure of scientific inexactitude, in which case progress of the SRP is measured by a reduction in inaccuracy over broader range of novel facts.

For any given SRP, the success probability for each stylized fact is updated at the end of each period based on the experimental successes and failures of the associated researchers. It is assumed that updated success probabilities are constructed using equation (1):

$$p_{j,d,t+1} = \gamma p_{j,d,t} + (1-\gamma) \frac{\sum_t \text{Success}_{j,d,t}}{\sum_t \text{Attempts}_{j,d,t}} \quad (1)$$

where  $\gamma \in [0,1]$ . Equation (1) describes updated success probabilities as a convex combination of the past success probability and the current accumulated success rate. It is assumed that a researcher attempts a stylized fact with some exogenous probability,  $\delta \in (0,1]$ . At any period, the number of attempts for any stylized fact,  $\text{Attempts}_{j,d,t}$  depends on the probability,  $\delta$ , and the number of researchers in the SRP at the time. The number of successes,  $\text{Success}_{j,d,t}$ , depend on research outcomes which are partially (but not completely) influenced by the SRP's success probability,  $p_{j,d,t}$ . If  $\gamma = 1$ , the success of the core is predestined and cannot be improved beyond its design. If  $\gamma = 0$ , the success of the core is entirely determined by the practice of science itself. When  $\gamma \in (0,1)$ , practice augments science sluggishly over time. Since the state of research influences the probability research is successful and vice versa, the process generating scientific success features history-dependence. The 'protective belt' of the SRP is defined as a record of success probabilities, a record of research successes and a record of research attempts for each stylized fact which are updated each period.

### 3.3 Making Science

#### 3.3.1 Agent Types and SRP Choice

When conducting research, each agent in the virtual world will use a protocol ( $P_{i,d,t}$ ). The protocol they use will depend on the SRP they chose and the type of researcher they decide to be. A protocol maps a specific research practice to the probability that some researcher,  $i$ , will successfully reproduce stylized fact,  $d$ , in period  $t$ .

I consider three agent types which are consistent with stability of the SRP paradigm: imitators, innovators and mutators. Agents with these features are somewhat common in computational studies of learning and

evolution in the social sciences (see Nelson and Winter (1982) or Macy and Willer (2002) for examples) and aid in the stability of the SRP paradigm. Imitator agents act as follows:

1. Using received reports from the mailbox, form private (perceived) success rates for each SRP for each stylized fact ( $\tilde{p}_{i,j,d,t}$ ).
2. Valuing successes at unity and failures at zero, derive the expected perceived value of each SRP ( $\text{Value}_{i,j,t} = \sum_d \tilde{p}_{i,j,d,t}$ ).
3. Choose the SRP with the highest expected value,  $j^*$ .
4. Perform science by following the SRPs protocol ( $P_{i,d,t} \rightarrow p_{j^*,d,t}$ ). Note that *perceived* SRP success rates determine the protocol, but the protocol's use in practice involves the *actual* success probabilities of the SRP which is unobserved to the agent.

Imitator agents generate inflow of interest for relatively successful SRPs and outflow for relatively unsuccessful ones. In other words, they jump on the best bandwagon they have knowledge of. Whether or not the SRPs are 'progressing' or 'degenerating' then depends not only on the success probabilities of the SRP but also on the number of new researchers (or, abstractly, new research projects) that it attracts. Imitation, however, eventually leads to a stale state of science. New ideas or techniques never enter the SRP, therefore the potential for matching new novel facts is constrained and improving scientific success is limited.

To add the potential for semi-stochastic evolution of SRP successes, innovator agents are incorporated. Innovator agents behave in a similar way to imitator agents with the exception that they only partially follow the SRP protocols. They use a mixture of their natural ability with the SRP to perform science:

1. Using received reports from the mailbox, form private success rates for each SRP for each stylized fact ( $\tilde{p}_{i,j,d,t}$ ).
2. Valuing successes at unity and failures at zero, derive the expected value of each SRP when mixed with private ability ( $\text{Value}_{i,j,t} = \sum_d [\lambda \tilde{p}_{j,d,t} + (1-\lambda)p_{i,d}]$ ,  $\lambda \in (0,1)$ ). The SRP choice by innovator agents would be identical to that of imitator agents, since private ability is independent of the SRP.
3. Choose the SRP with the highest estimated point yield,  $j^*$ .
4. Perform science using a convex combination of the SRPs protocol and private ability ( $P_{i,d,t} \rightarrow \lambda p_{j^*,d,t} + (1-\lambda)p_{i,d}$ ). Again, the protocol's use in practice involves the *actual* success probabilities of the SRP which is unobserved to the agent.

Although innovator agents introduce the potential for improving SRP success, they may diminish it with equal likelihood. Because of this, it is extremely unlikely that the protocol for innovators will ever yield perfect success (even when current success probabilities for an SRP is 1) or perfect failures (even when current success probabilities for an SRP is 0). We therefore expect the protocol used by innovators in practice to generate, on average, a lower success probability of matching novel facts than imitators since private ability is a random draw from a distribution with a mean of 0.5.<sup>9</sup> In addition, researchers are more likely to choose innovation when science is young or if the SRP is not well developed. Few stylized facts have been attempted at the start, so the agent's perceived success from including some of their private ability is greater than without. Together, the distortionary nature of innovation and its appeal early in science suggest that science may become limited in early development.

To improve the state of science, mutator agents are next added. These agents mix the most successful features of a randomly chosen SRP with their own. The end result is a scientific protocol that yields an expected point score that is at least as good (if not strictly better) than either SRPs. They then adopt the SRP they've used most often in their mix:

1. Using the SRP they joined in the previous period,  $j^*$ , and a randomly selected SRP,  $j'$ , recombine protocols to maximize expected success for each stylized fact.
2. Adopt the SRP used most often across stylized facts
3. Perform science using the new protocol ( $P_{i,d,t} \rightarrow \max\{p_{j^*,d,t}, p_{j',d,t}\}$ ).

Intuition suggests that a world consisting entirely of mutator agents would yield the highest success probability for any stylized fact on average. The presence of mutator agents allows SRPs to adapt to criticisms or failures by absorbing the successes of others. The notion of mutation implemented here also allows for the possibility of intra-disciplinary research to enhance scientific successes.

Agent types as described above relate to the role that different sorts of research play in the development of science. Imitator agents represent the status quo, which may be preferable if success probabilities are already high. Innovators either improve or diminish the current state of science depending on their natural abilities. Mutator agents drive scientific progress through adaption and intra-disciplinary research.

### 3.3.2 Performing Science and Disseminating Results

Once agents have decided which type of researcher to become and which SRP to join, matching a set of stylized facts is attempted. First, whether or not the agent chooses to attempt a stylized fact at all is the result of a random draw from a uniform distribution (recall the probability of attempting to match any stylized fact is  $\delta$ ). If a novel fact is attempted, scientists receive either a success result or a fail result when doing research. This result is the product of yet another random draw from a uniform distribution. To determine the probability of receiving a success result for an attempted stylized fact,  $d$ , the probability associated with the researcher's protocol,  $P_{i,d,t}$ , is used (recall this probability relates to the actual success probability of their chosen SRP as described by equation (1)). If the outcome of the draw is less than the value of the protocol, the researcher earns a 'success' result. If not, the researcher earns a 'fail' result. For every success they achieve, they receive a 'point'. Points accumulate for the researchers through the duration of the simulation and reflect the agent's success or prestige.

After performing science, all results (attempts, successes, failures) are recorded by the SRP. Probabilistic success of the program is updated using the results from performed experiments and equation (1). Scientists draft reports that *only contain successes* which they then mail out to colleagues. Other scientists receive the reports in the appropriate pigeonholes in their private mailboxes with probability  $\mu \in (0,1]$ . This process then repeats. Private information, which individuals use to make decisions on type and SRP membership, is an incomplete subset of global information produced by the scientific community. Further, scientific success is itself a function of past success and agent types.<sup>10</sup> These two mechanisms induce feedback between individual agents and the virtual environment.

### 3.4 Agent Interaction

Agents choose between imitation, innovation and mutation through interactive behaviour. While it is potentially optimal to ask agents to make this decision by maximizing their expected point score based on their private information, this assumption would sterilize the model of social activity. Instead, agents will both optimize expected scores *and* engage in social learning to determine their type. Social learning is implemented using a rule and meta-rule. The rule is:

1. Weight the scores for each agent type,  $\tau \in \{\text{imitator, innovator, mutator}\}$  by  $\omega_{i,\tau,t}$ .
2. Choose the type with the highest expected weighted score.<sup>11</sup>

The rule for each agent is a triple and a procedure for calculating weighted expected scores. To follow the rule, each agent must be born with a *perception* of weights. For simplicity, it is assumed that agents are born with a random draw from a uniform distribution for each agent type:  $\omega_{i,\tau,0} \sim U[0,2]$ . The average perceived weight at the outset for each type is unity. In lieu of standard points for successes, scientists receive ‘weighted points’ based on the type that they choose. This can loosely relate scientific prestige to research activity and, in turn, affect the agent’s type decision. If, for example,  $\omega_{i,imitator,t} < \omega_{i,mutator,t}$ , then any stylized fact achieved by an imitating agent would receive less prestige than the same stylized fact matched by a mutating agent for any agent,  $i$ . A mutator is thus a relatively more desirable vocation for the agent. At the end of each period, once science has been completed, an average weight for each type will be computed to represent a community-wide measure of prestige:

$$\Omega_{\tau,t} = (1/I) \sum \omega_{i,\tau,t} \quad (2)$$

Agents will then receive an actual weighted score (real prestige) depending on their scientific successes, the type they have chosen and the associated community-wide weight described by equation (2). When  $\Omega_{\tau,t} = 1$ , prestige is based purely on research output (i.e. a success is a success regardless of the type of research done).

Rules are updated according the following meta-rule:

1. When called to choose type, select another agent at random,  $i'$ .
2. Ask  $i'$  what their weighted score was in the previous period.
3. If  $i'$  received a higher weighted score, adopt their perceived triple from the previous period, else use own triple from the previous period.

Note that the agent only asks his colleague what his score was and what his perception of the available type was, never about the private information of his peer. This activity can be thought of as pedagogy or networking. The evolution of perceived weighting schemes thus depends on the scientific successes and weighting regimes of previous periods making the meta rule history-dependent. By looking at the relative values of  $\Omega_{\tau,t}$ , we can evaluate how researchers learn to adopt different types over time and whether or not preferences over types benefits or harms scientific success.

## 4 Simulations

### 4.1 Baseline Simulations

Figure 1 briefly describes the process of making science for a typical agent.

Figure 1 here: The Structure of the Agent-based Model

Before performing simulations, parameters and initial values in the model must be calibrated. First, the initial set-up of the model is generated:

1. Agent ability for any given stylized fact,  $p_{i,d,0}$ , is a random draw from a uniform distribution,  $U[0,1]$ .
2. A subset of researchers, calibrated to 10% of the total population, are selected at random to start SRPs. They use their private abilities to set the SRP's initial success rates.
3. Structures are created to hold data for all researchers and the SRPs.
4. Initial researchers perform a round of research, update the success probability of the SRP, and mail out results. Initial researcher type is set as 'imitator'.

Given this set up, we would expect the initial average success rate of SRPs at the start of the simulation to be approximately 50% across stylized facts.

There are many other parameters within the ABM that the programmer can adjust to perform 'experiments'. In the ABM described above, they are  $\delta$ ,  $\mu$ ,  $\gamma$ ,  $\lambda$  and the weighting regime. These are left as free parameters. One can also experiment with the length of the simulation and the number of agents, although these tend to increase the amount of processing time. 200 simulations, each consisting of 100 agents who live for 200 period are conducted in the experiments.<sup>12</sup> I set the total number of stylized facts available during the simulation to 100 and consider a benchmark calibration with  $\delta = 5\%$ ,  $\mu = 50\%$ ,  $\gamma = 0.5$  and  $\lambda = 0.75$ . The 'endogenous variables' of the model are the ending number of SRPs, the ending average success probabilities of surviving SRPs and the average ending social values assigned to each type. The ending number of SRPs after 200 periods captures the degree of pluralism which exists within the scientific community in the presence of each type of agent, with 1 being a state of *consensus*. The average success probability (across stylized facts) for surviving SRPs is reported which measures the 'progress of science' given the actions of the

agents. The average ending rule (across researchers) reports the value the scientific community places on each research practice.

#### **4.1.1 Preset Agent Types**

I first predetermine the agent types to ensure the model produces expected behaviours before adding agent-to-agent interaction. When agent types are fixed, the researcher's only decision is which SRP to join. Simulation results indicate that mutation yields the highest average probability of scientific success, followed by imitation and then innovation. This suggests that mutation is the most preferred type (in terms of success probability) while innovation is the least preferred type. The simulation data also indicate that innovation induces the largest degree of pluralism, followed by imitation and then mutation. Mutation results most often in consensus and achieves unanimity faster than the other two types.

Table 1, columns [1-3], and figures 2 and 3 report the simulation outcomes. Figure 4 shows 'typical' simulations for the preset types. The results for average ending SRP success are expected. In the random science algorithm, we would expect imitator agents to achieve scientific success greater than 50% (since they choose the best SRP to imitate) but not by very much (since the number of stylized facts to match is large). Innovators impose their private ability on the SRP protocols used by imitators. Since their ability is a random draw from a uniform distribution with a mean of 0.5, we would expect the innovator protocol to, on average, perform worse than imitator protocols and pollute science. Since mutator agents take the best core processes of existing SRPs and incorporate them into their own, we would expect relatively high success.

The results for number of ending SRPs are also somewhat expected. Mutator agents produce highly successful core processes and therefore attract researchers quickly. This results in frequent consensus. Imitator agents move to the most productive research programs resulting in eventual consensus. The process, however, takes more time than mutation. Innovation induces persistent pluralism via independent action.

#### **4.1.2 Endogenous Agent Types**

When allowing agent types to be endogenously chosen based on rules and social learning, I focus on two different methods for updating private weights on different types of research. First, I consider the case where agents update their private weights based on received points for successes taken at their face value (or 'unweighted scoring' - a 1 for each success and a 0 for each failure). In this scenario, only the number of successes matters when agents trade their rules on weighting research types. Over time, the individual values

assigned to different research activities (and thus, the social values) reflect those of the most successful practitioners. Scientific success impacts the values of the research community, but not vice versa.

I then consider the case where private weights are updated based on prestige points (or ‘weighted scoring’). In this scenario, agents receive points for successes based on the type of research they’ve chosen and the *social value* of that type. Researchers essentially ask one another how much prestige they’ve accumulated over time and use this information to set their rules on weighting research types. Over the course of the simulations, individual values assigned to different types of research (and thus, the social values) reflect both the values of the most successful practitioners and how the scientific community assigned prestige to different practices in the past. Scientific success impacts the values of the research community and vice versa in this case. It is also possible for the scientific community to maintain a research practice with less success potential, provided it receives enough peer esteem.

Table 1, columns [4-5] along with figures 2 and 3 report the results. Figure 4 shows ‘typical’ simulations for the endogenous types and how the population of researchers are divided into the available SRPs over time. We can compare results from each weighting scheme to those found under fixed agent types to identify the impact of social organization on the construction of science. Many striking relationships are indicated by the simulations.

The scientific community places a low weight on innovation, a medium weight on imitation and a high weight on mutation. This result is as expected since mutation yields high success rates and innovation can pollute science. The community as a whole, therefore, favours types associated with high scientific success and abhors types associated with low scientific successes by the end of the simulation period. This occurs in both weighting schemes, indicating that the scientific community reinforces good practices with high prestige even though the possibility that the society places a high weight on inferior practices is present.

Simulations will likely end with a population of mutators, but will likely start with a population of innovators. This is due to the relative attractiveness of the innovator type when science is young. When few stylized facts have been matched, a researcher tends to believe his private ability can yield more successes on average than the SRP alone and, as a result, innovates. Over time, as private knowledge accumulates, it becomes relatively more lucrative to switch to mutation as a means for achieving scientific success. Figures 5 and 6 corroborate these findings.

Table 1 and figure 3 suggest that ending success rates are 0.542 (unweighted) and 0.528 (weighted) on average with type endogeneity, which slightly higher than those for either imitator or innovator presets but far lower than that for the mutator preset. Early innovation results in a significant reduction in scientific success potential, even if mutation is prevalent later on.

Table 1 and figure 2 show how often simulations result in consensus. Consensus resulted in the majority of simulation (67% of simulations for unweighted scoring and 52% for weighted scoring). Pluralism occurs in 11% or fewer simulations in any specific form (in bins of 2 - 9 SRPs, a maximum of 11% in each bin). When types are endogenous, the simulation results in a unique surviving SRP more often than any other individual sort of pluralism. While ending with multiple SRPs is frequent when types are endogenous, no specific number is particularly common (unlike type presets where ending with 2, 3 or 4 SRPs is common for mutators and imitators). It might be the case that agent type selection acts as a mechanism for achieving consensus quickly. In cases where consensus is never reached, it is due to the influence of innovator's independence on scientific practice.

Table 1 here: Simulation Results

Figure 2 here: Histogram - Ending # SRPs

Figure 3 here: Histogram - Ending Avg. Success Rates

Figure 4 here: Simulation Samples - SRP Membership

Overall, the simulations suggest that there is a negative relationship between pluralism and average scientific success rates (as reported in table 2). There is, in turn, a negative relationship between the ending number of SRPs and the prevalence of the mutator type and a positive relationship between the ending number of SRPs and the prevalence of the innovator type. Further, higher scientific success rates are positively associated with simulations containing mutators by the end of the simulation period while lower scientific success rates are associated with simulations containing innovators by the end of the simulation period. The social structure of the scientific community weights scientific successes appropriately to guide practice towards the most successful techniques, which most often results in consensus through intra-disciplinary approaches.

Figure 5 here: Population of Types - Unweighted Sum Score Sets Rule

Figure 6 here: Population of Types - Weighted Sum Score Sets Rule

Table 2 here: Simulation Results

## 4.2 Robustness

### 4.2.1 Augmented Parameterization

To test the models performance under different parameter assumptions, the following additional simulations are performed:

1. A benchmark calibration with endogenous types and a weighted period score affecting the agent's rule.
2. High information frictions (low probability of receiving mail,  $\mu = 0.1$ ).
3. Low information frictions (high probability of receiving mail,  $\mu = 0.9$ ).
4. Sluggish adaption of proven methods (low impact of experiments on SRP success rates,  $\gamma = 0.9$ ).
5. Quick acceptance of proven methods (high impact of experiments on SRP success rates,  $\gamma = 0.1$ ).
6. Low impact of innovator ability on scientific success ( $\lambda = 1$ ).

Because the scheme in which the 'prestige' associated with a particular type of agent is most interesting, this case will serve as the benchmark (and correlates to column [5] of table 1). Recall that this model produces a fairly high instance of consensus relative to all other forms of pluralism, but with a low ending success probability.<sup>13</sup> The simulation results for these alternative calibrations appear in table 3 and in figures 7 and 8.

Table 3 here: Simulation Results - Robustness

Figure 7 here: Histogram - Robustness - Ending # SRPs

Figure 8 here: Histogram - Robustness - Ending Avg. Success Rates

Agents in the model react to changes in these parameters exactly as expected. The occurrence of consensus is reduced when there are information frictions (low probability of receiving a report) or when science adjusts sluggishly to new research. As the amount of information that scientists receive about contemporary experiments rises, there are reductions in the *degree* of fragmentation (the number of ending SRPs). Simulations that involve increasing the degree to which an SRP's protocols adjust to new research do not generate results which dissent from the benchmark model. In all of these cases, the ending success rate is similar to that of the benchmark (approximately 0.53).

There is improved frequency of consensus and a higher ending success rate for science (0.575) when innovator ability has less of an impact on research. When the impact of one's own talents is low we can expect a faster adoption of intra-disciplinary techniques. In other words, scientists still prefer to become innovators when science is young but switch to mutators more quickly as research is published. This results in higher scientific success rates and less fragmentation within the field and is achieved with less 'downward pressure' on the social value of the innovator type. Innovators act more similarly to imitators as their influence on science falls.

#### 4.2.2 Augmented Initial Values

In all simulations, there is a correlation between the age of science and the development of the social weighting regime. Recall the initial assumption: for each agent type,  $\tau = \text{imitator, innovator, mutator}$ , the starting weighting rule for any agent  $i$  in the initial period is  $\omega_{i,\tau,t} \sim U[0,2]$ . When science is new, agents always have evenly distributed initial weights pre-assigned to them which are random draws from a  $U[0,2]$  distribution. From the benchmark calibration described in the previous section, that the ending average weight for imitators, innovators and mutators respectively were 0.975, 0.503 and 1.106. Over time, society has come to value mutation and punish innovation. We might be interested in the results from a model in which researchers are indoctrinated with a specific distribution of social weights *prior* to the start of science. Using the results from the benchmark case, I impose the following initial distribution of weights for the agents:

$$\omega_{i,\text{imitator},0} \sim U[0,1.95]$$

$$\omega_{i,\text{innovator},0} \sim U[0,1.006]$$

$$\omega_{i,\text{mutator},0} \sim U[0,2.212]$$

With this calibration, agents start science with the idea that innovation is a less-preferred type. The results are reported in table 4 and figure 9. As expected, early innovation still occurs but is eliminated quickly by agents. The average number of ending SRPs is low (1.39 with a standard deviation of 0.64) and the average ending scientific success rate is high (0.70 with standard deviation of 0.05). These results are closer to those from an environment consisting of only mutators. More strikingly, the value of the mutator type is *positively* correlated with the value of the other two types, whereas it was negatively correlated in the benchmark calibration. In other words, it is no longer necessary for the scientific community to place such a high weight on mutation and a low weight on innovation to secure high scientific success; resetting the initial values has accomplished much of the work of the scientific community. This highlights the role of the scientific culture in constructing knowledge: to determine and encourage the practices that result in successfully explaining novel facts when inferior techniques are available.

Table 4 here: Simulation Results - Augmented Initial Values

Figure 9 here: Histogram - Robustness - Ending Avg. Success Rates

## 5 The Agent-based Model and Scientific Progress

What does the agent-based model say about scientific progress given that the actual practice of science is suppressed within the ‘random science’ algorithm? The simulations suggest that social interactions between agents facilitate the construction of private information (via disseminating work) and social valuation of different types of research activity (via exchanging rules for choosing agent types). These social activities influence how agents perform research and how scientific research programs develop. As a result, innovation and pluralism are prevalent early in science while intra-disciplinary approaches and consensus are a learned phenomena eventually emerging over time. It can be argued that this pattern enriches the ideas of Kuhn (1970) on paradigms of science even though the model abstracts from the actual construction of knowledge.

In Kuhn’s model, a community of researchers perform scientific activities in an effort to understand a selection of novel facts. Over time, this community establishes a research tradition, which defines a state of “normal science”. The practice of normal science attempts to construct matches between the paradigm’s

predictions and the data, but inconsistencies arise. These inconsistencies may be ignored early on, but as they accumulate they begin to undermine the scientific tradition which defines the paradigm. Eventually, a critical point is crossed, the prevailing paradigm is rejected (a crisis), and a new one takes its place.

Because social phenomena are a product of social systems, which themselves are continuously changing and evolving, we would expect new stylized facts to reveal themselves in the social sciences. In the computerized model above, the crisis point can be thought of as the early period of the simulations when few stylized facts are explained. During this time, researchers are more likely to explore a variety of research programs and engage in innovation. Over time, the SRPs that best capture these stylized facts adapt and survive while those that are unpopular or have low success potential die. The state of normal science can be interpreted as the state of consensus (a unique SRP) that eventually results at the end of the simulation. This consensus features a particular scientific tradition which is constructed by pulling approaches from other SRPs (mutation) in an attempt to maximize the program's success potential. Eventually, new stylized facts come to light and the process repeats.

## **6 An Application: Modern Macroeconomic Theory**

The agent-based model of scientific research programs developed in this paper can be used as a device for reflecting upon trends in current practice. One of the most poignant examples of conflict and consensus amongst researchers in the field of Economics has been within the discipline of macroeconomics. A brief history of macroeconomic theory and how it has changed during the past century can provide a colloquial substantiation of the simulation model described in the previous sections.

This 'story of macroeconomics' is a common one among economists.<sup>14</sup> Prior the 1930s, the received macroeconomic models of the day emphasized the self-correcting nature of the economy (so-called 'Classical macroeconomics') and the immediate adjustment of prices to clear markets. These models were unable to explain the severity and the duration of unemployment that had defined the Great Depression which had spread across the globe during the 1930s. By allowing for price rigidities in models of economic performance, John Maynard Keynes was able to show that the economy could persist in a state of recession, warranting policy intervention by the state. Keynes and his students developed an approach in which assumptions about the relationships between economic variables were incorporated into a mathematical model of the economy which was then empirically estimated. This model was shown to fit the macroeconomic data well and sparked

nearly 40 years of research in the field of business cycle research. These early innovations dramatically deepened our understanding of how the economy behaves dynamically.

In the 1950s and 1960s, economists became dissatisfied with the Keynesian models and the policies they supported. When a group of economists (the Monetarists) criticized the Keynesian's reliance on fiscal policy and claimed that monetary policy could have substantial impacts on the economy in the short run (a possibility that was omitted from the original Keynesian framework), Keynesian models began to adapt to accommodate a stronger role of money. Also at this time, some noted the models were constructed using 'ad hoc' assumptions about the relationships between economic variables which were inconsistent with behaviour at the micro level. When Keynesian models started to fall to pieces in the 1970s, being unable to explain and resolve the high inflation and high unemployment of the Oil Shock years, many economists returned to theories consistent with classical microeconomic theory which incorporated more flexible prices ('Neo-Classical macroeconomics').

The Neo-Classical research program (again, an innovation) emphasized the construction of models from the fundamentals of economic behaviour ('micro-foundations') and involved wide-scale market clearing within the economy (general equilibrium). This process was aided by assuming that agents within the economy were rational and could form rational expectations about future variables, making market clearing a dynamic process during business cycle movements. These models employed calibration techniques to compare simulated results to the real economy. The Neo-Classical SRP grew alongside the struggling Keynesian SRP throughout the 1980s and 1990s. Many Keynesian practitioners responded to the criticisms made by Neo-Classical researchers by adapting Keynesian models to include rational agents and micro-foundations (so-called 'Neo-Keynesian macroeconomics'). By finding logical reasons why markets may not clear (efficiency wages, menu costs, market power, etc.), Neo-Keynesian economics built elaborate micro-founded models of the short-run economy which were capable of matching many stylized facts seen in macroeconomic data.

Since, a class of models built on micro-foundations designed to study the economy over the business cycle (known as 'dynamic-stochastic-general-equilibrium' models or DSGE) has dominated the field. These models are coupled with empirical underpinnings and are often relied on for policy analysis. Although there are still differences between Neo-Classical DSGE models and Neo-Keynesian DSGE models when it comes to assumptions about price flexibility, both model types are built on similar assumptions about basic economic behaviour and are accepted by the research community as 'appropriate'. Both calibration and empirical

estimation is employed in testing how well these models fit the data, resulting in a fairly complete picture of economic activity at the aggregate level.

As the history suggests, innovation and fragmentation of the scientific community is common at the out-set, after new stylized facts come to light (the start of the Keynesian SRP after the Great Depression, the re-emergence of the Classical SRP during the Oil Shock years). When innovation no longer guaranteed a high probability of scientific success, researchers came to prefer synthesis. By recombining the successful traits of known SRPs, better science is produced. Since the 1950's, the Keynesian SRP has continued to adapt to criticisms by absorbing the approaches of others (the Monetarist's role of money, the micro-foundations of the Neo-Classical SRP). In the modern era, the development of a 'synthesis' between Keynesian and Classical models to explain periods of both price rigidity and price flexibility, and the emergence of the DSGE framework as a unified practice in both fields. As suggested by the simulation model, crisis leads to a period of pluralism and innovation, followed by consensus and mutation. These are brought about via social interaction that takes place among researchers; through the 'culture of science'.

## **7 Conclusion**

By utilizing a framework such as agent-based modelling for evaluating methodological theories<sup>15</sup>, those who believe that action and practice are, in part, a social construct can study the implications of their ideas by creating and experimenting with simulated realities. Methodological theory comes to life within a virtual environment and it becomes possible to study the sorts of behaviours we as researchers exhibit from the 'outside'. As agent-based models permeate many fields within the social sciences, this technique becomes an increasingly popular method for the study of societies.

The goal of the agent-based model in this paper was to pin down the relationship between social behaviour and scientific practices. The Lakatosian SRP gives structure to the scientific community. Agent-to-agent interaction is added to disseminate research results and to establish norms and values over different research techniques. In the model, an SRP must attract scientists and produce frequent matches to a series of stylized facts to progress. Which SRP a scientist chooses to join depends on their private assessment of each program's success potential. Further, the success of the research a scientist conducts depends partially on the program they choose and partially on the type of research activity selected (join the bandwagon, innovate, mix successful approaches). Social dealings between members of the scientific community affect the private information of the researchers and value placed on different activities. As a result, research programs develop

over time in a history-contingent fashion. Further, social values for different research activities evolve so that the most lucrative actions receive the most prestige by the scientific community.

Although the actual construction of knowledge is suppressed within the 'random science' algorithm, the computerized model tells us a story about how science develops as a social phenomenon. Intra-disciplinary research is preferred to imitation and innovation within a field of study. However, early in science when there are many novel facts to explain, the scientific community is fragmented and researchers often choose to innovate on different techniques (although this may result to a worsening of scientific success). Over time, a dominate research program will frequently emerge in which the agents practice intra-disciplinary approaches; a state of consensus. When a large number of new facts to explain emerge, the process repeats. This story is consistent with both past methodological theory and some colloquial evidence from the field of Economics. Agent-based modelling provides a flexible approach to the study of social organization and allows us to evaluate our ideas on the progress of science from a new perspective.

## References

- Blackhouse, R. E. (2004), 'The methodology of scientific research programs', in *The Elger Companion to Economics and Philosophy*, eds. J. B. Davies, A. Marciano, and J. Runde, Edward Elgar: Northampton, 181-196.
- Blaug, M. (1992), *The Methodology of Economics*, 2<sup>nd</sup> ed., Cambridge University Press: Cambridge.
- Eastman, T. E and Mahootian, F. (2009), 'Complementary frameworks of scientific inquiry: Hypothetico-deductive, hypothetico-inductive and observational-inductive', *World Futures*, 65(1), 61-75.
- Edmonds, B. (2007), 'Artificial science: A simulation to study the social processes of science', in *Social Simulation: Technologies, Advances and New Discoveries*, eds. B. Edmonds, K. G. Troitzsch, and C. H. Iglesias, IGI Publishing: Hershey, 61-67.
- Epstein, J. M. (1999), 'Agent-based computational models and generative social science', *Complexity*, 4 (5), 41-60.
- Gilbert, N. (1997), 'A simulation of the structure of academic science', *Sociological Research Online*, 2, URL <http://www.socresonline.org.uk/2/2/3.html>.
- Hands, D. W. (1994), 'The sociology of scientific knowledge', in *New Directions in Economic Methodology*, ed. R. E. Blackhouse, Routledge: New York, 75-106
- Hands, D. W. (2004), 'Constructivism: The social construction of scientific knowledge' in *The Elger Companion to Economics and Philosophy*, eds. J. B. Davies, A. Marciano, and J. Runde, Edward Elgar: Northampton, 197-212.
- Hands, D. W. (2008), 'Popper and Lakatos in economic methodology', in *The Philosophy of Economics: An Anthology*, 3rd ed., ed. D. M. Hausman, Cambridge University Press: Cambridge, 188-203.
- Hodgson, G. M. (1998), 'The approach of institutional economics', *The Journal of Economic Literature*, 36(1), 166-192.
- Jalladeau, J. and Kuhn, W. E. (1978), 'Research program versus paradigm in the development of economics', *Journal of Economic Issues*, 12(3), 583-608.

- Keesing, R. M. (1974), 'Theories of culture', *Annual Review of Anthropology*, 3, 79-97.
- Kuhn, T. (1970), *The Structure of Scientific Revolutions*, 2<sup>nd</sup> ed., The University of Chicago Press: Chicago.
- Lakatos, I. (1969), 'Criticism and the methodology of scientific research programmes', *Proceedings of the Aristotelian Society*, 69, 149-186.
- Lakatos, I. (1970), 'Falsification and the methodology of scientific research programmes', in *Criticism and the Growth of Knowledge*, eds. I. Lakatos and A. Musgrave, The Cambridge University Press: Cambridge, 91-196.
- Landreth, H. and Colander, D. C. (2002), *History of Economic Thought*, 4<sup>th</sup> ed., Houghton Mifflin: Boston.
- Macy, M. W. and Willer, R. (2002), 'From factors to actors: Computational sociology and agent-based modelling', *Annual Review of Sociology*, 28, 143-166.
- Mankiw, N.G. (2006), 'The macroeconomist as scientist and engineer', *The Journal of Economic Perspectives*, 20(4): 29-46.
- McCloskey, D. N. (1983), 'The rhetoric of economics', *The Journal of Economic Literature*, 21(2): 481-517.
- Morgan, M. S. (2009), 'Models, stories and the economic world', *The Journal of Economic Literature*, 8(3), 361-384.
- Nelson, R. N. and Winter, S. G. (1982), *An Evolutionary Theory of Economic Change*, Harvard University Press: Harvard.
- Peirce, C. S. (1967), 'Note on the theory of the economy of research', *Operations Research*, 15(4), 643-648.
- Shapere, D. (1964), 'The structure of scientific revolutions', *The Philosophical Review*, 73(3), 383-394.
- Sun, R. and Naveh, I. (2009), 'Cognitive simulation of academic science', *Proceedings of the 2009 International Joint Conference on Neural Networks*, 418-424.
- Tesfatsion, L. (2002), 'Agent-based computational economics: Growing economies from the bottom up', *Artificial Life*, 8(1), 55-82.

## Notes:

<sup>1</sup>Department of Economics, University of Otago, dan.farhat@otago.ac.nz. Special thanks to Viktoria Kahui and Shruti Kapoor for helpful comments. All errors are my own.

<sup>2</sup>Eastman and Mahootian (2009) provide a concise description of the hypothetico-deductive model: [A] First observe phenomena from data. [B] Construct testable hypotheses explaining the phenomena from accepted axioms. [C] Design experiments to test the hypotheses. [D] If the hypothesis is falsified, repeat [B]. Otherwise, await corroboration via [C]. In lieu of repeated experimentation, as is done in the physical sciences, parts [C] and [D] in the social sciences usually involve designing some sort of empirical investigation or qualitative assessment.

<sup>3</sup>Focusing on the social framework within which scientific knowledge is derived, we can link together methodological inquiry with recent renewed interest in institutionalism, or ‘neo-institutionalism’. Taking general ideas from biology, neo-institutionalist approaches to understanding behaviour views individual action as dependent on habits, rules and their evolution, all of which are a function of shared psychology and society. See Hodgson (1998) for details.

<sup>4</sup>See Tesfatsion (2002) for a description of how this approach has been developed in the literature. The ABM method is also known as a ‘generative approach’ to social science (Epstein (1999)).

<sup>5</sup>Metarules are themselves rules that guide how the behavioural rules change over time

<sup>6</sup>SRPs adopt a notion of hypothetico-deductivism as a standard of practice. A stream of hypotheses or models are constructed from a set of core axioms or assumptions (‘hard core’). What sort of hypotheses are constructed and the method used for testing them is determined by an agreed-upon heuristic. After rigorous testing, the results, be they good (passing) or bad (failing), contribute to the general success of the program (‘protective belt’). Social scientists rely on the ability of the core theory to match a myriad of existing and emerging novel facts identified in real data. SRPs are said to be "progressing" if the hard core increasingly produces a stream of testable hypotheses for a broader range of stylized facts (and "degenerating" otherwise). Although some would argue that the SRP model does not provide a completely accurate description of many fields at different moments in history, such as Economics (see Hands (2008), Blackhouse (2004)), they are a reasonable and widely accepted starting point for examining the scientific community.

<sup>7</sup>Agent's private information is incomplete.

<sup>8</sup>Explicitly modelling the construction of models from core axioms is computationally more demanding than modelling statistical processes. As aforementioned, see Gilbert (1997), Edmonds (2007) and Sun and Naveh (2009) for constructive models of scientific knowledge.

<sup>9</sup>Intuitively, innovators use a protocol that is a convex combination of the best SRP's success rate and their private ability. From the discussion on imitator agents, the SRP's success rate, given enough time, is likely to evolve to a number slightly greater than 0.5. We expect private ability to be, on average, 0.5. A convex combination, for  $\lambda < 1$ , to be lower than that of imitators, but still slightly greater than 0.5.

<sup>10</sup>In essence, the way that scientific activity is modelled (as 'random science') is related to the notion of Pólya's urn. Colloquially, for any scientist,  $i$ , and for any stylized fact,  $d$ , an urn containing an unspecified number of white and black balls exists. The fraction of balls in the urn that are white is equal to the researcher's protocol,  $P_{i,d,t}$ . The researcher draws a ball from the urn and either gets a white one ('success') or a black one ('failure'). Based on his draw and the draws of others in his SRP, the proportion of white balls change according to equation 1.

<sup>11</sup>Recall that mutators choose a mixing SRP randomly. To calculate expected scores, they average over perceived protocols for known SRP for each stylized fact:  $(1/N_{i,t}) \sum_{N_{i,t}} \max(p_{j^*,d,t}, p_{j',d,t})$ , where  $N_{i,t}$  is the number of known SRPs at time  $t$  to the researcher.

<sup>12</sup>A small number of agents is chosen to reduce computing time. All simulations are done in MATLAB.

<sup>13</sup>Due to the early presence of innovation, which distorts scientific success.

<sup>14</sup>See Blaug (1992), Landreth and Colander (2002) and Mankiw (2006) for a detailed account of the history of Keynesian macroeconomic theory.

<sup>15</sup>This approach to methodology is thus 'reflexive' (Hands (2004)) in the sense that it uses the approach which it itself suggests is ideal for guiding scientific inquiry.

## Tables:

Table 1: Simulation Results

	New Initial Conditions				
	[1]	[2]	[3]	[4]	[5]
Avg. # SRPs	2.415	4.190	1.555	2.210	2.990
St.Dev.	1.217	1.350	0.647	2.168	2.518
Avg. Success	0.516	0.500	0.775	0.542	0.528
St.Dev.	0.012	0.007	0.018	0.029	0.022
Rule (Imitator)				0.884	0.975
St.Dev.				0.544	0.563
Rule (Innovator)				0.380	0.503
St.Dev.				0.470	0.573
Rule (Mutator)				1.320	1.106
St.Dev.				0.512	0.588
Correlation (End SRP #, Avg. Success)	-0.313	-0.200	-0.173	-0.523	-0.630
[1] Set Types: Imitators					
[2] Set Types: Innovators					
[3] Set Types: Mutators					
[4] Unweighted Sum Score Sets Rule					
[5] Weighted Sum Score Sets Rule					

Table 2: Simulation Results

Unweighted Sum Score Sets Rule			
	End #	End #	End #
Correlation	Imitators	Innovators	Mutators
End SRP #	0.078	0.920	-0.897
End Avg. Success	-0.044	-0.446	0.438
Weighted Sum Score Sets Rule			
	End #	End #	End #
Correlation	Imitators	Innovators	Mutators
End SRP #	0.031	0.826	-0.771
End Avg. Success	-0.174	-0.608	0.694

Table 3: Simulation Results - Robustness

	Weighted Score Sets Rule					
	[1]	[2]	[3]	[4]	[5]	[6]
Avg. # SRPs	2.990	6.215	2.400	3.725	3.090	1.935
St.Dev.	2.518	3.355	1.736	2.769	2.711	1.527
Avg. Success	0.528	0.532	0.535	0.523	0.528	0.575
St.Dev.	0.022	0.035	0.029	0.025	0.030	0.010
Rule (Imitator)	0.975	1.031	0.989	1.057	0.988	0.945
St.Dev.	0.563	0.594	0.570	0.578	0.559	0.546
Rule (Innovator)	0.503	0.536	0.499	0.493	0.526	0.923
St.Dev.	0.573	0.572	0.555	0.530	0.581	0.537
Corr. w/Imitator	-0.037	-0.032	-0.050	-0.019	-0.078	0.081
Rule (Mutator)	1.106	1.104	1.098	1.081	1.189	1.114
St.Dev.	0.588	0.591	0.559	0.579	0.530	0.593
Corr. w/Imitator	-0.208	-0.203	0.018	-0.144	-0.174	-0.029
Corr. w/Innovator	-0.149	-0.141	-0.235	-0.266	-0.245	-0.138
Correlation (# SRPs, Avg. Success)	-0.634	-0.669	-0.515	-0.535	-0.508	-0.357

[1] Benchmark calibration

[2] Low probability of receiving mail ( $\mu=0.1$ )

[3] High probability of receiving mail ( $\mu=0.9$ )

[4] Low impact of experiments ( $\gamma=0.9$ )

[5] High impact of experiments ( $\gamma=0.1$ )

[6] Low impact of innovator ability on scientific success ( $\lambda=1$ )

Table 4: Simulation Results - Augmented Initial Values

	Benchmark	Low Weight on Innovation at Outset
Avg. # SRPs	2.990	1.390
St.Dev.	2.518	0.640
Avg. Success	0.528	0.700
St.Dev.	0.022	0.050
Rule (Imitator)	0.975	0.915
St.Dev.	0.563	0.541
Rule (Innovator)	0.503	0.226
St.Dev.	0.573	0.133
Corr. w/Imitator	-0.037	0.128
Rule (Mutator)	1.106	1.588
St.Dev.	0.588	0.454
Corr. w/Imitator	-0.208	0.132
Corr. w/Innovator	-0.149	0.335
Correlation (# SRPs, Avg. Success)	-0.634	-0.225

# Figures:

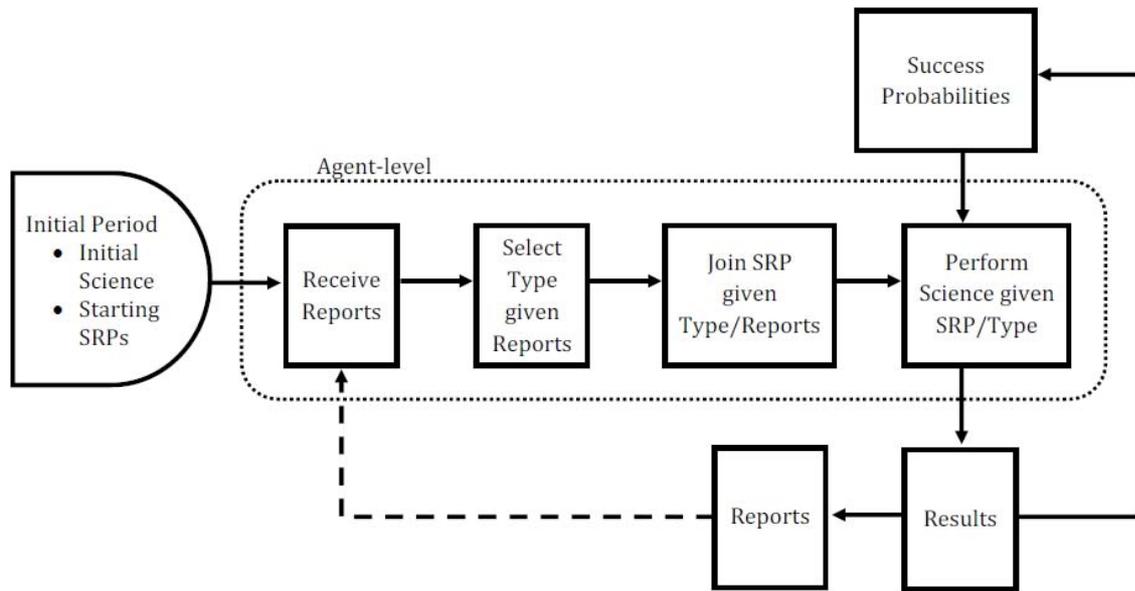


Figure 1: The Structure of the Agent-based Model

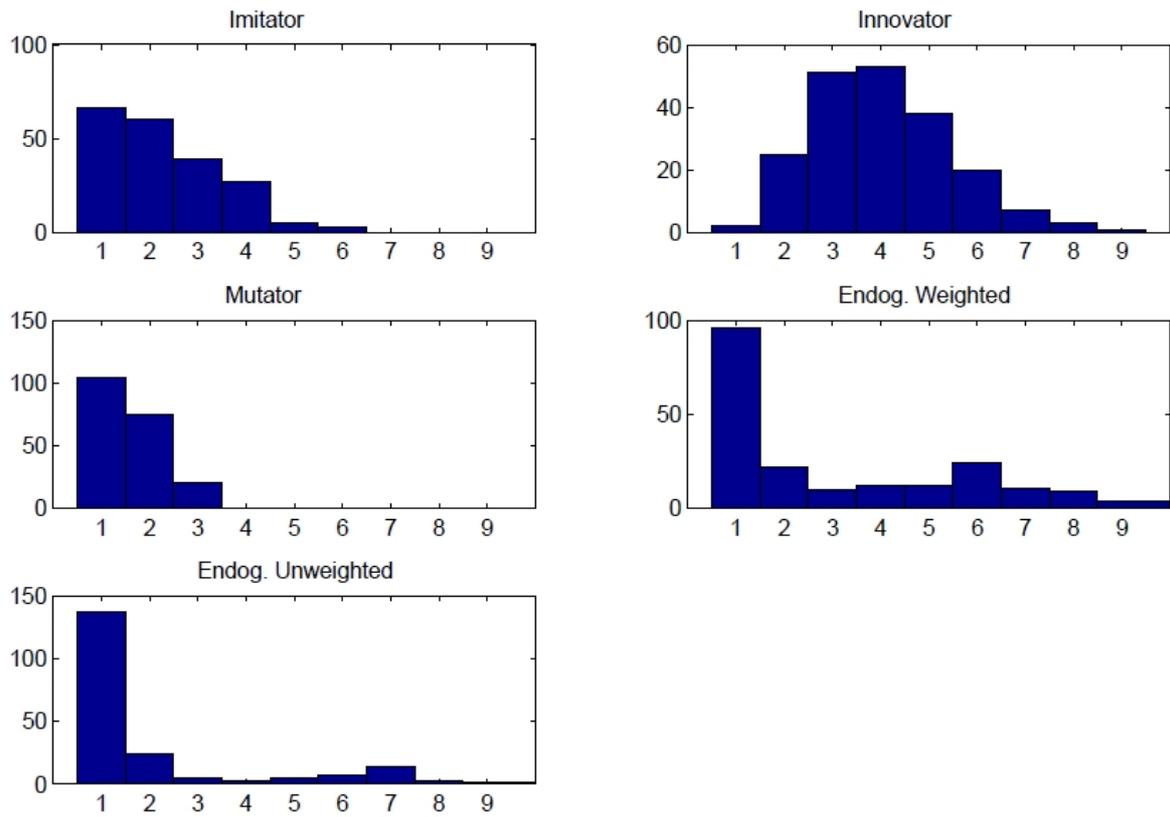


Figure 2: Histogram - Ending # SRPs

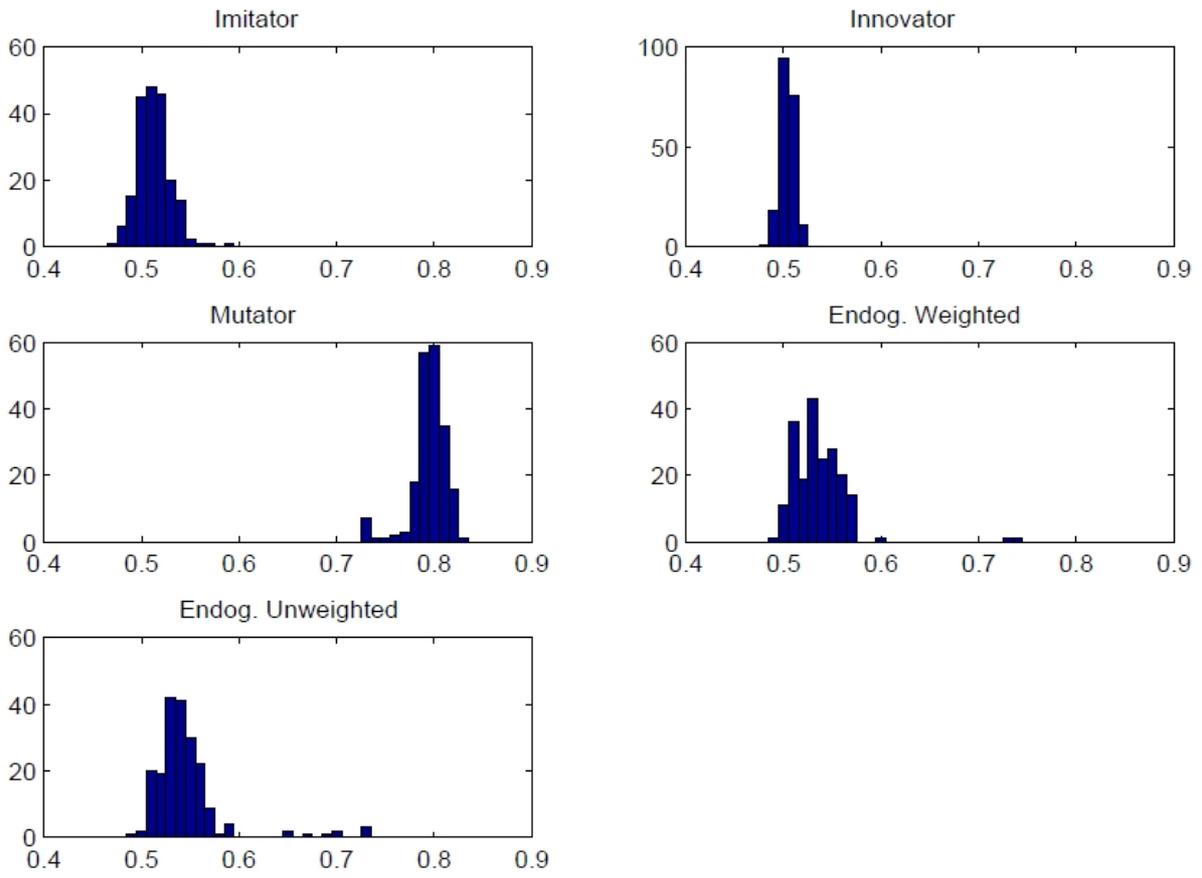


Figure 3: Histogram - Ending Avg. Success Rates

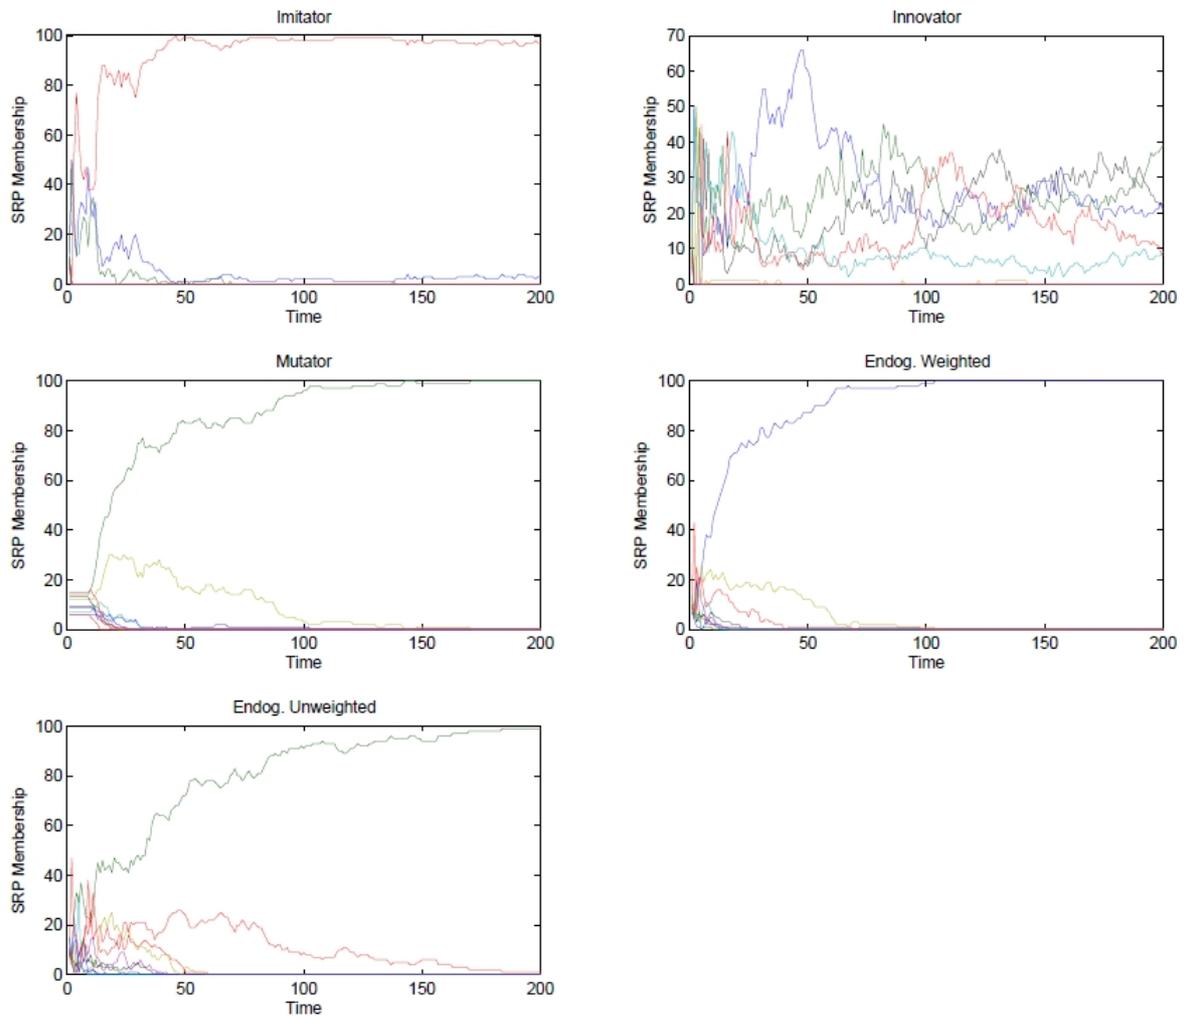


Figure 4: Simulation Samples - SRP Membership

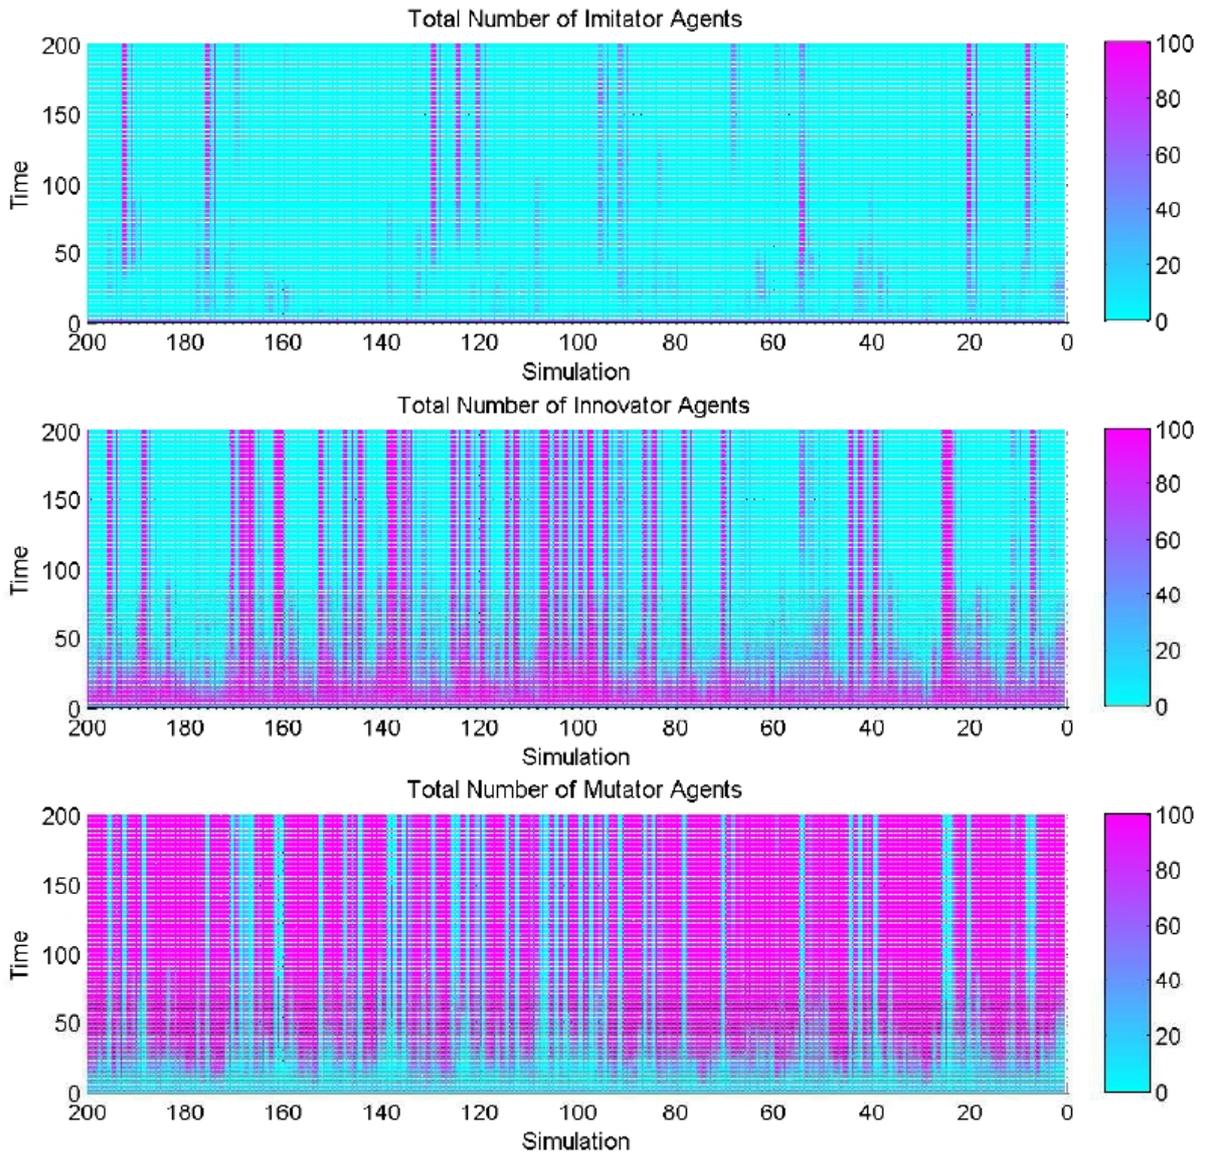


Figure 5: Population of Types - Unweighted Sum Score Sets Rule

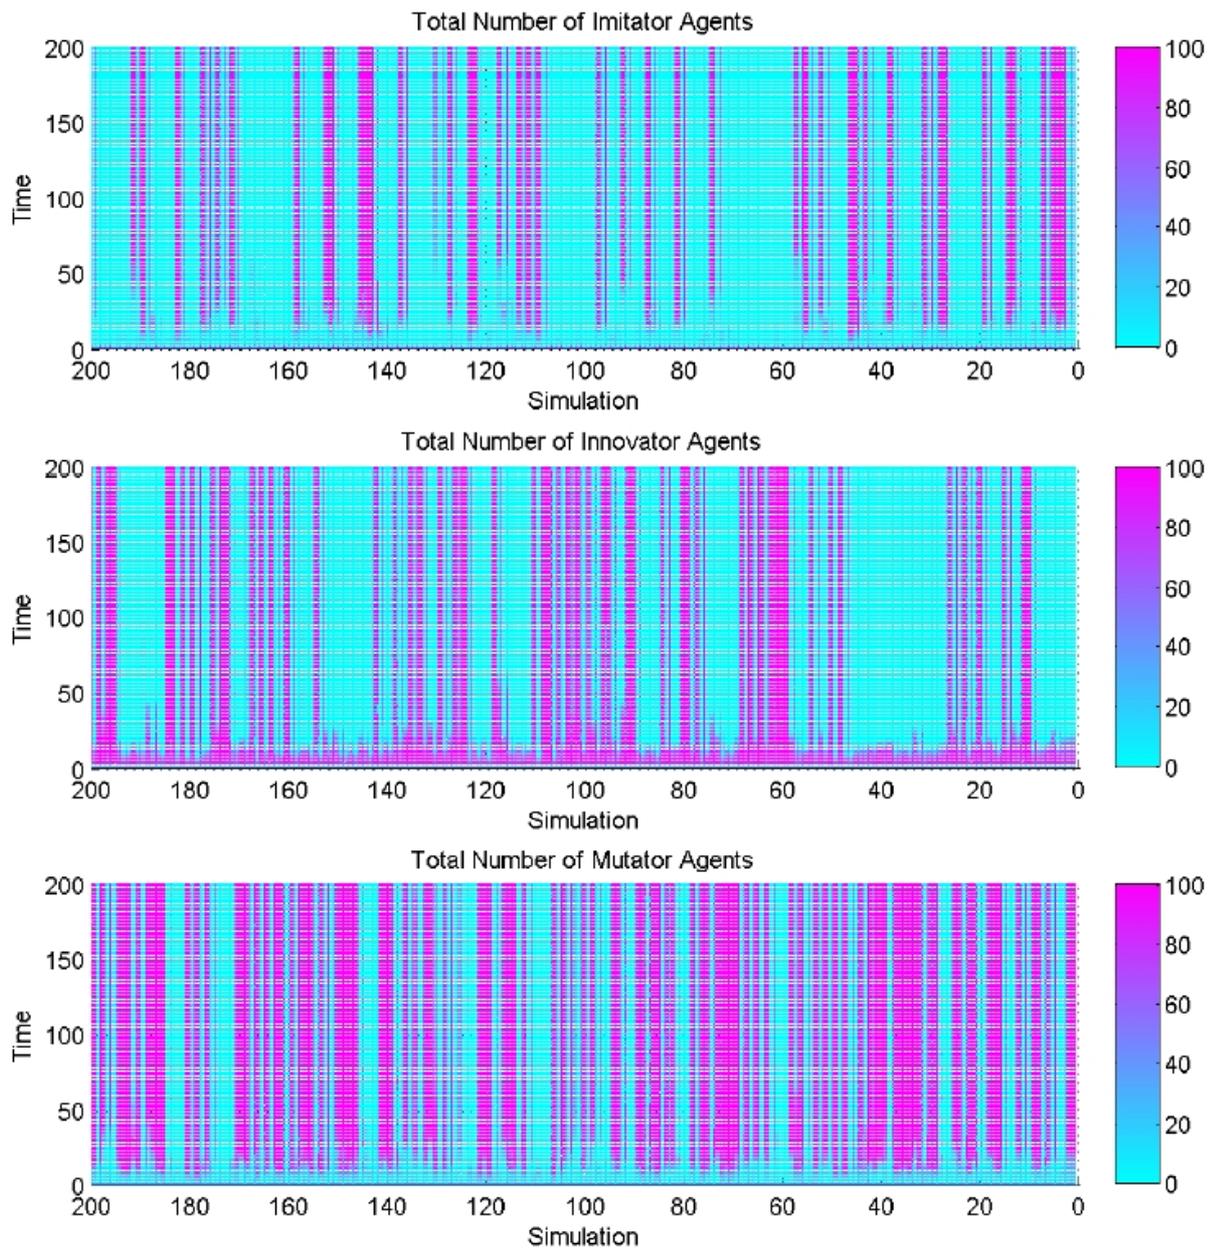


Figure 6: Population of Types - Weighted Sum Score Sets Rule

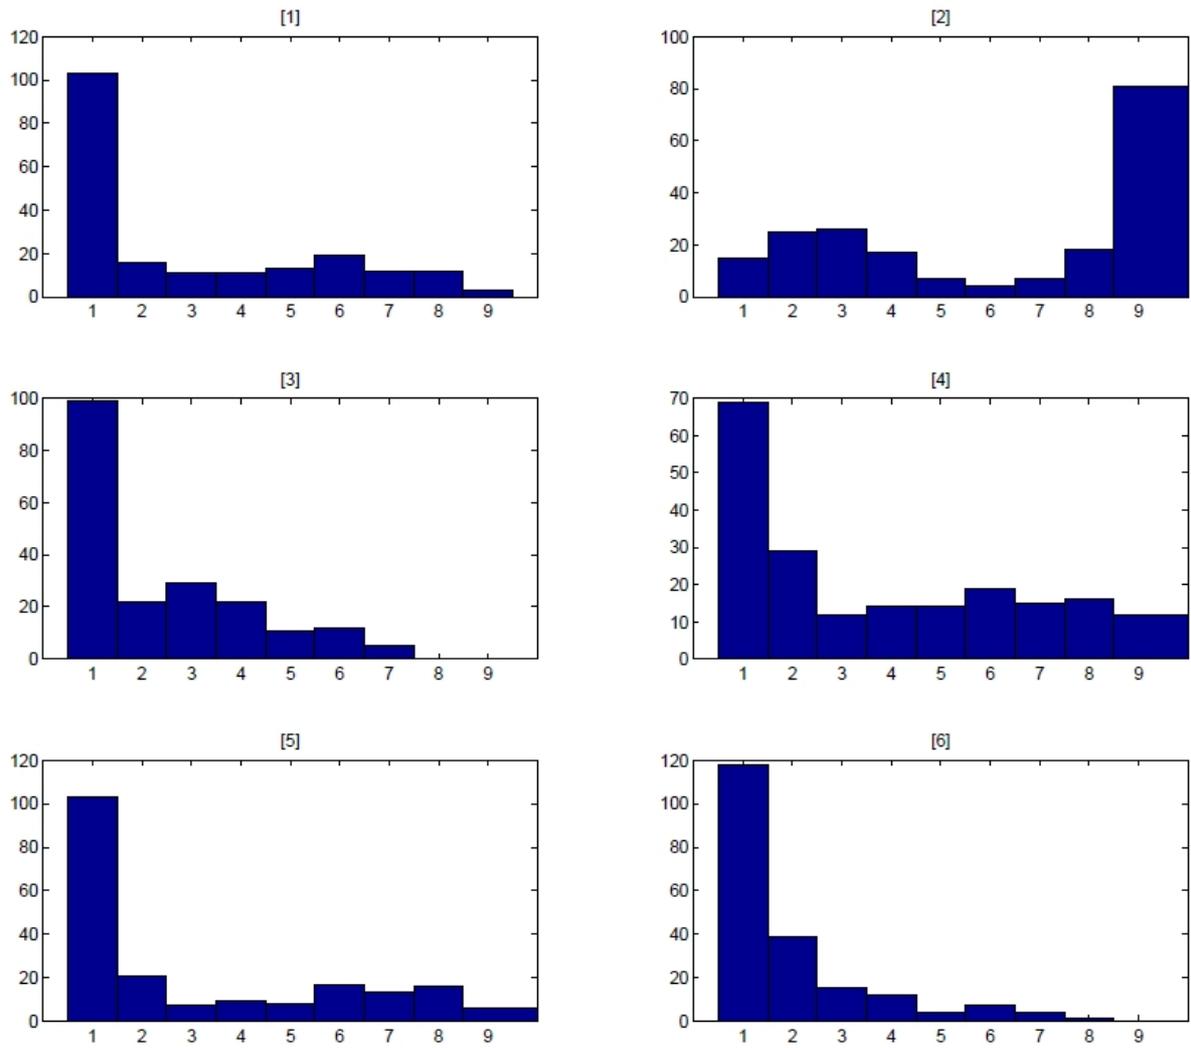


Figure 7: Histogram - Robustness - Ending # SRPs

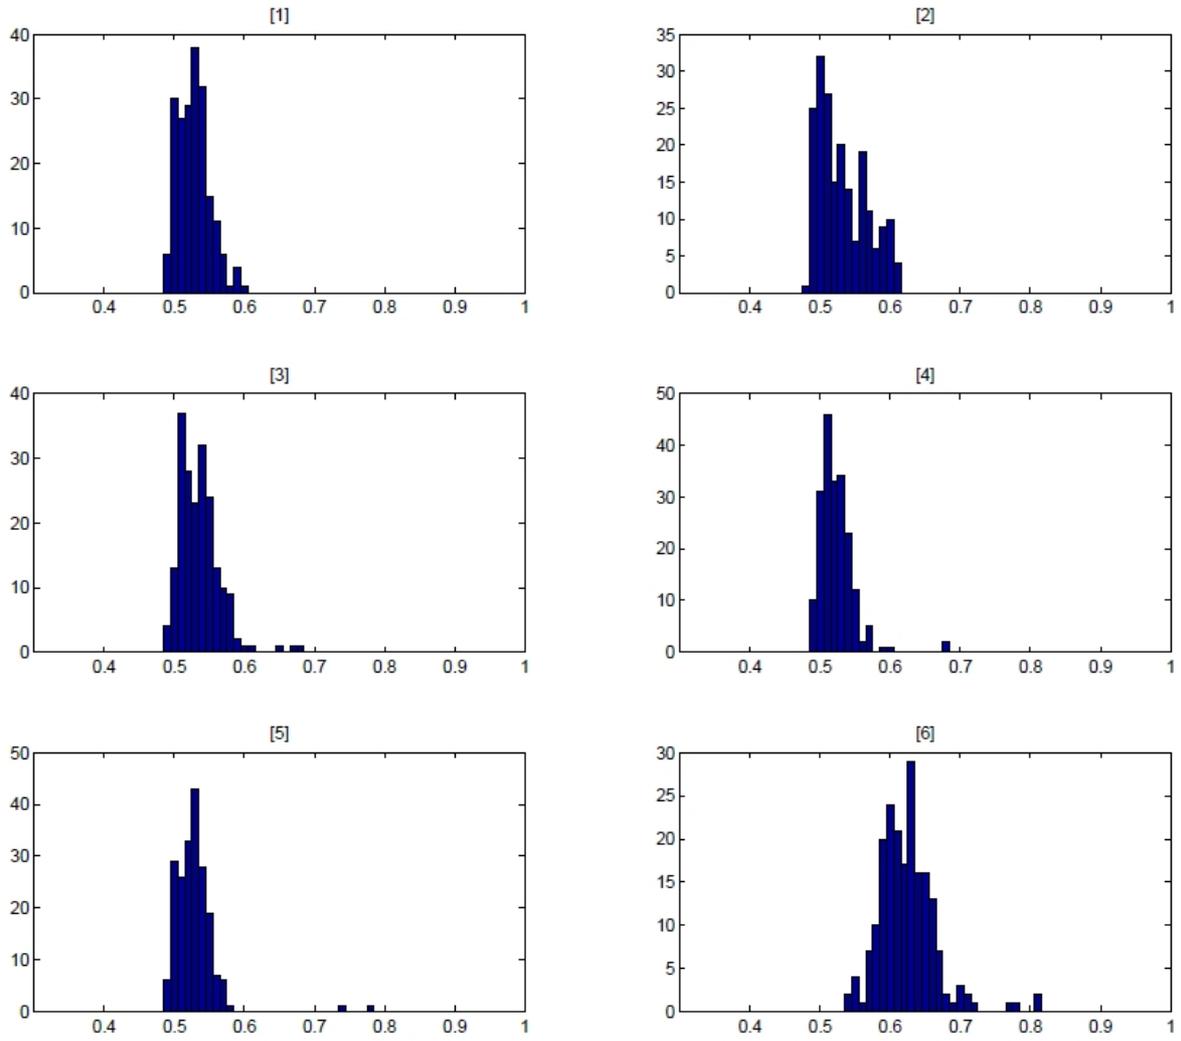


Figure 8: Histogram - Robustness - Ending Avg. Success Rates

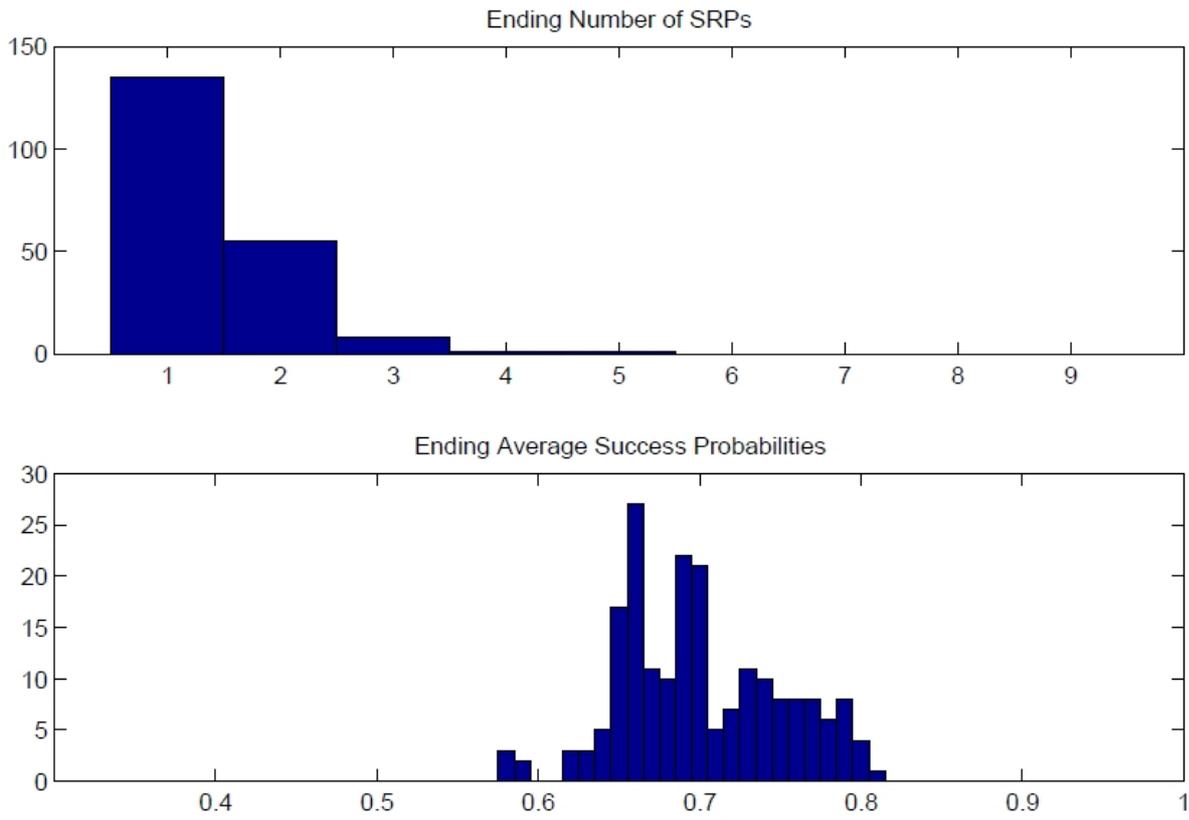


Figure 9: Histogram - Robustness - Ending Avg. Success Rates