# Agent-Based Modelling of Stress and Productivity Performance in the Workplace

by

Matthew Page

A Thesis

presented to

The University of Guelph

In partial fulfilment of requirements

for the degree of

Master of Science

in

Mathematics

Guelph, Ontario, Canada

# ABSTRACT

## AGENT-BASED MODELLING OF STRESS AND PRODUCTIVITY PERFORMANCE IN THE WORKPLACE

Matthew Page
University of Guelph, 2013

Advisor:
Dr. Daniel Ashlock

The ill-effects of stress due to fatigue significantly impact the welfare of individuals and consequently impact overall corporate productivity. This study introduces a simplified model of stress in the workplace using agent-based simulation. This study represents a novel contribution to the field of evolutionary computation. Agents are encoded initially using a String Representation and later expanded to multi-state Binary Decision Automata to choose between work on a base task, special project or rest. Training occurs by agents inaccurately mimicking behaviour of highly productive mentors. Stress is accumulated through working long hours thereby decreasing productivity performance of an agent. Lowest productivity agents are fired or retrained. The String representation for agents demonstrated near average performance attributed to the normally distributed tasks assigned to the string. The BDA representation was found to be highly adaptive, responding robustly to parameter changes. By reducing the number of simplifications for the model, a more accurate representation of the real world can be achieved.

## Acknowledgements

First and foremost I would like to express my deepest appreciation and thanks to my advisor Professor Dan Ashlock for his time, patience, encouragement, motivation, enthusiasm for difficult problems and wealth of knowledge. His guidance provided me the insight into modelling real world problems and tackling the difficult issues contained within this thesis. I could not have had a better mentor and advisor for my graduate studies. Thank you to Professor Ayesha Ali for your kind words on my work and encouraging me to look at many future possibilities for the research. Thank you to my examination committee, Professor Rajesh Pereira and Professor Gary Umphrey for the great feedback and comments on the thesis and the wonderful discussion following my defence. Thank you to all the members of the Ashlock lab who supported me through my studies and encouraged me when things were tough. Furthermore I would like to thank my family and friends for their unending support throughout my studies. To my dad, Chris, thank you for instilling in me the desire to tackle all things analytical and for teaching me about business. It was this knowledge that made this research possible and accessible to the real world. To my mom, Shirley, thank you for showing me there is so much to life contained within the creative elements of our being. Finally I would like to thank my wife, Beth, for her love, support, encouragement, and understanding in the most trying of times over the years. It is her steadfast dedication to serving the vulnerable populations of society that inspires me to take on the hardest of problems and find solutions to the impossible. Without her encouragement and support I would not be the person I am today.

# Contents

# List of Figures

# List of Tables

# Chapter 1

# Introduction

The impact of stress on an individual and the associated ill-effects are ubiquitous. The negative effects of stress on an employee have a substantial impact on productivity loss. The Canadian Government recently released its first national mental health strategy to address the growing prevalence of this issue in Canadian society [58]. This strategy acknowledged that the impact of mental health illnesses cannot be addressed through treatment alone and must also be mitigated through healthy lifestyles including through work-life balance. This study is the first of its kind to examine the effects of fatigue and overwork as potential contributing factors to the overall stress of individuals within an organization using evolutionary agent-based techniques framed as a game that models employee behaviour. Furthermore, it attempts to establish a framework and the necessary technology for future research into the predictive modelling of stress. The model is built on the results of qualitative and quantitative social-science and health research to encode maximum hours worked, differentiated response, and the impact of the compounding effects of stress on individuals.

The productivity loss due to stress related ill-being translates into a sizeable loss of potential revenue for corporations [33, 55]. The well-being of an individual is seen as the fulfilment of five distinct factors: career well-being, social well-being, financial well-being, physical well-being and community well-being [73]. These five core aspects of well-being are reviewed in depth in a series of studies conducted by Dr. Jim Harter from Gallup's workplace management practice [73]. Career well-being is seen as how much an individual enjoys their time spent at work each day, feeling appreciated as an employee and feeling a sense of pride and appreciation when speaking about one's employer [73]. Social well-being is considered as having strong relationships and love in one's life, including having a support system to rely on when encountering problems in life [73]. Financial well-being is perceived as the effective management of one's economic life prudently, being aware of costs and in control of expenditures [73]. Physical well-being is described as having good health and additionally having enough energy to get tasks done on a daily basis [73]. Finally, community well-being is the sense of engagement one has with the area in which they live [73]. In the United States amongst individuals with the lowest well-being, the costs associated with productivity loss were found to be $28,800 USD annually per employee [62]. When investigating the economic impact of fatigue alone the costs to corporations are still significant. Rosekind et al. tested 4188 employees at four US corporations and found that fatigue-related productivity losses were approximately $1967 per employee each year [55]. While the exact financial impact of stress and fatigue on individuals varied depending on the definitions and parameters used in the study, the research consistently found a measurable negative correlation between stress and fatigue and the associated economic cost to organizations. These studies demonstrate the economic significance of this research area.

The majority of research on the impact of stress on an individual over the last few decades has been conducted from a medical or psychological perspective. The ill effects of stress range from obesity [10,37], to sleeplessness [1], and heart disease [44,64]. Stress is a significant contributor to employee ill-being and has many adverse effects on an individual [70,71]. In 2009, Marcora et al. established a link between mental fatigue and impairment of physical performance in humans [71]. A statistically significant correlation between the psycho-biological state of mental fatigue and the negative impact on physical and cognitive efforts was determined when testing people under controlled conditions [71]. Lupien et al. conducted research on the effects of stress on behaviour and cognition [70]. Findings included that exposure to chronic stress negatively impacted the brain's development capabilities. Fevere et al. allege that approximately 27 per cent of adults in Europe experience at least one form of mental ill-health in any given year [50]. Moreover, depression is on track to become the highest ranking cause of disease in the developed world by 2020 [59]. The findings of these researchers contributed to the structure of the model tested in this study, including the incorporation of diminishing agent productivity as stress levels increase and fatigue sets in.

Stress is defined as the physiological or psychological response of the body to any given stressor, real or perceived [69]. Stressors are deemed to be any external, chemical, biological or environmental demand placed on an organism that prompts a response. Depending on the nature of a given stressor, an individual may have a positive or negative response [45]. Minimizing the negative stress (distress) and maximizing the positive stress (eustress) on an individual can yield optimal work performance from an individual [80]. The impact of stress on work performance has been studied by psychologists Robert Yerkes and John Dodson. Their research led to the creation of the Yerkes-Dodson law which is an empirical relationship between stress and performance in an attempt to obtain optimal performance from an individual [80]. This law states that individual performance increases with physiological or psychological eustress to a unique maximum for each individual. If the maximum is exceeded, performance begins to decrease. This law is often depicted as a normally distributed curve on a graph and is commonly referred to as the pressure performance curve as seen in Figure 1.1.

The well-being of individuals at work is increasingly being recognized by governments and employers as important to the overall profitability of organizations and gross domestic product of nations [34]. Consequently, both governments and employers are investigating and implementing strategies to mitigate the ill-effects of work, including

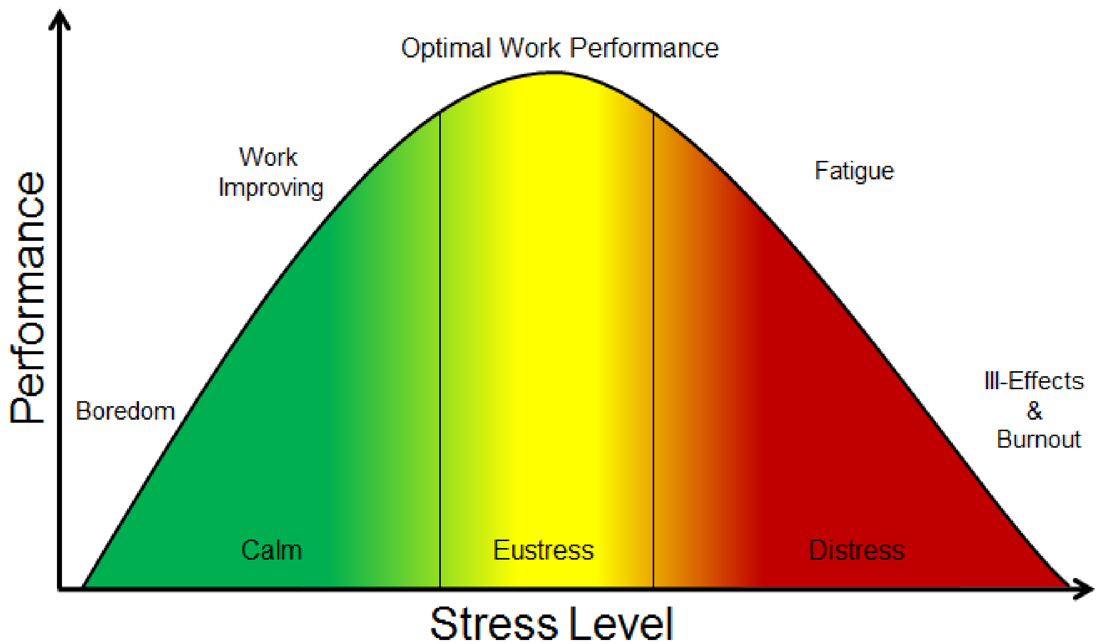


Figure 1.1: The Yerkes-Dodson pressure performance curve that establishes and outlines optimal work performance conditions.

stress, on individual employees [58]. Organizations are developing business cases that quantify the economic benefit of an employee's well-being, accounting for the loss of productivity and revenue due to the ill-being of employees [62].

There has been significant qualitative research over the last few decades in an attempt to optimize employee well-being [67] and demonstrate the importance of work-life balance [38, 52, 74]. Barsh, Cranston and Lewis' qualitative research on centred leadership in the workplace found a correlation of employee fatigue to productivity; establishing that employees with more rest, namely the relief of stress, attained higher productivity and were more effective at problem solving [12]. In the 1980's, A.F. Sanders attempted to establish a model of stress and human performance. Sanders concluded that "performance measures are not proper indicants for stress" and "performance measures are only needed as a control measure, ascertaining that sufficient effort is allocated to keep performance at an optimum" [67]. These results led to the embedding of the notion of minimum productivity required to not get fired in the model tested.

In designing the model some simplifications were necessary to establish a preliminary framework to mathematically model stress using computational intelligence. The following is the overall layout of this study: Chapter 2 provides the necessary background and context for the research. Chapter 3 is a presentation of the preliminary trinary String representation agents, the first attempt at modelling stress in the workplace. Chapter 4 contains a revised agent encoding, replacing the String representation with Binary Decision Automata to improve the adaptability of agent behaviour. Chapter 5 presents the enhanced Binary Decision Automata model which further reduces the number of simplifications in the model. Chapter 6 provides further mathematical review of the near average performance of the String representation. Chapter 7 outlines the conclusions and possibilities for future research.

# Chapter 2

# Evolutionary Computation

The mathematical and computational concepts that were utilized throughout this work are outlined and analyzed below. First, evolutionary algorithms, a subset of evolutionary computation, are explained followed by an outline of the notions of evolutionary fitness and agent representation with relevance to the social aspect of modelling stress and a description of the encodings used for this model. Evolutionary algorithms were selected to conduct the agent-based modelling for this research because the population sizes used for realistic tests are small enough that a plausible global behaviour is unlikely to be established. Other forms of mathematics such as differential equations may be useful to investigate global behaviour for large population sizes such as in finance for risk analysis [14] and epidemiology for disease transmission [40].

## 2.1 Evolutionary Computation History

Evolutionary algorithms are a subset of the field of evolutionary computation. There have been many founders of different components of Evolutionary Computation, however no single person has been credited with the discovery of this field. The origins of the discipline began in the late 1950's and early 1960's with G.E.P. Box who looked at increasing industrial productivity with evolutionary operations [15]; R.M. Friedberg researched learning machines for IBM [35, 36]; and H.J. Bremermann investigated optimization through evolution [16]. Following this foundational research, the field remained predominately unknown for several decades until computational processing capabilities were increased. In the 1970's, three distinct approaches of evolutionary computation were established and co-evolved together shaping the field today. These approaches are *genetic algorithms*, *evolutionary programming*, and *evolution strategies*.

Genetic algorithms were initially invented by John Holland [41] as a means of mimicking Darwinian evolution to solve optimization problems. These algorithms are established as heuristics which in turn are used to speed up the iterative process used to attain a solution for a given problem. Evolutionary programming was discovered by Lawrence Fogel [31] as a means of creating artificial intelligence. Fogel's intent was to take finite state machines (FSM) to encode behaviour into Boolean predictions driven by input and output states. Evolution strategies were developed by Ingo Rechenberg [43] and Hans-Paul Schwefel [68] and were designed to solve discrete and continuous parameter optimization problems. The aforementioned authors have all made significant contributions to the field of evolutionary computation and enabled the design of modern day evolutionary algorithms [3]. Evolutionary algorithms are utilized today in a broad spectrum of fields ranging from cellular automata in bioinformatics [25,30,76], structural design in engineering [24, 28, 81], modelling game theory in economics [2, 13, 17], molecular design in chemistry [46, 61], algorithm optimization in computer science [39, 54], optimizing manufacturing processes [24, 60], multi-objective optimization problems in mathematics [32, 53], and quantum mechanics and computing in physics [57, 66, 72].

## 2.2 Evolutionary Algorithms

Evolutionary algorithms are meta-heuristic optimization algorithms that use the concepts of natural selection and biological evolution to arrive at possible solutions for optimization problems [11, 26, 29]. A candidate solution is a member of the set of all possible solutions to a given problem. These candidate solutions act as agents (individuals) in the population. The overall structure of evolutionary algorithms remains relatively similar throughout different applications. In terms of pseudo code, the general approach is:

```
Initialize a random population of individuals
Establish and evaluate the fitness of each individual
  REPEAT :
        Select the most-fit individuals for reproduction              (SELECTION)
        Generate new offspring through mutation and crossover         (MUTATION)
        Evaluate the fitness of offspring                             (EVALUATION)
        Replace least-fit population with offspring                   (REPLACEMENT)
  UNTIL EXIT CRITERIA
```

Figure 2.1: The standard structure of an evolutionary algorithm.

The above generalized approach is the same that is applied to the model used in this research. Two fundamental concepts typically outline the overall performance of evolutionary algorithms: representation and the fitness function [3]. The components of genetic algorithms are described below.

### 2.2.1 Population

The **population** in an evolutionary algorithm is the collection of possible solutions for a specific problem [11]. Each solution is represented by an individual in the population. The role of the population is to provide individuals for the mating and creation of offspring to allow a computational search to take place [3]. While population size can be varied for all problems, choosing a suitable size enables a more appropriate representation of solutions. This is due to the fact that 'small' population sizes can potentially lead to poor solutions [9, 51, 78] and 'large' populations require more computational effort and consequently time and may not produce a more optimal result [48, 78]. Therefore there is a trade-off between time required to reach the solution and computational effort. The goal is then to establish an algorithm that runs quickly and provides a sufficiently accurate result for the outlined problem.

The population used in this model represents the collection of individual employees. Consequently the population can be viewed as the staff of the company for which an employee works. Thus metrics such as overall productivity and stress for the organization are ideally to be maximized and minimized, respectively. It should be noted that only overall productivity is explicitly coded into the model to be maximized. In order to accomplish this, the collection of candidate solutions (individuals) needs to be evolved to locate good solutions. While stress for individuals should ideally be minimized, it is possible that a candidate solution may contain high stress. Consequently, a good solution for the stress problem outlined in this model is an individual that has optimally low stress and high productivity.

### 2.2.2 Individual

The **individuals** in a population are the unique data structures that portray elements of the given problem's solution space [75]. Biologically speaking, individuals are the elements within the algorithm that contain a chromosome or gene that represent one possible solution. Typically a gene is established similar to the string structure that DNA is encoded with, however instead of using CGAT as representative characters in each loci (position), a binary representation is used. This means that for each position in the string, a value of 0 or 1 could exist. The values that can exist at any given loci will be discussed later as these are problem specific and refer more to the representation. In the case of the research developed in this thesis, values of 0, 1, and 2 are possible. The data structure of the individual is not limited to a string and can take on more complex forms such as finite state machines or artificial neural networks.

The data structure that the individual represents varies depending on the problem.

For any given individual in the population, the data structure will also outline a specific solution for the problem. The data structure that encodes an individual's behaviour is referred to as the representation for evolutionary algorithms.

### 2.2.3 Representation

A significant component of an evolutionary algorithm is the **representation** used [3,63] as depending on the nature of the problem, the outcome can be determined by the chosen representation [5]. There are many different types of representation such as arrays, strings, tree structures, neural nets and finite state machines. A representation is used to encode a specific problem, mapping the real world to one that can be calculated within a computer program. Additionally, how the representation is designed and encoded has a dramatic outcome on the ability of the algorithm to effectively establish candidate solutions and ultimately its ability to search the solution space [22].

Two different representations were explored in the development of this model. As the application of evolutionary algorithms to the modelling of stress is new, a simple representation was selected for the initial modelling followed by a more enhanced model. A trinary String representation was selected due to its simple structure and ease of incorporation to establish the framework for an employee. The use of Binary Decision Automata (BDA) was then employed to expand on the cognitive and adaptive capabilities of the modelled employees. This was done through establishing multiple states for each BDA and mutating the states to create variations in the structure. Examples and further information regarding the encoding of the string representation and the BDA are provided in Section 2.4 and 2.5 respectively.

### 2.2.4 Mutation & Variation Operators

Stochasticity is introduced into the model by the inclusion of *mutation* through *variation operators* [29]. The mutation acts on the individual's data structure to alter a randomly selected component to allow for genetic diversity. Ultimately, it is variation that allows for the implementation of Darwinian evolution into the algorithm. The mutation acts on the individuals within the population after reproduction has taken place. When two selected individuals from the population mate, offspring are created representing a new individual in the population. This new individual will have adopted traits from both parents but how these traits are selected is dependent on the type of crossover that takes place. Crossover is the method by which offspring inherit the genes from their parents and consists of the exchange of structures between individuals [3, 30]. This exchange of information can be seen as a binary variation to the data structure itself [3]. In this model inaccurate replication of a mentor's behaviour pattern is used instead of crossover. In this study, crossover was not used as the training of individuals was meant to mimic one individual learning from another. This is carried out through the inaccurate mimicking of the behavioural patterns (by copying the data structure) of successful individuals in the organization (population). A mentor is selected to train an under-performing individual at which time the trainee will adopt the mentor's actions and work patterns. It then applies mutation to account for the inaccuracies of mimicking one's behaviour.

On the creation of the offspring, probabilistic mutation is applied to allow for diversity in the new gene's structure. The mutation represents a unary change to an individual structure as it only affects a single locus in the data structure's gene. In this model, for both the string representation and the BDA's, the value $\mu$ denotes the probability that at a single locus in the genetic structure of the agent will mutate. As with representation, there are many types of mutation and variation operators that ultimately affect the outcome of the algorithm and its ability to arrive at a solution.

### 2.2.5 Fitness Function

The measure by which a population is evaluated in an evolutionary algorithm is known as the *fitness function*. A fitness function tests an individual in the population and assigns a quality measure to that individual and its genotype [3]. The fitness function is established from the requirements outlined in the problem by which an individual and population must adapt. It provides and mathematically defines the notion of improvement for individuals whether this be maximization or minimization.

The collection of benchmarked candidate solutions obtained from the range of the fitness function is known as the *fitness landscape*. The fitness landscape is the graph, over the space of all the genes, of the fitness value. The population picks out a subset of the fitness landscape and allows for a visualization of the relationship that exists between a candidate solution and its position amongst all possible solutions [3,30]. Much like the graph of a function, a fitness landscape will have peaks and valleys. If many peaks and valleys occur for a problem, the smaller values (in relation to each other) are known as local optima for the problem and the larger values are known as global optima [3,30]. Consequently, if high fitness is required for the problem (maximization) then the smaller peaks are local optima and highest peaks are global optima. If low fitness is required, then the lowest points would be the global optima and values higher would be local optima. The measure of fitness for any problem will depend on the representation chosen for the individuals.

The fitness function in this model is the test of productivity for a given corporation. A certain minimum level of productivity is set in the model which must be achieved by an employee so as to maintain employment. This establishes a measurable test by which the employees can be compared. Successful individuals will have high productivity and unsuccessful individuals will have low productivity. Consequently, the test for the fitness function is to maximize the productivity of individuals and establish which agents are successful. Ultimately the goal of any fitness function is not only to establish an optimal solution but also to calculate these solutions as quickly as possible since evolutionary algorithms must be run for numerous iterations to generate any non-trivial results.

### 2.2.6 Exit Criteria & Termination Condition

Most problems will have many local optima and potentially global optima. This is due to an evolutionary algorithm enumerating all possible solutions within the time set for the problem. This time period might be a certain number of iterations of the algorithm or could be when a certain initialized value is reached for fitness. These conditions are known as the *termination condition* or *exit criteria* for the algorithm [3, 11, 26]. As evolutionary algorithms are stochastic, there is no guarantee that the candidate solutions generated are in fact an optimum for the problem. This can cause a problem as the algorithm may never stop if the condition is set to obtain a certain value. Thus, there are several commonly used options for termination conditions and exit criteria [11, 26]:

1. There is a maximum allowable time for CPU processing.

2. There is a limit to the total number of fitness evaluations conducted.

3. In a specific period of time (i.e. a certain number of generations or fitness evaluations), the improvement to the fitness is lower than a set value.

4. The diversity of the population is insufficient compared to a set value.

This model consisted of employing a specific number of generations as the exit criteria for the algorithm. This method was applied to both the String representation and the BDA, and was thoroughly tested in parameter studies with 10 generations for short term simulations and 100 generations for long term simulations [6, 7].

### 2.2.7 Generation

A *generation* is the population of individual data structures that lasts for each iteration of the evolutionary algorithm [26, 75]. The selection, replacement, and mutation taking place on the current generation, the new generation of individuals established are what exist for the next step in the iteration of the algorithm. The number of generations necessary for a problem is typically established by the number of iterations the algorithm is set to run. In this study, the number of generations (iterations) covered 10 years (120 iterations) and 100 years (1200 iterations).

### 2.2.8 Selection

The role of *selection* is to establish a hierarchy and rank order among individuals to allow an optimal pair of parents to be selected for mating [26]. A selected individual is designated a *parent* if it undergoes variation through crossover in order to generate offspring, new individuals. It is the role of selection that allows for improvement among individuals. In the context of this study, after the generation of the population and the establishment of an initial benchmark of the fitness of individuals, the population is sorted in rank order based on productivity, from highest to lowest. The mentors (parents) are then selected from the training tiers denoted by upper and lower bounds that establish a per cent range. The top performing individuals are mentors and lowest performing individuals are trainees.

The values selected for the bounds were decimal values chosen from the set $\{0, 1\}$ with 0 being the highest and 1 being the lowest. For example, if the upper bound was 0 and the lower bound was 0.1, this would denote the top 10 per cent of individuals in the population. Once the per cent from which mentors can be chosen is established, the selection of mentors is probabilistic as there may not be a one-to-one relationship between the number of mentors and the number of individuals being retrained. The values of selection were varied in this study to establish the performance of the model under different circumstances [6, 7].

### 2.2.9 Replacement

The role of *replacement* is similar to that of selection; however, it denotes the individuals that are to be replaced instead of who should mate. The individuals chosen for replacement are probabilistically selected from the lowest performing individuals in a population. For this model, replacement denotes the individuals that require retraining due to poor performance, a consequence of too much stress. The tier from which individuals are selected is denoted by a real value in the range $[0, 1]$. This range establishes the replacement fraction from which individuals are selected. For example, if the replacement fraction is set at 0.1, the lowest 10 percent of employees would be selected for retraining. If the value were set at 1, all employees would be retrained.

The training used is similar for both the String representation and the BDA. For String agents, a simple training model of imitation is used. This is implemented using probabilistic mutation. Each trainee inaccurately copies the time assignments of the mentor that is assigned to them. The parameter $\mu$ mentioned previously then denotes the probability that an agent will adopt their mentor's value at a given loci in the gene.

## 2.3 Model Simplifications

In this model an agent's performance is evaluated on a weekly basis and is comprised of the agent's base task and special project work performance. These weekly totals are summed to attain a monthly figure that enables the adaptive algorithm to update the agent population. The stress response curve is comprised of the parameters $\alpha$ and $\beta$. The rate at which stress impacts an agent's performance is denoted by the $\alpha$ parameter with smaller values representing greater stress tolerance. The alpha parameters in this study are selected for agents from a normally distributed range with the minima and maxima set as fixed parameters in the model. The $\beta$ parameter denotes the floor of the curve and represents the maximum possible work a maximally stressed agent will achieve. When the stress response curve is established, the maximum value of 1.0 represents no stress and values near $\beta$ represent the maximum stress for an agent. The weekly figures are created through the following steps:

1. The number of individual actions the agent undertakes that are not of type 0 are tallied. A standard work day is regularized to be 8 hours. If more than 8 hours are worked by the agent then the daily stress score is $(hours - 8)$

2. The total stress number, $t$, is a number in the range 1-1024 and is calculated by multiplying the daily stress scores for the 5 days of the work week.

3. The *stress factor* for the week is computed from $t$ as

$$S(t) = \beta + (1 - \beta)\frac{1}{1 + (\alpha t)^2} \tag{2.1}$$

4. The performance on base work for an agent is calculated as the number of hours spent on their base job in a week multiplied by the stress factor. This is represented for an agent $A$ by $\mathcal{N}_A$ for an agent.

5. The performance on the special project for an agent is calculated as the number of hours spent on their special project in a week multiplied by the stress factor. This is represented for an agent as $A$ by $\mathcal{S}_A$.

The total performance of an agent is measured on a monthly basis. If the agent's productivity on their base task falls below an acceptable minimum performance level, M, the agent is fired and a new agent is introduced into the population. The new agent is introduced by cloning an existing random agent in the population and mutating any of the states in the BDA of the selected agent. The agents that are not fired are given a performance evaluation $P_A$, governed by:

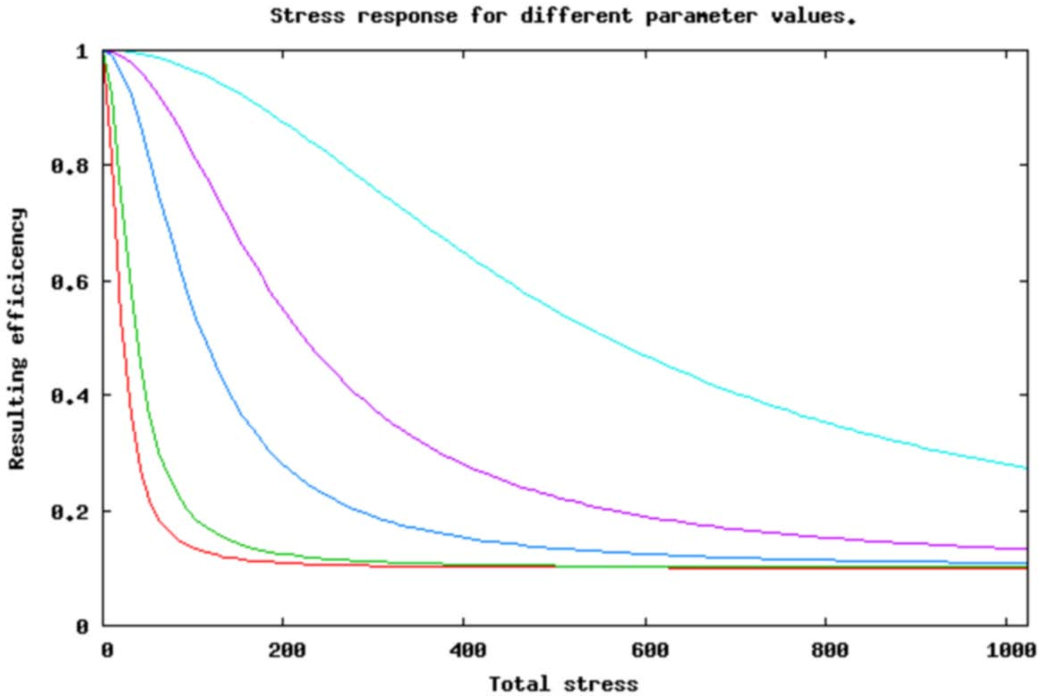


Figure 2.2: Response of the stress factor curve to the parameter $\alpha$. Shown are curves for $\alpha = 0.05$, 0.03, 0.01, 0.005, and 0.002 with the stress floor parameter set to $\beta$=0.1. The smaller the alpha value, the slower the curve descends from 1 to $\beta$.

$$P_A = \delta \mathcal{N}_A + (1 - \delta)\mathcal{S}_A \tag{2.2}$$

The parameter $\delta$ in the aforementioned equation is the representation of management's perception of the relative worth of special projects. This perception is denoted as the coefficient $0 \leq \delta \leq 1$ for any given month. Consequently, this parameter introduces stochasticity into the model by allowing the agents to predict the importance of the special projects that they choose to undertake.

## 2.4 String Representation

The initial agent representation selected was in a basic chromosomal form that made up a simple String representation [6]. This can be seen in the results found in Chapter 3. This type algorithm is the simplest as members of the population are stored as fixed length strings comprised of a predefined alphabet. In the case of this study, the string alphabet was made up of values in the set $\{0, 1, 2\}$.

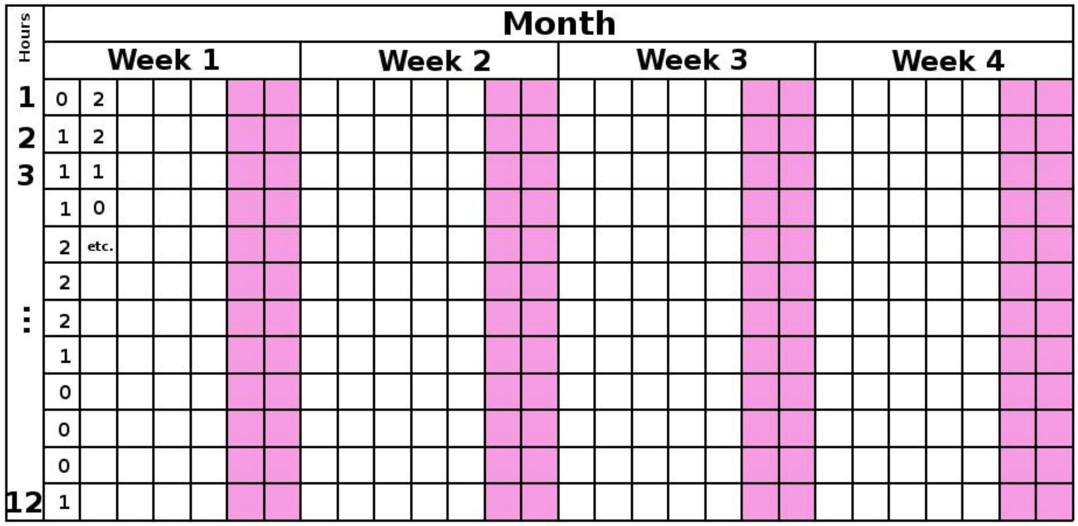


Figure 2.3: An example of the way time is structured in the chromosome. This describes an agent's time allocation to rest, base task, and special projects. The string describing this agent would start 0111222100012210... Coloured days are weekends.

The string that defines an individual is comprised of 240 characters from this set. Each of the loci in the string represents an hour of time. This hour can be spent working on one of three tasks (each of the elements of the set). The designation 0 represents rest which consists of all non-stressful activities for the agent. These could include taking a break, sleeping, going on a walk or engaging in low-stress conversation with co-workers. The designation 1 represents time spent working on the base task assigned to the agent. The base task is the work that the individual was hired to do and would typically appear in one's job description. The designation 2 represents time that is spent on special projects. A special project can be deemed any work outside of the agent's base job which can advance the company and gain status (reputation) for the individual.

## 2.5 Binary Decision Automata

In order to increase the realism of the model, the initial String representation needed to be improved to provide enhanced cognitive capabilities to the agents and furthermore to allow for basic decision making. Accordingly, Binary Decision Automata were used. Binary Decision Automata are augmented finite state machines with the underlying transitions driven by Boolean tests. There are three input variables that the BDA considers: the fraction of time the agent has spent on its base task, working on special projects, and the stress multiplier. The stress multiplier is represented by a real value between 0 and 1. In the model, a value of 1 represents non stressed and 0 represents a maximally stressed agent. In each state of the BDA, one of the input values is compared in a Boolean test using one of the operators (seen in Table 2.1 in Figure 2.4) to a constant. If the test result is true, then one possible action occurs and the next state results otherwise a different second action occurs and next state results. In this study the 'near' operator is set to have a default value of $\epsilon = 0.05$ for comparison.

| State | Test | If T | If F |
|---|---|---|---|
| 0 | if(Strss$<$0.874) | Rest$\rightarrow$ 2 | Base$\rightarrow$ 4 |
| 1 | if(Task2$\leq$0.779) | Rest$\rightarrow$ 7 | Base$\rightarrow$ 2 |
| 2 | if(Task1$\leq$0.289) | Spec$\rightarrow$ 1 | Spec$\rightarrow$ 1 |
| 3 | if(Task2$>$0.288) | Base$\rightarrow$ 8 | Rest$\rightarrow$10 |
| 4 | if(Task2$>$0.124) | Base$\rightarrow$ 8 | Rest$\rightarrow$ 6 |
| 5 | if(Strss$>$0.784) | Spec$\rightarrow$10 | Rest$\rightarrow$ 3 |
| 6 | if(Task1$\leq$0.439) | Spec$\rightarrow$ 9 | Rest$\rightarrow$ 9 |
| 7 | if(Strss$\approx$0.088) | Rest$\rightarrow$ 0 | Spec$\rightarrow$ 9 |
| 8 | if(Task1$\geq$0.792) | Spec$\rightarrow$11 | Base$\rightarrow$ 9 |
| 9 | if(Task2$\leq$0.217) | Spec$\rightarrow$ 3 | Base$\rightarrow$ 0 |
| 10 | if(Task1$>$0.990) | Base$\rightarrow$ 3 | Base$\rightarrow$ 0 |
| 11 | if(Task2$<$0.480) | Rest$\rightarrow$ 4 | Rest$\rightarrow$10 |

Figure 2.4: A binary decision table representing twelve states for an agent and the transitional logic.

| Symbol | Meaning |
|---|---|
| $<$ | less than |
| $\leq$ | less than or equal to |
| $\geq$ | greater than of equal to |
| $>$ | greater than |
| $\approx$ | near ($\lvert a-b\rvert < \epsilon$) |

Table 2.1: The possible choices of comparison operators for Boolean tests in BDAs. For the study, $\epsilon = 0.05$ and values for a and b are in the range $0 \leq x \leq 1$.

In the BDA model, probabilistic mutation is also employed but instead of manipulating a single loci in the gene, it inaccurately copies the states of the mentor's BDA and applies the mutation operator to one of seven components of the states. These components are:

- The input action (stress, task1, task2).
- The comparative operator used for the Boolean test.
- The fixed numerical value the operator is applied against.
- The action executed if the Boolean test is true.
- The state to proceed to after the true action is executed.
- The action executed if the Boolean test is false.
- The state to proceed to after the false action is executed.

The distribution of the number of mutations (applications of the mutation operator) for the BDA's was varied and is discussed further in Chapters 4 and 5.

# Chapter 3

# An Agent-Based Model of Stress in the Workplace

The preliminary attempt at modelling stress in the workplace is the focus of this chapter. It is based on "An Agent Based Model of Stress in the Workplace" by Daniel Ashlock and Matthew Page which appears in the proceedings of the 2013 IEEE Symposium Series on Computational Intelligence on Evolving and Adaptive Intelligent Systems (EAIS) [6].

The purpose of this model was to establish an initial evolutionary algorithm using a String representation and to conduct a parameter study. This piece of technology represents a novel contribution to human resources and economics as well as a novel application of evolutionary algorithms to modelling stress. It was the hope of the researchers that the tools and technology designed within could establish a framework to build on for further research and potentially assist organizations and government policy makers in proactively dealing with stress amongst individuals.

## 3.1 Outline and Assumptions

Individuals early in their career will often attempt to mimic the behavioural patterns of successful individuals they encounter. This mimetic behaviour is attempted in order to advance one's career through the recognition of achievements. This achievement is typically in the form of work completed and the associated recognition that comes from senior leaders in the organization. In order to gain recognition from senior leaders, new employees usually have to work additional time to complete tasks outside of their assigned duties. This increased time requirement places stress on individuals and prolonged exposure can lead to physical illness [12]. Consequently, as there is a finite amount of time in a given workday, an individual must adopt an effective work ethic that allows for a sustainable work-life balance. In this chapter the imitation of high-status individuals is investigated as a possible indirect source of stress.

This study uses simple agents to model individual employees early in their career. The representation selected was a String representation as was described in Chapter 2. The memetic aspect of agent behaviour is also seen in many experiments of spatial games and evolutionary computation [47]. Despite the simplistic representation of agents, the fitness evaluation used for the worker agents is more complex. This type of technology appears in numerous studies of agent based modelling [27, 39, 77, 79]. The agents in this study have their time divided month by month, hour by hour into time spent resting, working on the assigned base task or working on a special project.

As this piece of work is a preliminary attempt to develop a model of stress in the workplace, it encodes a number of features and simplifications. These features include:

1. The quality of an individual's work decays as the stress of the individual exceeds an minimum tolerance [65].

2. To avoid firing, there is a minimum productivity requirement on the base task of an individual.

3. There is a positive correlation between the productivity on the base task and special projects and the overall status of the agent in the company. As productivity increases, so does status, assuming the agent has not been fired due to low performance.

4. Agents that possess low productivity and are not fired at each time step are re-

trained in an attempt to improve their performance [19].

5. Agents that are fired are replaced by newly hired agents that contain similar behaviour to existing agents within the organization. This attempts to demonstrate the process of hiring people that meet qualifications for a job.

6. New agents and low performing agents are assigned to a mentor for training. The mentors are selected from a varying top percentile of performers from the company.

7. Management's perception of the relative worth of special projects and the base task requirement varies randomly on a week-to-week basis. This encapsulates the changing priorities and decision-making that exists within corporate culture.

In addition to the simplifications and features encoded above, two novel concepts were encoded: covert drug use and proactive management.

The concept of covert drug use was introduced into a select few individuals in the model. The actual drug use was modelled by giving individual drug users a higher stress tolerance denoted by a separate stress response parameter $\alpha_d$. The drug users are the same as other agents in all other ways. The adverse effects of long term drug use were not encoded into the model and remain something to be investigated in the future. Drugs are perceived as being helpful in the short term for handling stress however despite this, the use of drugs has an impact on corporate welfare.

Proactive management was added as the final novel feature into the initial model. Using the productivity metric "hours spent working", management would target a specific number of workers $N_m$ for increased productivity through working slightly more. The population of individuals is sorted by the stress factor $N_p$ and the lowest stressed individuals have one of the 0 loci changed to a 1 or 2. This represents an incremental increase in the time spent at work and yet provides insight into corporate practices and the correlation to overall population productivity.

There were a number of hypotheses that the researchers tested in the parameter study. First, selecting mentors from different status tiers will have an impact on the overall productivity of an agent. Second, the covert drug use of individuals to enhance their stress tolerance will cause a mismatch in stress tolerances between mentor and trainee and will indirectly cause substantial damage to the trainee due to the unreasonable expectation of performance. Finally, the variability of the importance of special projects and that of the base tasks will impact the overall corporate performance.

## 3.2 Design of Experiments

As this study was the first attempt of modelling stress, a parameter study was used to determine the behaviour of the model. The baseline study was conducted with the parameters outlined in Table 3.1. As was outlined in Equation 2.2, the base productivity is seen as *stress factor*$\times$*hours on base task*. The training cohort represents the top 10 per cent of performers in the company. For each of the experiments conducted in the parameter study, 10 replicates were performed.

| **Parameter** | **Values** |
|---|---|
| Population Size | 120 |
| Time Length | 10 years (120 months) |
| Stress Tolerance | $\alpha = 0.05$ |
| Stress Floor | $\beta = 0.2$ |
| Special Project Worth | $0.4 \leq \delta \leq 0.6$ with uniform distribution |
| Minimum Productivity Requirement | 60 |
| Training Cohort | 0-0.1 |

Table 3.1: Baseline parameters used in experiments

Following the establishment of the baseline, the individual parameters were varied to ascertain their effect on the model. Minimum base productivity was varied from a minimum of 45 to a maximum of 105 in increments of 15. The hypothesis for this test was that as the base requirement for work increased, the overall stress of individuals would increase to match the increased work requirements. Additionally, as the overall amount of time on base task increase, less time would be available to focus on special projects ultimately leading to a decreased overall status of individuals. Furthermore, it was expected that as the base requirement was increased for individuals, the number of firings would increase to match as the overall number of employees unable to meet the base requirement would be greater. Following the tests on the minimum base productivity, the stress drop-off parameter $\alpha$ was varied to determine its impact on the model performance. For this study, the values were set at $\alpha$= 0.05, 0.03, 0.01, 0.005, 0.002 and 0.0001. The expectation was that individuals with lower abilities to adapt to stress will have an overall higher total average of stress, quantified by the stress factor $S(t)$. This inability to adapt to highly stressful situations is also predicted to negatively impact the productivity of individuals. Management's perception of the relative value of the base and special projects were then varied with the values set at $\delta = 0 - 0.2$, $0.2 - 0.4$, $0.4 - 0.6$, $0.6 - 0.8$, $0.8 - 1.0$. The hypothesis was that higher levels of importance placed on special tasks by management will yield higher overall status and recognition for individuals that are evolved.

The selection process for mentors was then tested through varying the percentile from which mentors could be chosen. Recall that mentors ranges are established as a pair of decimal numbers specifying a range in the unit interval; the highest performing members of the population are at zero and the lowest at one. These ranges were set at $0 - 0.1$, $0.1 - 0.2$, $0.1 - 0.3$, $0 - 0.2$ and finally $0 - 0.5$. The overall hypothesis for selecting different ranges was that as the number of possible mentors available for selection increased, the overall productivity of the population would increase. This is due to the larger pool of candidates from which to re-train individuals. The stochasticity of the parameter $\delta$ means that a narrow training pool will learn management's current preferences. The replacement fraction of low-performing individuals was then varied. The expectation from this test was that the larger the group of individuals being selected for retraining,

| **Parameter** | **Values** |
|---|---|
| Stress Drop-off $\alpha$ | 0.05, 0.03, 0.01, 0.005, 0.002, 0.0001 |
| Minimum Productivity | 45, 60, 75, 90, 105 |
| Delta Lower-Bound | 0, 0.2, 0.4, 0.6, 0.8, 1 |
| Training Range | 0-0.1, 0.1-0.2, 0.1-0.3, 0-0.2, 0-0.5 |
| Replacement Fraction | 0.1, 0.2, 0.3, 0.5, 0.7 |
| Number of Drug Users | 0, 5, 10, 15, 25, 50 |
| Drug Users Stress Tolerance $\alpha_d$ | 0.01, 0.005, 0.0025, 0.0001 |

Table 3.2: Parameters used within the study to establish baseline comparison between models.

the more optimal the overall base productivity would be. This is a result of allowing for the evolutionary algorithm to establish more ideal solutions at a faster rate due to the rejection of low performing candidate solutions. The values selected for the replacement fraction were set at 0.1, 0.2, 0.3, 0.5 and 0.7. Table 3.2 provides a concise listing of all parameters and values selected for the experiments.

### 3.2.1 Experiments with Drug Users

With the introduction of the covert drug users into the population, a second set of experiments was performed. The experiment was conducted with the same initial parameters as outlined in the base study, Table 3.1. The drug users however were encoded with a separate value, $\alpha_d$, that represented a substantially better stress tolerance than the non-drug using agents. For this simulation, the number of drug users was assigned as 5, 10, 15, 25 and 50 out of the population of 120 individuals. The goal of this test was to ascertain whether a non-drug taking agent, that is one with a lower stress tolerance, when assigned for retraining could replicate the behaviour of a drug taking mentor with a substantially higher stress tolerance. The hypothesis for this test was that the agents being trained would be unable to replicate the behavioural patterns of the assigned mentor and the result would be more stochasticity introduced into the model and large variance in the outcome of total population productivity.

Building on the first experiment through varying the number of drug users, the stress tolerances for drug taking agents was varied to study variation in the effect from the first experiments and to create a larger disparity between stress tolerances of the trainee and mentor. The parameter for the drug user's stress was varied, $\alpha_d = 0.01$, $0.0025$ and $0.001$, compared to the tolerance of the non-drug using individuals whose stress drop-off was set as $\alpha = 0.05$. In this instance the expectation was that as the drug users could tolerate increasing levels of stress, the mimetic behaviour of non-drug taking agents would result in higher stress levels overall. It was also anticipated that having drug users in the cohort of mentors will introduce a source of stress by training non-drug users to emulate drug users' behaviour without the enhanced stress tolerance granted by the drug use.

Due to the small number of replicates in the first two experiments with drug users, a third experiment was conducted to more accurately quantify the impact of drug use. For this experiment 4 sets of simulations were performed and 400 replicates were used instead of the original 10. The standard baseline parameters were used, as outlined in section 3.5. The number of drug user selected to be 0 or 10, and the training cohort set to the top 10 per cent or the second 10 per cent. The aim was to quantify the effect of drug use's interaction with the training process. The stress tolerances for the drug taking agents were set at $\alpha_d = 0.005$.

### 3.2.2 Experiments with Proactive Management

The final component tested for the String representation was the aspect of proactive management. For this experiment, the tests were structured similar to the large replicate test for drug users. The four sets of runs were repeated with the addition of proactive management $N_p = 30$. This resulted in the 30 least stressed individuals having to work one more hour per month.

### 3.3 Results

Over the course of numerous simulations, running the model with a variety of parameter combinations, the model exhibited behaviour as was predicted by the stated hypotheses with only a few exceptions. In varying the minimum productivity required to avoid firing, the base productivity increased to a maximum average of 84.67±1.54 hours. This maximum occurred when the minimum productivity was set to 90. In addition to the base productivity increasing, the number of firings concurrently rose. Figures 3.1-3.3 show how the number of monthly firings responds to the change in the required base productivity. Of note, once the threshold for minimum productivity increased past 100, there was a drastic drop in performance resulting in a complete firing and re-training of the workforce due to the inability of agents to meet the unrealistic expectation for the base task.

As demonstrated by figures 3.1-3.3, the model initially has time of sharp productivity increase representing a rapid adaptation of the agents trying to learn how to work. For base productivity requirements of 60 and 75 the amount of firing drops to near 0 in a smooth, nearly linear fashion. The slope of these curves, which reflect the learning algorithm training the agents to meet their base workload, are mediated by the mutation rate and could probably be accelerated by increasing it. When the base workload is at 90 the burn-in period to learn to meet base workload is longer than the duration of the simulation for 7 of 10 replicates. This represents either an unreasonable base workload or evidence that the mutation rate could profitably be raised.

With the addition of 5 drug users into the population and the minimum productivity requirement set to 60, the model exhibited a much higher level of stochasticity than that observed in the simulations without drug users. This can be seen in the productivity plots shown in Figures 3.4 and 3.5. The effect is likely caused by drug users training people to unreasonable expectations.

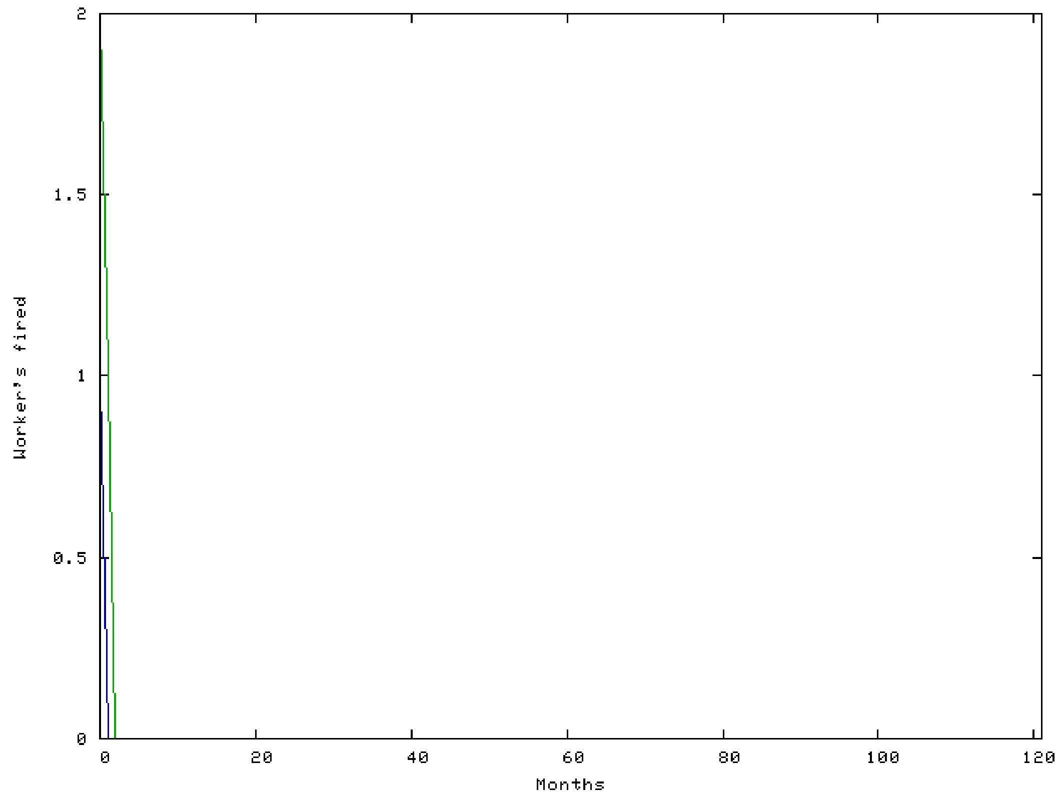


Figure 3.1: Number of monthly firings as a function of time for the experiments in the parameter study with minimum productivity set to 60.

The greatest effect on stress occurred in the variation of the stress drop-off parameter ($\alpha$). In general, as the rate at which stress drop-off decreased, the overall stress of the agents increased in conjunction with the number of hours worked. The maximum mean number of hours worked per day reached 8.76 and in several instances the maximum stress value reached the highest possible value of 1 (no stress) when the threshold for the stress drop-off was $\alpha = 0.0001$. This is to be expected as meeting high productivity goals is not difficult when agents have favourable stress response.

In the instances where management's recognition of special projects was varied more, the overall status of the employee increased when the values for were higher. Not surprisingly, this had the greatest effect on the awarding of recognition and the status outcome from the special project.

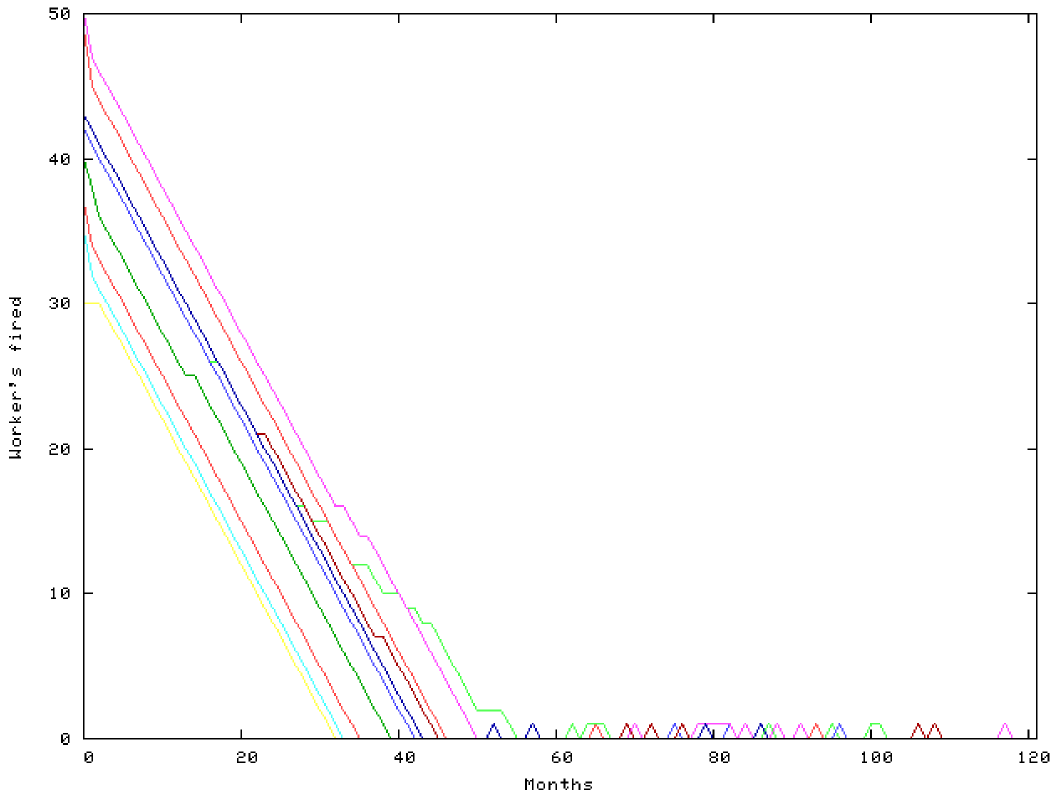


Figure 3.2: Number of monthly firings as a function of time for the experiments in the parameter study with minimum productivity set to 75.

When the replacement fraction of the population was varied, there was little impact to the model and very little fluctuation in the stress, base productivity, and status of the individual agents. This was not the expected result. The fraction of individuals assigned to retraining may matter more in simulations run against companies with smaller numbers of workers, a matter slated for future investigation.

When increasing the number of drug users in the simulations, a general result was that the maximum values attained for the stress factor continually peaked higher than 0.99 and the overall base productivity average was, in most cases, lower than non-drug user simulations. This corresponds with a higher firing average in general than the non-drug user simulations. When manipulating the training tier with drug users involved a more stable result can be achieved. When the training tier is set lower, as drug users

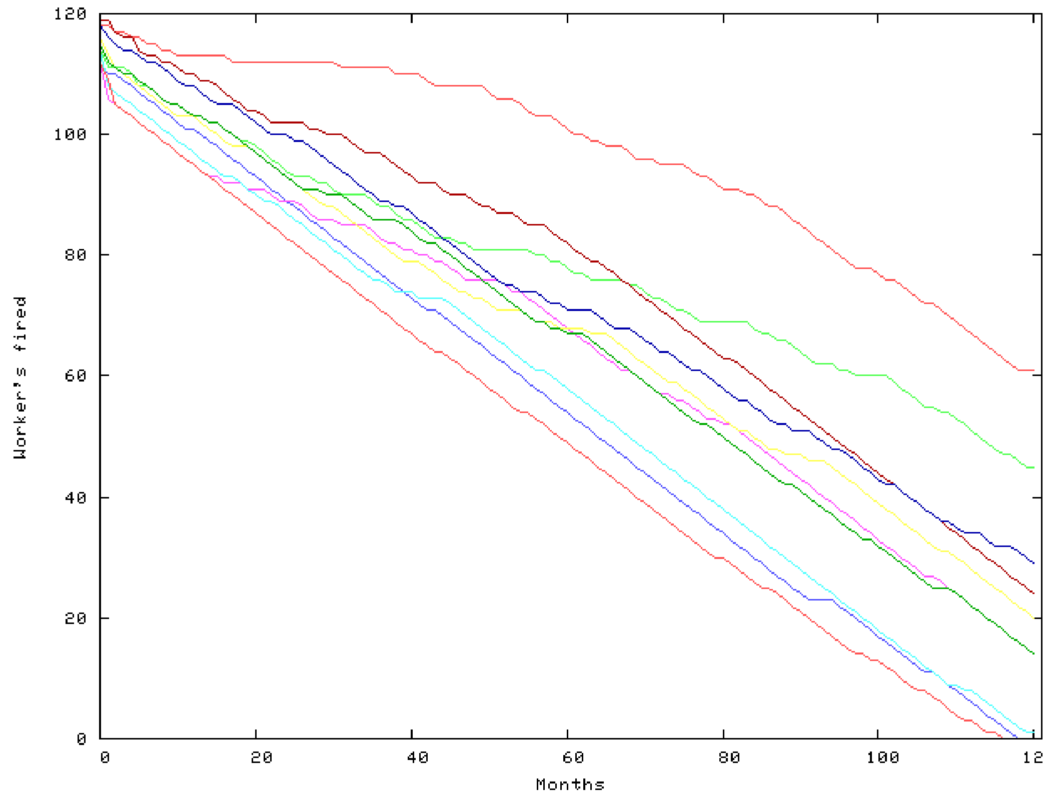


Figure 3.3: Number of monthly firings as a function of time for the experiments in the parameter study with minimum productivity set to 90.

are often high performers, this result is probably explained by having fewer drug users in the cohort. This can can be seen by comparing Figure 3.5 and 3.6.

It is worth noting that the parameter $\beta$ has very little impact on the simulation. Most agents stay far over on the left side of the stress response curve. This means that none get near the floor of the curve and so all effects can be modelled by varying alpha unless much smaller values of alpha are tested.

The curves shown in Figures 3.4, 3.5 and 3.6 have a similar shape. There is an initial sharp rise in productivity followed by a slow climb. The initial rapid climb represents the removal of patently absurd time allocations from the population. These include both those who do not work and those who work themselves to unsupportable stress levels. The continuing upward progression suggests that 10 years (120 months) is not

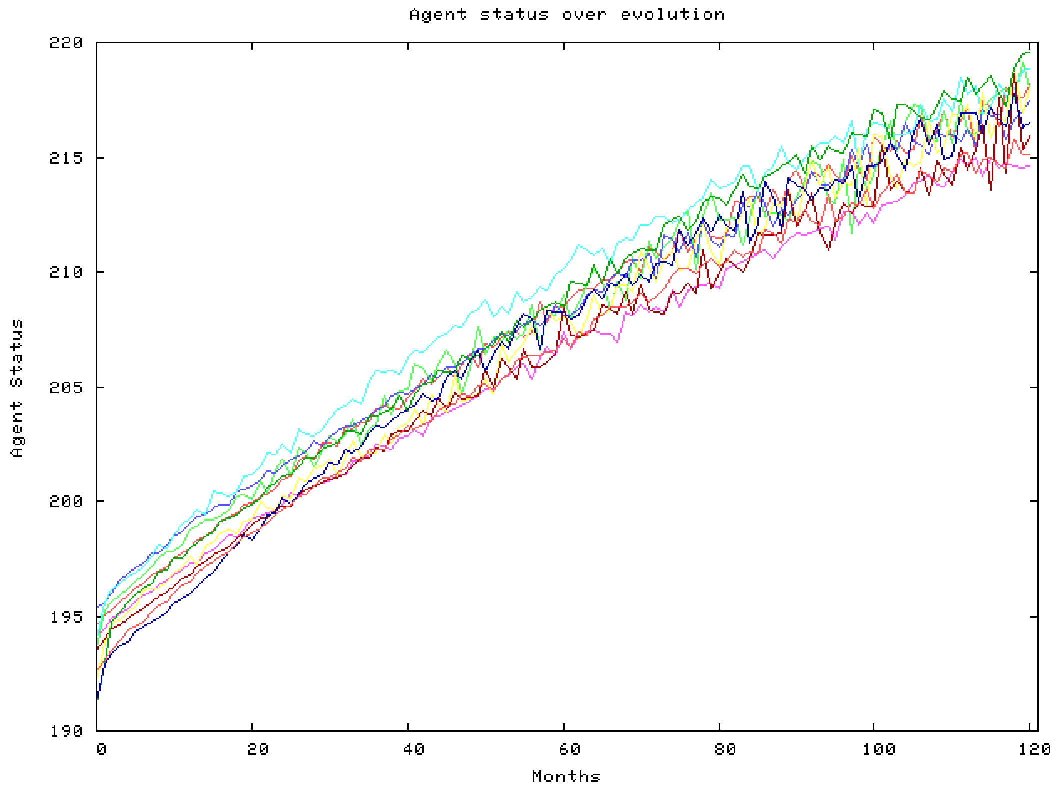


Figure 3.4: Response of 10 simulations with no drug users in a population of 120 with initial $\alpha = 0.05$, $\beta = 0.2$, and minimum productivity set at 60.

enough time for the system to come to equilibrium. The baseline experiment was re-run with the length of the simulation increased from 10 years to 100 years. The resulting productivity curves are shown in Figure 3.7. This is an unreasonably long simulated duration, but it shows that small amounts of improvement occur for a very long time. Coincidentally, there is a knee in the productivity curve near 120 months - suggesting that the simulation lengths chosen show the most vigorous period of adaptation of the system.

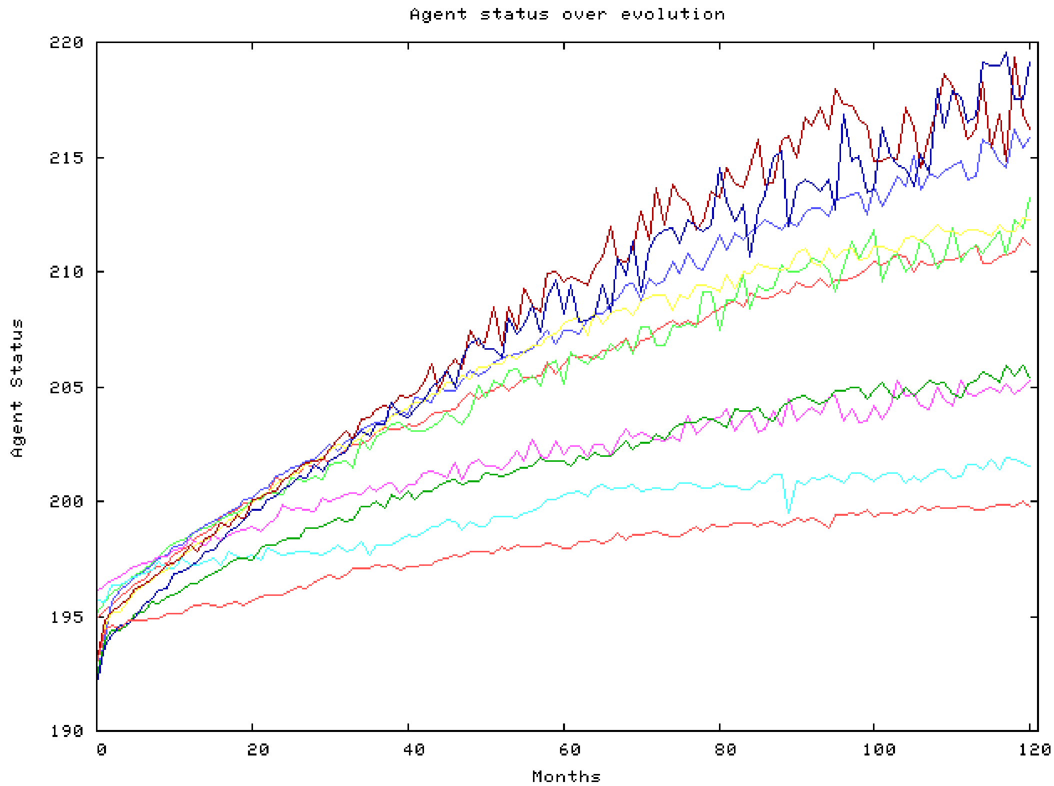


Figure 3.5: Response of 10 simulations to the addition of 5 drug users into a population of 120 with initial $\alpha = 0.05$, $\alpha_d = 0.01$, and minimum productivity set at 60.

### 3.3.1 Results for Large Replicate Study

Table 3.3 summarizes the results for the third experiment. This experiment compared using the top performing 10 per cent of the workforce as mentors with using the second 10%. These comparisons were made between a drug free company and in a company with 10 per cent of the workforce covertly using drugs.

The hypothesis that choosing a lower performing cohort for trainers would grant immunity to the negative effects of drug use is verified by these experiments. When the drug users are added into the population, overall performance increases at the trade-off of more firings. However, this result is not obtained when the training tier is selected to omit the drug taking mentors. Of note, no drug use still results in the highest overall productivity, despite the fact that drug use enhances the productivity of the users in the

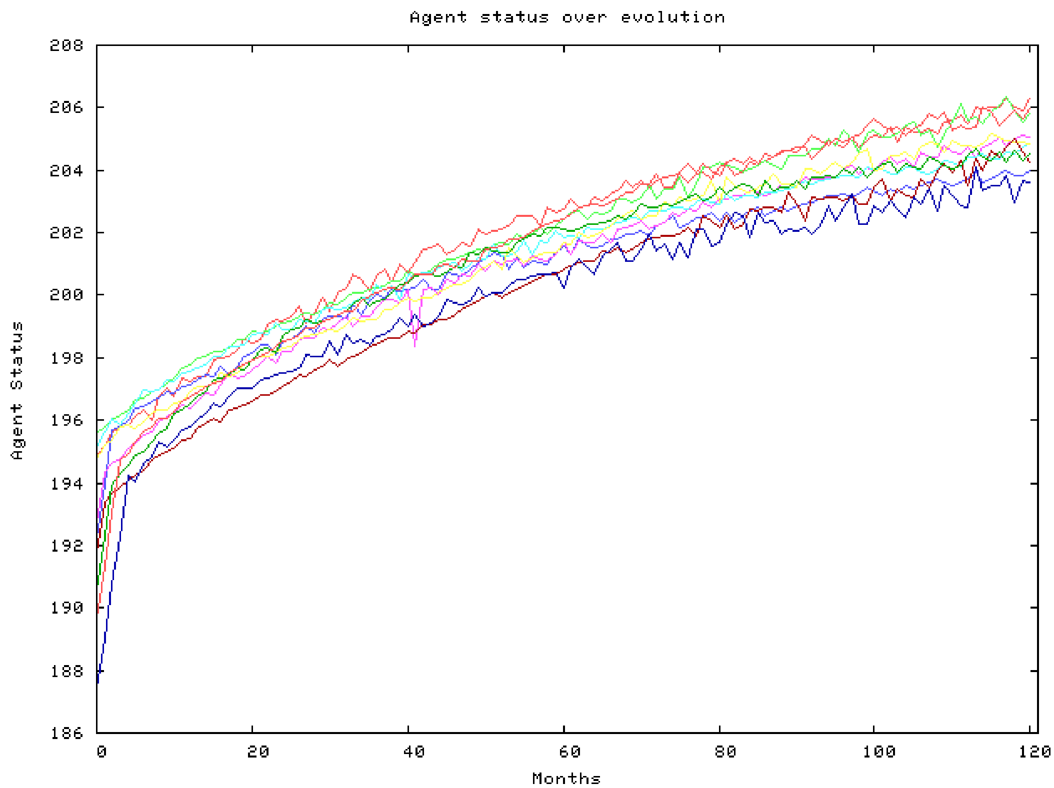


Figure 3.6: Response of adjusting the training tier for selection set at 0.2 to 0.3 with the inclusion of 5 drug users.

short term. Overall training of non-drug users to unrealistic expectations by drug users has a significant impact.

| Trainers | Drug Users | Base Task $\mathcal{N}_A$ | Performance $\mathcal{S}_A$ | Monthly Firings |
|---|---|---|---|---|
| 0.0-0.1 | 0 | 216.04±0.20 | 86.56±0.14 | 8.31±0.22 |
| 0.1-0.2 | 0 | 207.20±0.12 | 83.99±0.14 | 8.39±0.21 |
| 0.0-0.1 | 12 | 206.56±1.14 | 89.30±0.36 | 9.81±0.58 |
| 0.1-0.2 | 12 | 207.97±0.12 | 84.72±0.15 | 7.97±0.20 |

Table 3.3: Results on the agent's average base task, overall performance, and the number of monthly firings. Compared are drug free and 10% drug using worker populations in which the training cohort is drawn from the top or second 10% of the workforce. The values to the right of ± are the standard variance.

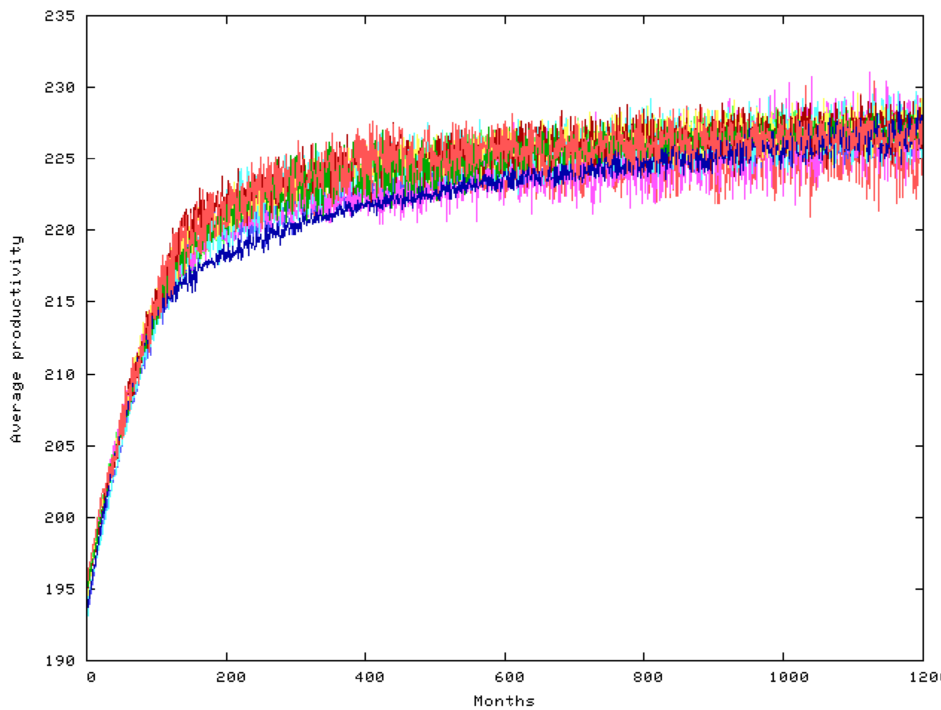


Figure 3.7: Productivity plots for ten replicates of the baseline simulation run for 100 years instead of 10 years.

### 3.3.2 Results with Proactive Management

Table 3.4 summarizes the results of the fourth experiment testing the impact of proactive management. The notable effect of these results is that proactive management increases low performance-related firing. This suggests that the net effect of proactive management is to push workers into high work (high stress) time allocations which lead to stress related degradation of performance. Ultimately it demonstrates that in the proactive management simulation, stressing workers though encouraging them to work extra time yields a bad result on overall corporate productivity.

The relative effects of drug use and choice of training cohort remain fairly similar. This resulted in an increased number of monthly firings as was experienced in the long term replicate simulations. This was slightly unexpected as proactive management should push the relatively low stress drug users to take on absurd workloads at which

| Trainers | Drug Users | Base Task $\mathcal{N}_A$ | Performance $\mathcal{S}_A$ | Monthly Firings |
|---|---|---|---|---|
| 0.0-0.1 | 0 | 215.77±0.19 | 86.91±0.41 | 9.04±0.22 |
| 0.1-0.2 | 0 | 206.78±0.13 | 84.24±0.14 | 9.19±0.22 |
| 0.0-0.1 | 12 | 207.41±0.91 | 84.84±0.32 | 10.25±0.51 |
| 0.1-0.2 | 12 | 207.68±0.15 | 84.60±0.16 | 8.58±0.20 |

Table 3.4: Results on the agent's average base task, overall performance, and the number of monthly firings when management is targeting $N_p = 30$ low-stress workers for working one more hour per month. Compared are drug free and 10% drug using worker populations in which the training cohort is drawn from the top or second 10% of the workforce.

point one would expect the experiment in which the drug users are dense in the training cohort to have a much larger effect from training to unrealistic expectation. The most probable explanation is that management always proactively urges the drug users to take on more work until they can no longer cope and stress becomes overwhelming. In a sense the higher rate of firing ends up affecting the drug users as well as the non-users so that proactive management somewhat ameliorates the effects of the presence of drug users.

## 3.4 Conclusions and Next Steps

The experiments and results outlined in this chapter demonstrate the feasibility of the trinary String representation agent-based model of stress in the workplace. The incorporation of covert drug use and a policy change (proactive management) were used to further establish how the model behaves and varies from the baseline parameters. The experiments conducted support the following conclusions:

- When no drug users are present, using the highest performing workers as mentors pays a significant bonus in productivity, exactly as one would expect.

- When drug users are present and the mentors are chosen from the top 10 per cent of performers (in all probability including many drug users), there is no significant productivity advantage to using the higher performers in the training cohort. There is a substantial disadvantage to using the best performers as trainers: the number of firings increases sharply.

- Comparing the experiments where the training cohorts are drawn from the second 10 per cent of performers, there is a small but statistically significant decrease in the number of rings when drug users are present. Since a few drug users have a high probability of being in the training cohort in this situation, this is most likely an interaction effect.

The most interesting results were in the experiments that explored the interaction of the presence or absence of drugs and the corresponding choice for training cohort. This demonstrated that by selecting a specific tier of mentors to be outside that of the tier containing the drug taking individuals ameliorates the stress tolerance mismatch that took place during the training of individuals. It was also seen that a substantial interaction appears in the studies exploring the drug users and the corresponding increase to the number of firings of non-drug users. This was due again to the inability of the trainees to match the difficult requirements and behavioural patterns of the mentors with a high stress tolerance.

While there are many simplifying assumptions in the String representation model, it serves to establish a framework for future research. It should be noted that the next step in the research process was to upgrade the representation to one that is more comprehensive and consists of nearly human-like behaviour. More complex agents that generate behaviour adaptively based on knowledge of their own current stress level, number of hours worked on each task, and possibly management expectation would generate more complex and, if well designed, more plausible agent behaviours. In modelling prisoner's dilemma [8, 21] a technique was developed to give agents emotions in the form of an artificial hormonal system. The agents with an emotional system were found to have a substantial competitive advantage. This emotional system could easily be adapted to manage the agent's stress response and even model effects like depression induced by working too hard. Another possible agent representation could be adapted from those used in the aforementioned prisoner's dilemma modelling [4], finite state machines with transitions driven by Boolean tests on inputs. Incorporating a $\lambda$-transition, that generates no action, into these binary decision automata permits them to make complex decisions based on multiple data items. The use of more complex agent representations provides more entry points for modelling additional phenomena. Use of binary decision automata with access to an emotional system would permit modelling, for example, of the impact of the work environment on the agents non-work life. While such models would require careful normalization against behavioural data to avoid representational sensitivity effects [23], the point where this type of complex system can be profitably modelled is approaching. This work was done in a second publication [7] and is explored in detail in the next chapter.

# Chapter 4

# Binary Decision Automata Modelling of Stress

This chapter outlines the second attempt made to model stress incorporating a new agent representation. It is based on Binary Decision Automata Modelling Stress in the Workplace by Matthew Page and Daniel Ashlock which appears in the proceedings of the 2013 IEEE Congress on Evolutionary Computation (CEC) [7].

The agent representation was modified in the second paper from a trinary String representation to a Binary Decision Automata (BDA) representation. The aim of this was to enhance the representation of the agents to incorporate a more diverse behavioural pattern through the incorporation of BDA. BDAs possess more cognitive capabilities in the form of knowledge of their current work states for base task and special projects as well as their current stress level. This allows the BDA to base the outcome of their behaviour on the Boolean tests against their known states.

## 4.1 Outline

### 4.1.1 Agent Representation

Worker agents are represented as BDAs. BDAs are augmented finite state machines with transitions driven by Boolean tests. The tests operate on three input variables: the agent's fraction of time thus far spent doing their base task or working on special projects and their stress multiplier. The stress multiplier 1 represents no stress and 0 represents total stress; the stress multiplier is subsequently explained. In each state 1 of the 3 variables is compared, using 1 of the operators in Table 2.1, to a constant. If the result is true one possible action and next state result otherwise, a second action and next state are used. The "near" operator uses a default value of $\epsilon = 0.05$ in this study. There were 12 states used for each BDA in this study. An example of a BDA can be seen in Figure 4.1.

| State | Test | If T | If F |
|---|---|---|---|
| 0 | if(Strss<0.874) | Rest→ 2 | Base→ 4 |
| 1 | if(Task2≤0.779) | Rest→ 7 | Base→ 2 |
| 2 | if(Task1≤0.289) | Spec→ 1 | Spec→ 1 |
| 3 | if(Task2>0.288) | Base→ 8 | Rest→10 |
| 4 | if(Task2>0.124) | Base→ 8 | Rest→ 6 |
| 5 | if(Strss>0.784) | Spec→10 | Rest→ 3 |
| 6 | if(Task1≤0.439) | Spec→ 9 | Rest→ 9 |
| 7 | if(Strss≈0.088) | Rest→ 0 | Spec→ 9 |
| 8 | if(Task1≥0.792) | Spec→11 | Base→ 9 |
| 9 | if(Task2≤0.217) | Spec→ 3 | Base→ 0 |
| 10 | if(Task1>0.990) | Base→ 3 | Base→ 0 |
| 11 | if(Task2<0.480) | Rest→ 4 | Rest→10 |

Figure 4.1: A Binary Decision table representing the twelve states for an agent. The symbol $\approx$ is the nearness operator; two numbers are "near" one another if $|a - b| \leq NR$ where $NR = 0.05$ in this study. Transitions have the form **action→ next state**. The BDA above represents the highest status achieved amongst the BDAs at $476.26 \pm 14.72$ which was calculated by $P_A = \delta\mathcal{N}_A + (1 - \delta)\mathcal{S}_A$.

As described in Chapter 3, there are three tasks an agent can choose to undertake: rest, base task, and special project.

## 4.2 Design of Experiments

As this study continues past research using a more complex BDA agent representation, a parameter study is performed on the new model to establish a baseline of comparison to the String representation used in the earlier study. The following parameters were used to establish the baseline: population of 120 agents was evolved for 120 months (ten years) using the parameters $\alpha = 0.05$, $\beta = 0.2$, $0.4 \leq \delta \leq 0.6$ with a uniform distribution. Base per month productivity to avoid firing was set to 60. Agents used 1-6 mutations during training with the number of mutations selected uniformly at random.

The parameters were varied from the baseline, one at a time, to establish how the model would react. The values used are summarized in Table 3.1. The minimum base productivity factor was varied in increments of 15 beginning at 45 and increasing to 105. The hypothesis was that as the base work requirement increased, the mean base productivity of agents would adapt to the changing requirement. Moreover, the agent would have to spend more time on their base work task reducing available time for rest and special projects. Additionally, it was expected that there would be a strong correlation between the increase in base productivity requirements and the number of individuals being fired.

The $\alpha$ parameter was varied to establish the agent's response to diverse stress tolerances and the impact on their performances. The values used for this were $\alpha = 0.05, 0.03, 0.01, 0.005, 0.002$ and $0.0001$ with the expectation that individuals with higher tolerances to stress would have lower average levels of stress as quantified by the stress factor $S(t)$. The $\delta$ parameter, management's perception of the relative value of the base job and special project, was varied with a lower bound of $\delta = 0, 0.2, 0.4, 0.6, 0.8$ with a range of 0.2.

The percentile from which mentors could be assigned was varied. The ranges are established by a pair of decimal numbers specifying a range in the unit interval where the highest performers are at 0 and the lowest at 1. The ranges for mentor selection tested were 0-0.1, 0.1-0.2, 0.1-0.3, 0-0.2 and finally 0-0.5. The underlying hypothesis was that the larger the base of mentors to select from the less diversity would be reduced. Varying the tier from which mentors are drawn represents the degree of challenge to the new agents being trained. The stochasticity of the parameter $\delta$ means that a small training pool drawn from high performers will over-learn management's current preferences. Finally the replacement fraction of the low-performing individuals was varied with the expectation that the larger the group being retrained, the greater the overall base productivity of the group. The values for the replacement fraction were set at 0.1, 0.2, 0.3, 0.5 and 0.7.

### 4.2.1 Experiments with Drug Users

A second set of experiments allowed for the introduction of a limited number of covert drug users into the population with the same initial parameters as defined in Table 3.1 with the exception of a higher stress tolerance resulting from taking drugs. For these simulations, the number of drug users introduced were set at 5, 10, 15, 25 and 50 out of the base population of 120 agents. The aim of this was to establish if the behaviour of the drug taking agents with a higher stress tolerance could be replicated by agents that are not taking performance enhancing substances. The hypothesis for this was that the agents would be unable to replicate this behaviour due to the inability to tolerate stress to the same degree as their mentor. This introduces additional stochasticity into the performance of the model. The drug users stress tolerance was varied with $\alpha_d = 0.01$, 0.0025 and 0.001 compared to the tolerance of the non-drug using individuals whose stress drop-off was set as $\alpha = 0.05$. As a consequence of increasing the tolerance of the drug taking agents, it was expected that the overall stress of the non-drug taking agents would be increased.

In order to better quantify the impact of drug use, a third experiment was conducted using 4 sets of simulations. These simulations used the baseline parameters with the exception that the number of drug users was set to 0 and 10 and additionally the training cohort was set to the top 10 percentile and second 10 percentile. The goal of this was to quantify the effect of drug users interaction with the training individuals. There were 400 replicates performed for this experiment as opposed to the ten replicates used in previous experiments. The baseline parameters and $\alpha_d = 0.005$ were not varied for this study.

## 4.3 Results

Based on the previous results and with the inclusion of the parameter study performed using agents represented as BDAs the model behaved as was hypothesized, with some exceptions. When the base productivity requirement to avoid firing was varied, it was found that the maximum average productivity attained was 85.02±9.70. This maximum occurred when the minimum productivity requirement was set to 90. Overall, as the minimum productivity requirement was increased the number of firings also rose. Of note, the BDAs only began to achieve higher status and productivity than the String representation when the minimum productivity was set higher than 75. Additionally, as the requirement was increased past 100, the base task productivity of the BDAs continued to increase whereas in the String representation it decreased. Moreover, the number of firings for the BDAs steadily increased even at productivity requirements of over 100, whereas the String representation agents had a drastic increase to the firing and re-hiring of the entire population for base work requirements over 100. Furthermore, there was only a slight drop in productivity of the BDAs past a base work requirement of 100 going from the aforementioned 85.02±9.70 to 84.22±9.78. These results confirm that the BDAs are more capable of dealing with stress and adapt much better to more demanding working conditions. The $\beta$ parameter had no noticeable effect on the model, as the agents consistently did not reach the stress floor.

With the introduction of 5 drug users into the population and with the minimum productivity requirement set to 60, the BDAs exhibited much less stochasticity than the String representation agents. This is demonstrated with the productivity plots shown in Figure 4.2.

The greatest effect on the model was when the stress drop-off parameter $\alpha$ was varied.

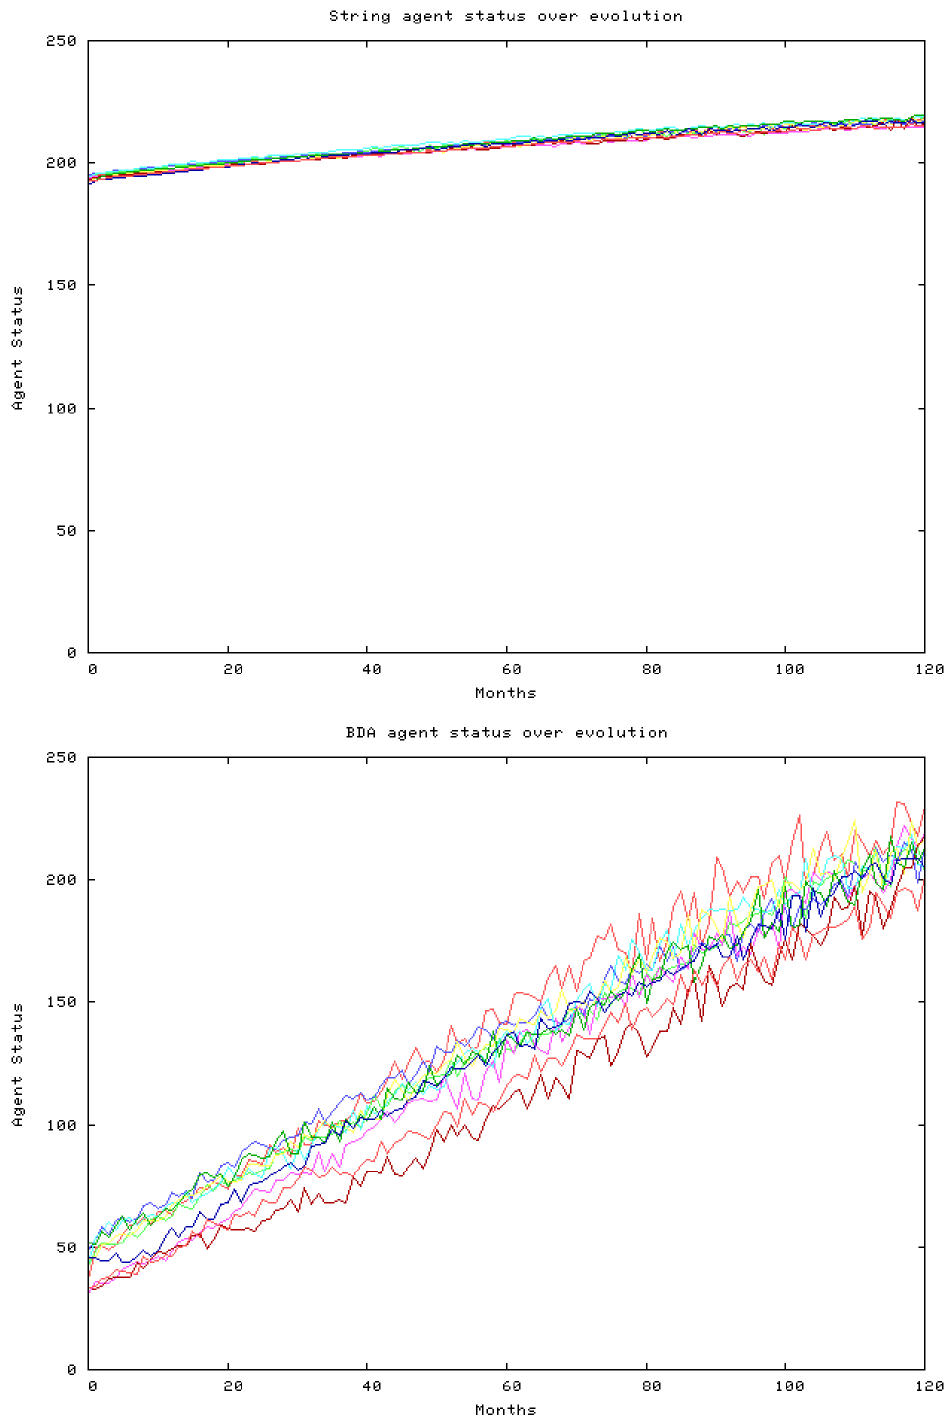


Figure 4.2: Model comparison of 10 simulations with no drug users in a population of 120 with initial $\alpha = 0.05, \beta = 0.2$ and a minimum productivity set at 60. Note that final productivity is higher for agents using the BDA representation.

Generally, as the rate of efficiency decay with additional hours worked was increased, the overall stress of the agents increased along with the number of hours worked per day. The maximum number of hours worked per day for the BDAs reached an average of 11.99 hours in contrast to the old String representation where the maximum reached 8.76 hours. This correlates with the expectation that with a greater ability to adapt to stressful situations, the BDAs are capable of working longer and harder. Note that with the increased maximum of hours worked, the productivity was substantially higher than that of the String representation (169.69±12.81 versus 90.49±1.83). As the $\alpha$ parameter was reduced to smaller and smaller values, the overall stress of the BDAs decreased attaining the least stress values of 1 repeatedly.

When management's recognition of special projects was varied, the overall status of the agents increased for higher values of $\delta$. Not surprisingly, this had a significant effect on the awarding of special recognition for individuals on special projects. The highest status achieving BDA is represented in Figure 4.1 where it achieved 476.26±14.72.

In the instances where the replacement fraction of the population varied, there was little impact to the model. The stress of individual agents and base productivity and status of the individual agents remained consistent. This was expected as it was found that, similar the String representation, the top mentors are always in the highest percentile of agents. Consequently when accounting for this, the overall productivity of agents does not substantially increase as the selection of mentors would still be the same.

Varying the number of covert drug users in the population generally decreased the status of agents and increased the average number of hours worked per day and number of firings. Of surprise to the researchers was that there was only a small impact to the total stress of the agents under the BDA model. This is another demonstration that the BDAs are reacting to their situation more effectively than the String representation agents.

Examination of BDA states show that they are evolving useful tests for enhancing

| State | Test | If T | If F |
|---|---|---|---|
| 0 | if(Task2≥0.979) | Spec→ 1 | Rest→ 4 |
| 1 | if(Strss≤0.797) | Base→10 | Base→ 2 |
| 2 | if(Task1≥0.289) | Rest→ 1 | Spec→11 |
| 3 | if(Task1≈0.235) | Base→10 | Spec→10 |
| 4 | if(Strss>0.314) | Base→ 9 | Base→11 |
| 5 | if(Task1>0.038) | Base→11 | Rest→11 |
| 6 | if(Task1≈0.552) | Spec→ 2 | Base→11 |
| 7 | if(Strss<0.088) | Rest→ 7 | Spec→ 5 |
| 8 | if(Strss≥0.953) | Spec→ 0 | Rest→ 7 |
| 9 | if(Task1>0.140) | Base→ 1 | Rest→ 3 |
| 10 | if(Task2>0.990) | Base→ 6 | Rest→ 5 |
| 11 | if(Strss≥0.645) | Base→ 5 | Base→ 6 |

Figure 4.3: The BDA above represents the lowest status achieved amongst the BDAs at 12.39±7.47.

an agent's success. For example in Figure 4.1, the best performing agent, the initial state is a Boolean check that stress $< 0.874$. Upon testing for this, if its stress level is too high, it will rest and then move to state 2. Comparing this with other successful agents, the stress Boolean test being first seems to be a logical choice as it governs the further actions. Thus, if an agent is stressed it ideally would want to rest otherwise it will work on its base work. Additionally, if it is not stressed, it does its base assigned task or a special project. In contrast to the worst performing agent, shown in Figure 4.3, one can see its initial state is to test that its special task $\geq 0.979$, then it works otherwise it rests. Moreover, once it moves to state 4 as per the diagram, it decides to work more. This type of self destructing behaviour causes the agent to over stress itself and typically results in the firing of the agent in the model.

### 4.3.1 Results for 400 Replicate Study

The results for the third experiment are summarized in Table 4.1. In this experiment, the top ten percentile was compared with the second ten percentile as groups for mentor selection and experiments were run with no drug users and 12 drug users. Of note while the BDAs begin at a lower overall productivity, their capabilities of learning surpass the String agent's as seen in Figure 4.4. This strong adaptive capability allows the BDAs to consistently out-perform the String agents in long term simulations.

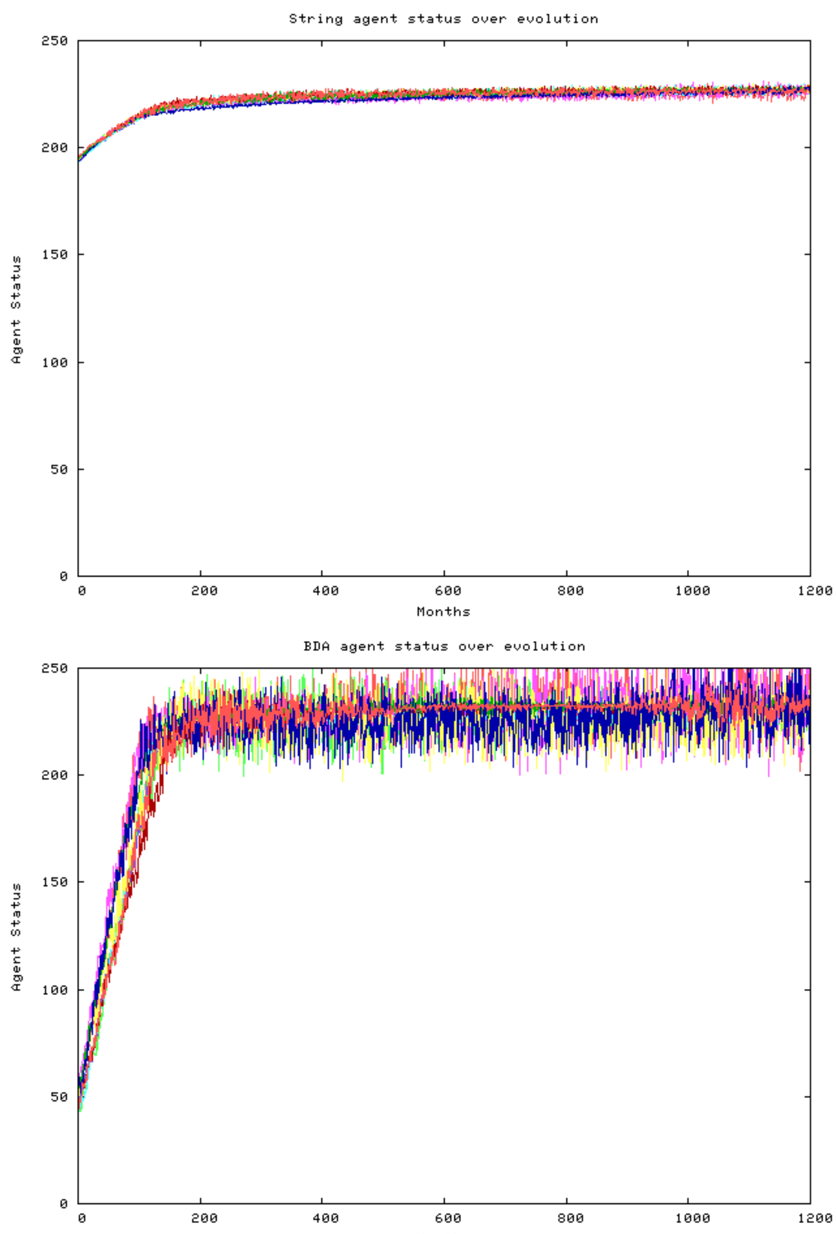


Figure 4.4: Model comparison of the String Representation versus the BDAs over 1200 months with the same aforementioned baseline parameters and $\alpha_d$=0.005. Agent status is the fitness value used in selection by the adaptive algorithm.

| Trainers | Drug Users | Base Task $\mathcal{N}_A$ | Performance $\mathcal{S}_A$ | Monthly Firings |
|---|---|---|---|---|
| 0.0-0.1 | 0 | 138.12±5.67 | 44.25±5.66 | 0.52±0.32 |
| 0.1-0.2 | 0 | 111.87±0.78 | 65.89±0.59 | 0.3±0.29 |
| 0.0-0.1 | 12 | 126.10±4.25 | 55.78±4.25 | 0.2±0.26 |
| 0.1-0.2 | 12 | 105.53±1.49 | 72.36±1.55 | 0.1±0.19 |

Table 4.1: Results on the agent's average base task, overall performance, and the number of monthly firings. Compared are drug free and 10% drug using worker populations in which the training cohort is drawn from the top or second 10% of the workforce.

### 4.3.2 Representational Impact

The relatively high early status (fitness) of the String agents and their lack of adaptability in contrast to the BDA based agents arises from the combinatorial character of the strings. A random string of substantial length must start with close to one-third of its loci assigned to each action. This happens, for the default parameter settings of the algorithm, to be an agent quite unlikely to be fired unless the actual arrangement of the characters leads to high stress and hence degraded productivity. The adaptive algorithm is able to rapidly sort out the order of actions and so we rapidly arrive at agents that have acceptable performance. Given that a relatively small number of changes are made to new agents generated by training, the string based agents lack the ability to change quickly. The multinomial distribution of tasks for the String representation are further explored in Chapter 6.

The BDAs are, essentially, evolvable pieces of code. They can easily generate all their actions with a very small number of states and so avoid the "I must be near average" effect that forces most strings to be uninteresting but acceptable virtual employees. A BDA can easily model an agent that does very little work (most actions are rest) or an agent that works to unrealistic expectations (burn-out), e.g. the BDA in Figure 4.3. One perspective on the change of representation is that the Strings and BDAs place a markedly different distributions on the space of behaviours. It is far easier to encode various types of mistakes with a BDA, something that is reflected in their low performance early in runs of the adaptive algorithm. Conversely, BDAs are adaptive and can generate a novel adaptive behaviour in a few mutations. Since the states encode responses to situations, these behaviours are transferable and far easier for the variation operators to modify.

## 4.4 Conclusions and Next Steps

This study builds on a previous study of a simple agent based model in the workplace. Using BDAs instead of a String representation, simple decision making capabilities are encoded into the agents. Additionally, the model shows the impact of the incorporation of covert drug use. In comparing the String agent's to the BDAs, the substantially better performance from the BDAs is likely due to their increased adaptive capability. This can also be attributed to the increased cognitive capabilities of the BDAs - they are aware of their own stress level and work performance thus far. As a result of having "memory" in the form of the descriptive statistics they use to make decisions, the BDAs have a more accurate representation of human behaviour and stress mitigation capabilities for more demanding working environments. While the base productivity begins at a lower state for the BDAs compared to the String representation, as time progresses, the BDAs overtake the String agents in productivity and generally have a lower number of firings and much higher status (fitness) than the String agents. When running the simulations for longer periods of time, 1200 time steps instead of 120, it is interesting to note that the BDAs begin their evolution around 90 to 100 monthly firings and consistently end with 0 to 1 monthly firings. This is important as it demonstrates the superior learning capabilities over the String agents who after 1200 time steps still have on average 5 to 10 monthly firings.

This model incorporates many simplifications but does allow for easy modifications and extensions. The encoding of more complex agents enables not only decision making but the capabilities of agents to more accurately adopt the behaviour of mentors as demonstrated by the differences in agent representation. This is of real interest as it begins to more accurately mimic human behaviour. The changes to the results based on varying workforce sizes have yet to be explored. This would allow for more accurate study of an individual's behaviour and their positive or negative effect on the workforce population.

# Chapter 5

# REDUCTION OF GENERALIZATIONS OF BINARY DECISION AUTOMATA

This chapter continues research on modelling stress. It reduces simplifications in the model by including stress decaying on weekends instead of resetting; drug-use behaviour is able to be learned and unlearned; and the individual stress tolerances are variable instead of static for all agents. Additionally the number of states in the BDA were varied to better understand the BDA representation and whether more states enables a more accurate representation of behavioural patterns.

This chapter continues to use a Binary Decision Automata as the representation for the agents. The goal of reducing the simplifications in the model is to allow for a more accurate representation of the real world workplace environment. This enhanced realism allows for more effective study of the interaction of stress and its effects on productivity in the workplace.

## 5.1 Design of Experiments

To establish how the model behaves a parameter study was undertaken with the fixed base values. A population of 120 agents was evolved for 120 months (10 years) using the parameters $0.03 \leq \alpha \leq 0.07$, $\beta = 0.2$ and $0.4 \leq \delta \leq 0.6$ with both $\alpha$ and $\delta$ chosen with a uniform distribution. The initial base productivity requirement was set to 60.

| **Parameter** | **Values** |
|---|---|
| Number of BDA States | 4, 12, 36, 108 |
| Minimum Productivity | 45, 60, 75, 90, 105, 120, 135, 150 |
| Max New Agent Mutations | 3, 5, 10, 25 |
| Max Training Mutations | 6, 10, 15, 25 |
| Probability of Drug Use | 0.01, 0.1, 0.25, 0.5 |
| Weekend Stress Reduction | 0, 0.1, 0.25, 0.5, 0.75, 1.0 |
| Prob. of Adopting Drug Use | 0.01, 0.025, 0.05, 0.1, 0.25, 0.5 |

Table 5.1: Parameters of the model that were tested.

Both agents that are retrained and those that are newly hired are created by selecting a mentor and then adopting a mutated copy of the mentor's BDA into the agent's. Agents had the possibility of 3 maximum mutations during re-training and 6 maximum mutations when they are new hires. The actual number of mutations is chosen uniformly at random in the range from 1 to the maximum. The weekend stress reduction was set to 0.5.

The model was then tested with the parameters outlined in Table 5.1 to ascertain how each parameter effects the outcome of productivity and stress on agents. The design of experiments is such that values are individually changed one at a time from the initial baseline parameters. The number of states that each BDA possesses was changed to the values 4, 12 36 and 108 respectively. The expectation was that each agent would need a minimum number of states to ameliorate the effects of stress and attain optimal productivity. The effects of varying the minimum base productivity required to avoid firing were then tested. These values were set at 45, 60, 75, 90, 105, 120, 135 and 150 out of a possible 240. The hypothesis was that as minimum required productivity increased the stress on the agents would increase as agents would have to spend more time working to meet minimum requirements.

The maximum number of mutations allowed for hiring and training were then varied separately. For the hiring of agents, the values were set at 3, 5, 10 and 25. The expectation was that the greater the number of mutations on an agent upon hiring, the longer the training process for the agent to obtain optimal performance. Secondly, the maximum number of training mutations were varied to include 6, 10, 15 and 25. It was anticipated that the greater the number of mutations taking place during the training process, the greater the noise introduced into the system and consequently the lower productivities and higher stress for the agents. The weekend stress reduction was then varied to determine the significance of more rest. The values were set beginning at 0 and were then varied to 0.1, 0.25, 0.5, 0.75 and 1. It was predicted that the greater weekend stress reduction values would correlate with higher productivity and lower stress values.

### 5.1.1 Experiments with Drug Users

In previous research, the effect of drugs was established demonstrating that in general the individual productivity of a drug taking agent was found to be higher due to the substantially higher stress tolerance. However, this increased productivity was at the cost of increased firings and stress on the agents trying to match the drug taking mentors capabilities. In this paper, the ability for drug taking behaviour is adapted to allow adoption and rejection of drug taking behaviour. This addition to the model is governed by the probability of adopting or rejecting drug use behaviour and the probability of using drugs when an agent. The parameters for adopting the behaviour were set at 0.01, 0.025, 0.05, 0.1, 0.25 and 0.5. The expectation, based on past results, was that with an increased percentage of drug users, the mean stress of the overall population would rise and furthermore that the number of firings would increase as well. Secondly, the probability of using drugs when possessing the drug taking behaviour was set to 0.01 and varied to 0.1, 0.25 and 0.5. It was expected that the greater the usage of drugs, the higher the performance of the agent as well as the lower the stress tolerance. Furthermore, it was anticipated that when selected as mentors, the higher drug using agents would cause a drastic increase in the stress level of the training agents and consequently increase the number of firings.

## 5.2 Results

Upon completing the parameter study and establishing how the model behaved, the outcomes were accurately hypothesized with a few exceptions. When varying the number of states that each BDA possessed the results to stress appeared to be valley shaped with the peaks of the valley representing the 4 state and 108 state BDA. The productivity and status of the BDA's was highest at 12 states achieving an average base productivity of 63.04±10.49 and average status of 86.47±14.11. Finally the average number of firings for the 12 state machine was 67.02±22.72. By comparison to the worst BDA, the 108 state machine, the average productivity was 28.89±6.08, average status was 36.76±7.77 and average number of firings was 95.89±18.26. The maximum status attained by any BDA was achieved with 12 states and was 398.71. The average age of the agents in the population varied with the change of number of states. The results of average stress were valley shaped as well with the minimum occurring at 12 states and was 41.47±12.20 and the maximum again at 108 states and was 64.68±13.18.

When the minimum productivity requirement was varied, the results were similar until the productivity requirement was set above 135. The average productivity of 12 state agents steadily decreased beginning at 63.12±10.44 for a requirement of 45, and decreasing until 62.69±10.88 for a requirement of 135. However, when the requirement was set above 135 the productivity dropped to 7.21±2.04. This drop in productivity correlated with a substantial increase in the average number of firings from 69.61±21.78 at productivity requirement of 45 to 115.16±15.41 at requirement of 135. This was due to the entire population being fired and re-hired repeatedly due to the inability to meet the excessive productivity requirements. The status of the individual steadily decreased from 87.43±13.99 at a productivity requirement of 45 to 82.02±14.29 at a requirement

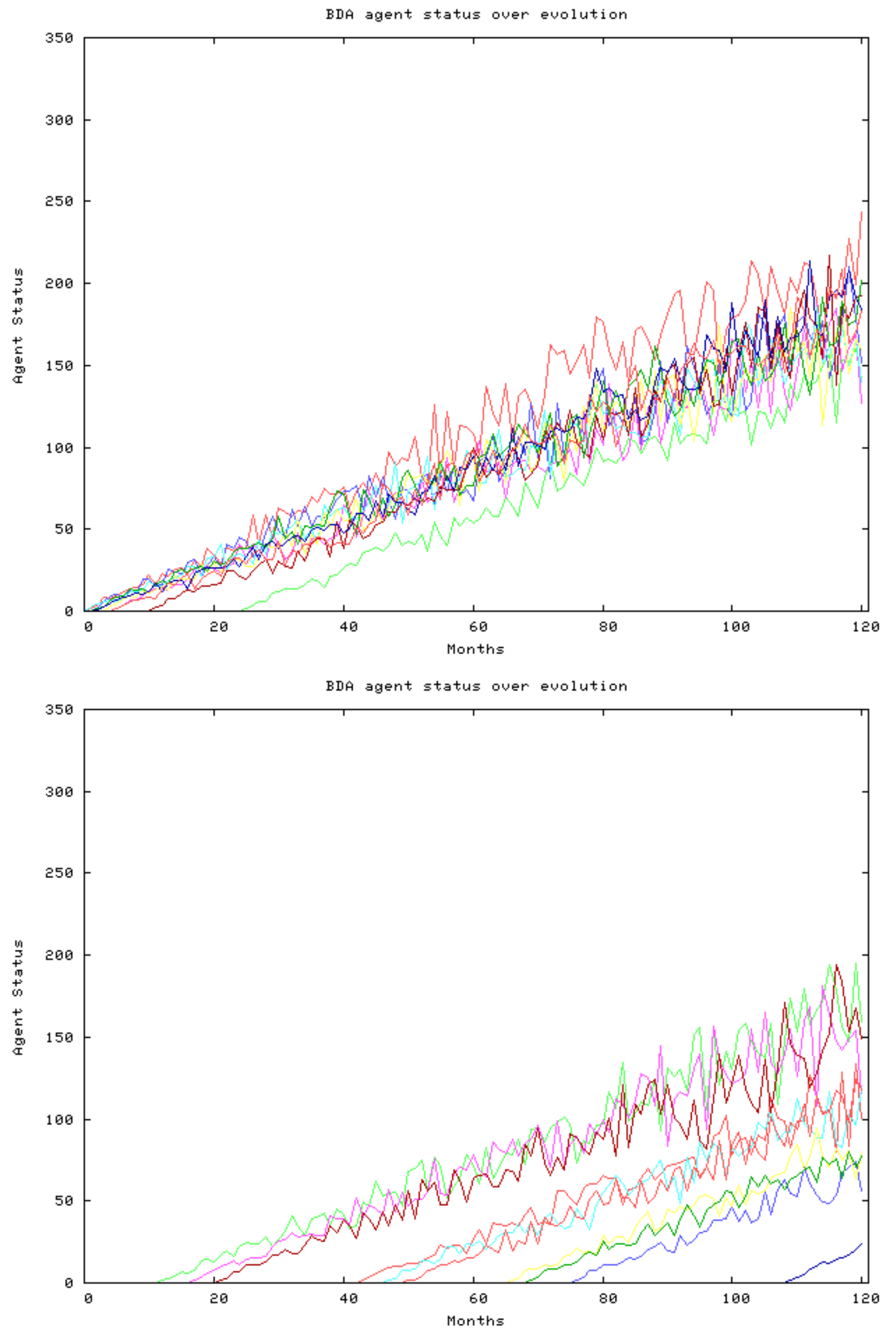


Figure 5.1: Outcome of productivity comparison of 12 states on the top to 108 states on the bottom using the base parameters of the model.

of 135. Again when the productivity was set above 135 the status of the individuals drastically dropped to 9.31±2.23. The average number of hours a week that agents worked also decreased as the productivity requirement rose. The most hours worked per week on average occurred at a productivity requirement of 45 where the maximum value of 4.79 was achieved. Furthermore, when the productivity requirement was set to 150, the number of hours the agents work dropped to 0.42 correlating with the sharp increase in firings at this productivity requirement.

Evaluating the effects of the number of retraining and hiring mutations on the BDA seemed to match expectations. For low numbers of mutations in hiring, the model would adapt with a small range of variation in terms of productivity change and stress on agents. With 3 mutations during hiring, the average productivity was $63.03\pm10.49$ decreasing to $58.99\pm10.80$ at 25 mutations. Additionally, the average stress of agents was $0.88\pm0.03$ at 3 mutations compared to $0.90\pm0.03$ at 25 mutations. Investigating the effects of the retraining mutations, the effect was greater than that of the hiring mutations. A low number of retraining mutations preserved the states of the BDA more accurately and correlated to a much better productivity and status result for the agent. At 6 mutations during retraining, average productivity was $63.04\pm10.49$ and average status was $86.47\pm14.11$. Increasing the number of mutations to 25 for retraining yielded an average productivity of $46.65\pm10.02$ and average status of $70.15\pm14.50$. These results were anticipated. For the hiring process as the agents would evolve through the duration of the experiment the starting point would be less important. However for the retraining process, maintaining the states of the mentor is more critical as it is the aim of the agent to preserve good behaviour as accurately as possible. Consequently increasing the number of mutations set the population back in terms of evolving to an optimal productivity and stress tolerance due to the stochastic changes of the mutation.

Increasing the probabilities of drug use and adopting or rejecting drug use behaviour had little effect on the model. This was surprising as in past work, the effect of a low number of drug users compared to a high number of drug users was noticeable. In this study, the probability of spontaneously starting drug use had no effect. When the

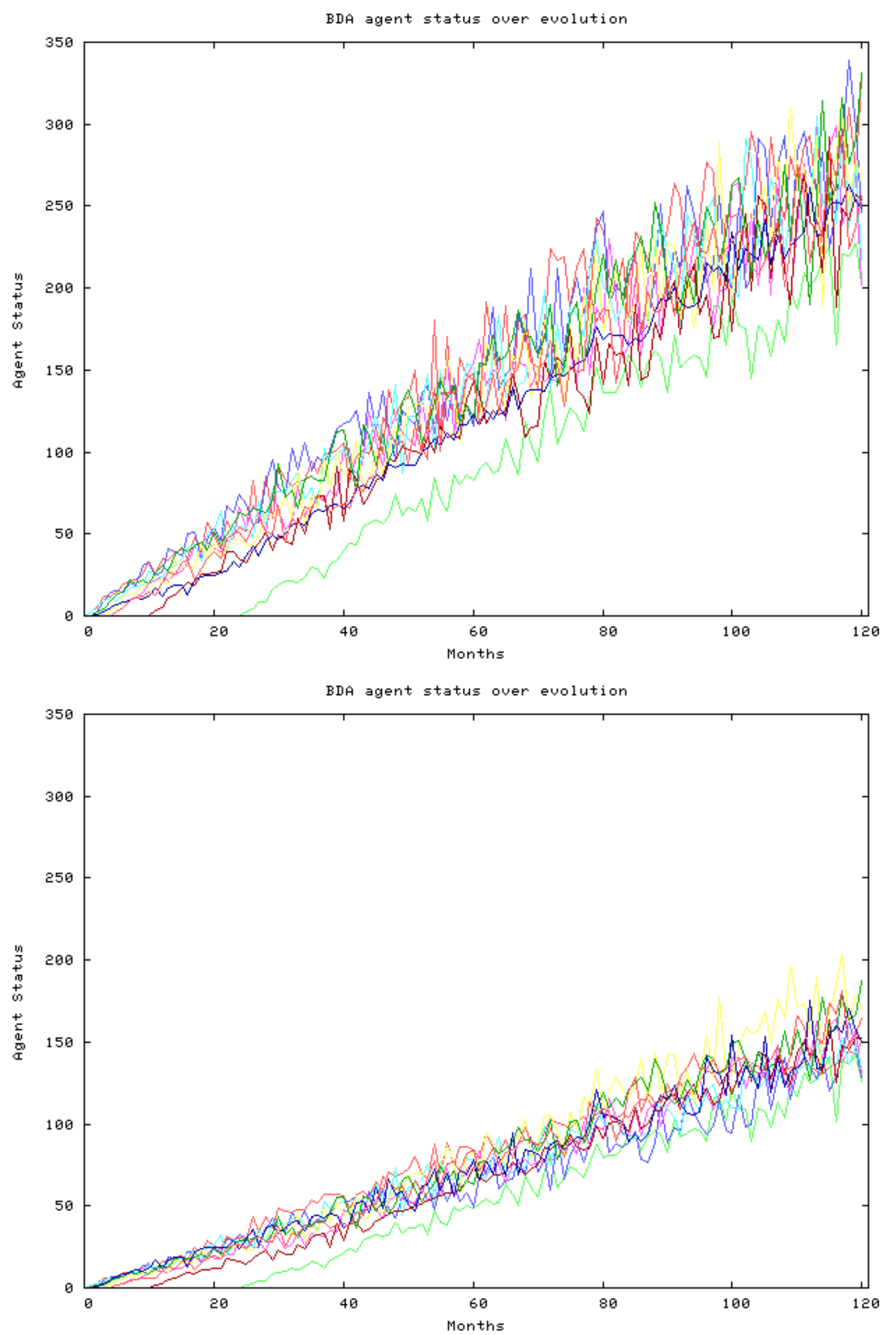


Figure 5.2: The effects on overall productivity for weekend stress reduction values of 0 seen on the top and 1 on the bottom.

probability of adopting drug use through training was increased, it increased the number of drug users in the population as expected however, had little to no effect on stress, productivity, and status of individuals.

The most substantial effect on the model occurred when the weekend stress reduction of the agents was varied. For low stress reduction values, that is those greater than 0.75, the stress and the productivity of agents was much lower than for the high stress reduction values of less than 0.25. In these instances, the productivity decreased from 98.10±16.13 to 47.46±8.17 corresponding to values of 0.1 stress reduction and 1.0 stress reduction respectively. The effects of weekend stress reduction on stress can be seen in Figure 5.3 and on productivity in Figure 5.2. Furthermore, the stress of the agents decreased from 1.00±0.00 to 0.86±0.03 for the aforementioned stress reduction values. Note that higher values of stress denoted less stressed individuals. Thus for high weekend stress reduction values, the stress of individuals is reduced to its lowest possible value. Another interesting reduction was that of the maximum possible status attained by agents. For a weekend value of 0.1 the maximum possible status peaked at 398.71 representing the best agents in the system. For a weekend value of 1, the agent's status achieved the lowest possible status observed at 300.07. The number of hours an agent worked on average also correlated to the values of weekend stress reduction. For higher reduction values, the agent would work more hours on average per week achieving the highest possible values in the system of 5.21. With low reduction values, the agents worked less hours and attained a minimum hours worked per week of 3.89.

## 5.3 Conclusions and Next Steps

The number of assumptions used in the model were reduced in an attempted to upgrade the model into a more accurate representation of the real world. Consequently, building on the previous research on this model [6, 7], the additions in this study have created some useful parameters that will have to be studied further to ascertain their full effect. Employing the use of BDA's to depict agent behaviour has proven to be more successful

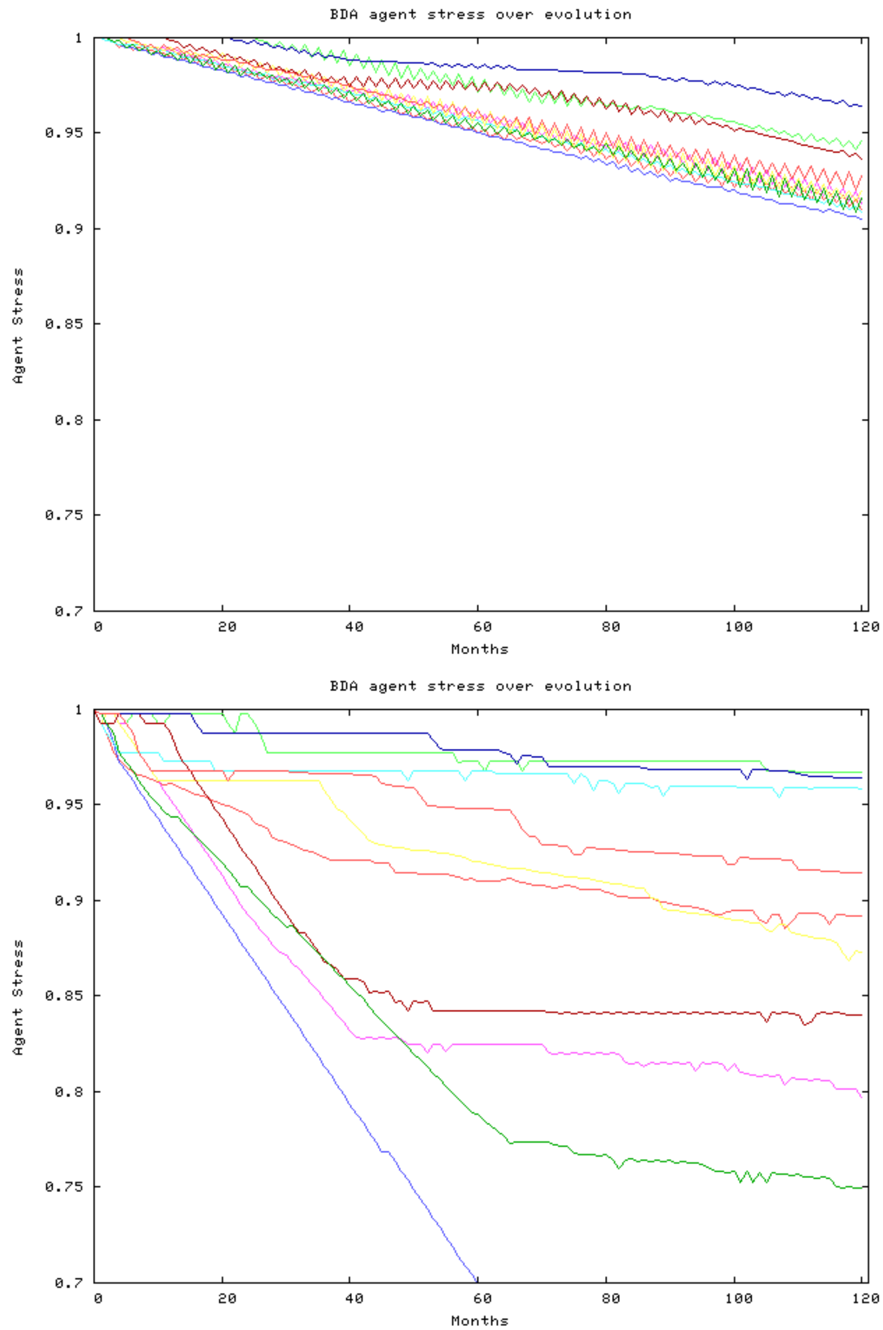


Figure 5.3: The effects on overall stress for weekend stress reduction values of 0.25 as seen top, and 1 as seen bottom.

in comparison to past String representations of agents [6]. The BDA's are proving quite adaptive to greater stress tolerances and work productivity requirements. This can

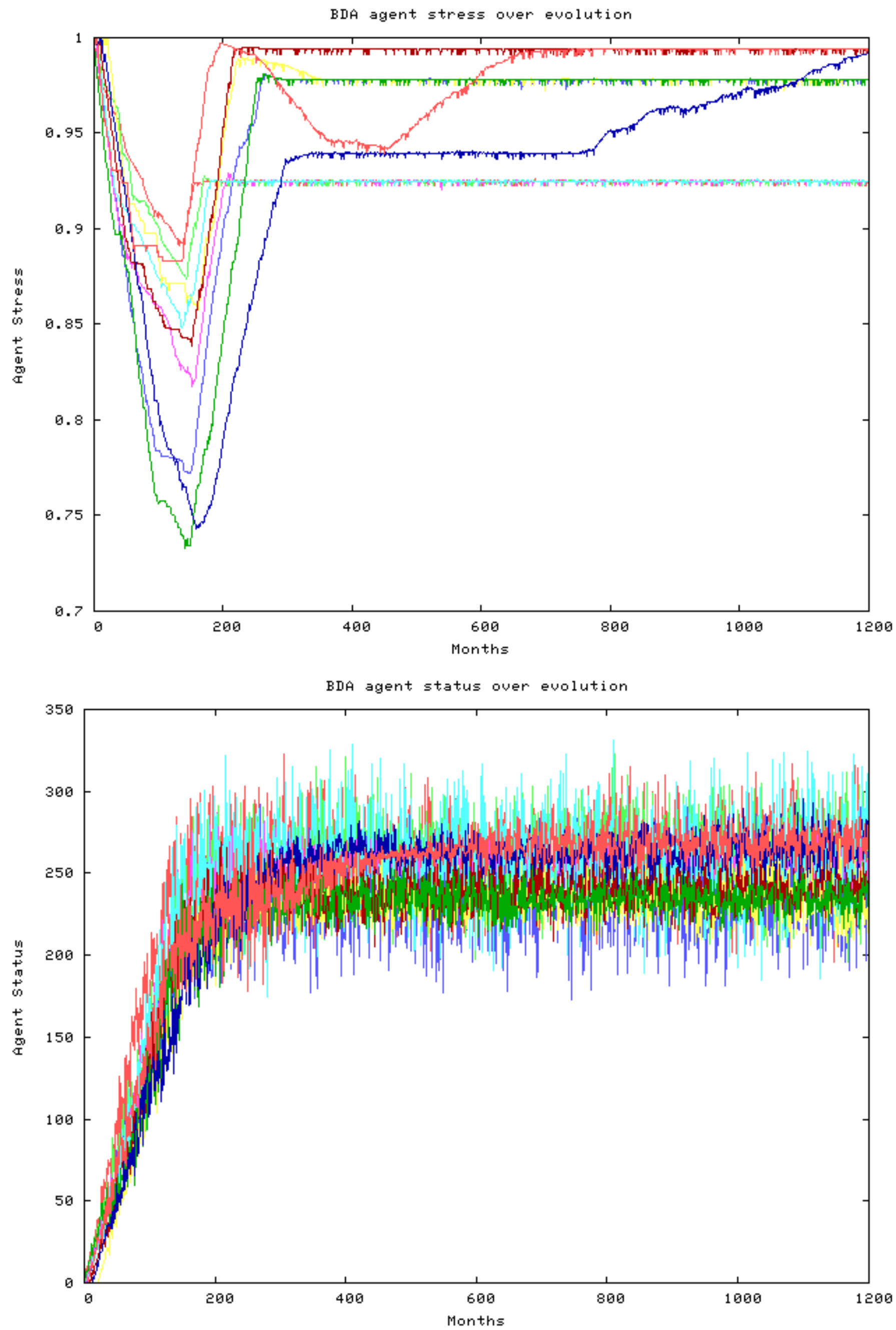


Figure 5.4: Stress and productivity of agents for base parameters of the model run for 1200 months.

be accredited to BDA's adaptive behaviour and memory component contained within the inputs for the states. The number of states appears to drive the outcome of the agents behaviour as is seen in Figure 5.1. It would seem that an agent requires enough states to make a basic decision but not so many that it takes longer to finish useful retraining than the duration of the simulation. Using this optimal number of states allows for increased adaptive and cognitive capability of the agents and establishes a more human-like representation of behaviour in the workplace. It is suspected that the increased adaptive capabilities of the BDA's ameliorates the effect of drug users in most situations. There appears to continue to be a strong correlation between stress and work productivity as well as work productivity and status which suggests that although the model is simplified it seems to encapsulate many aspects and nuances of the workplace. This can be seen in Figure 5.4 where the effects of stress drop-off for long term simulations as agents fully adapt to their working environment.

# Chapter 6

# Further Analysis of String Representation

The performance of the String representation agents is investigated in further detail in this Chapter. As was noted in the experiments and parameter study conducted in Chapter 3, and subsequent representational impact analysis conducted in Chapter 4, String agents possessed average performance with only a few outliers. It was suspected this had to do with probability of the possible actions for each loci having a one third chance of being selected when the individual was initialized.

To further understand the impact of changing the probabilities for each action being initially selected, this study was conducted to ascertain the effect on status, productivity and stress of individuals. The results were then compared to the previous performance of the String agents described Chapter 3.

The probability of obtaining any given number of outcomes can be calculated. In the case of the String representation used, there were 3 possible choices to select from: 0 for rest, 1 for base task and 2 for special project. These three values comprised the alphabet from which the possible actions for each agent can be selected [3]. The total length of the string for each agent is 240 characters as outlined in Chapter 2.

| Parameter | Values |
|---|---|
| Simulation Length | 120, 1200 |
| Initial Rest | 0.25, 0.3, 0.33, 0.4, 0.5 |
| Initial Base Task | 0.25, 0.3, 0.33, 0.4, 0.5 |
| Initial Special Project | 0.25, 0.3, 0.33, 0.4, 0.5 |
| Base Productivity Requirement | 60, 75, 90, 100 |
| Mutation rate for Training | 0.05, 0.1, 0.15, 0.25, 0.5 |
| Number of Drug Users | 0, 5 |

Table 6.1: Parameters of the model that were tested.

## 6.1 Design of Experiments

For this study the initial probabilities were modified in the code from a fixed one third probability for each possible action to two parameters: *Prest* and *Ptask*. From this the probability of the special project was established by $1 - Prest - Ptask$. The baseline parameters for the experiments conducted are the same as those found in Table 3.1. Following the variation for the distribution of tasks, the number of drug users were modified between 0 and 5. The base productivity requirement was also varied beginning at 60 and increasing in intervals of 15 to a maximum of 100. The mutation rate for training was also varied with values of $\mu$ = 0.05, 0.1, 0.15, 0.25, 0.5 to determine its impact on the training with different distributions. Finally the simulation lengths were selected as 120 and 1200. A concise listing of all the parameters tested can be seen in Table 6.1.

## 6.2 Results

The near average performance of the String representation seen in Chapter 3 could be overcome when the initial distribution of the base task was varied. When the initial distribution of actions is varied, the results significantly change the performance of the agents. With low initial values for the base productivity, noisy oscillatory behaviour is exhibited as seen in 6.4. When the parameter $Ptask$=0.25, the agents achieved very low status and productivity. In many cases the status achieved was 0 and the productivity achieved a low for the experiments at 59.57±1.19 when the productivity requirement was set to 75 and the initial distribution for $Ptask$=0.25. This is likely do to the fact that the agents need to adapt their poor working behaviour to match those expected of the minimum productivity requirement. In contrast to this the stress of the agents was lowest on average when the initial Prest values were highest achieving a maximum stress value of 0.99±0.01. In general when the Ptask values were higher, the overall number of firings decreased to a minimum of 0 when $Ptask$=0.5. When the parameter for retraining mutation was varied, results remained very close suggesting that mutation did not have a significant impact on the retraining of agents. With the addition of 5 drug users into the population when the initial distributions was set to a fixed one third, the results were the same as observed in Chapter 3. However, with values for Ptask set to 0.5 there was little variance observed in all the results suggesting that if the initial base task distribution is high enough, the effect of adding drug users can be ameliorated due to the agents possessing an sufficient capacity to meet the stress and work requirements of the drug mentors.

When the initial base task productivity was set to a one third probability for each task as per the original structure for the String representation in Chapter 3, the expected result occurred. This consisted of a steadily improving population as seen in Figure 6.2. Overall stress for agents increased when the base productivity requirement was set higher and additionally it was discovered that when values for $Prest$ were decreased, the stress increased as well. Note that for the averages of the special tasks, when values for base productivity were set to 60 and 75, the results appeared to be normalized around the distribution of $Ptask$=0.33. However when the productivity requirement was set to 90 the results reversed and formed a u-shape, seen in Table 6.2.

When covert drug users were added into the simulations, the result of changing the distribution seemed to ameliorate the effects seen before. The model demonstrated less

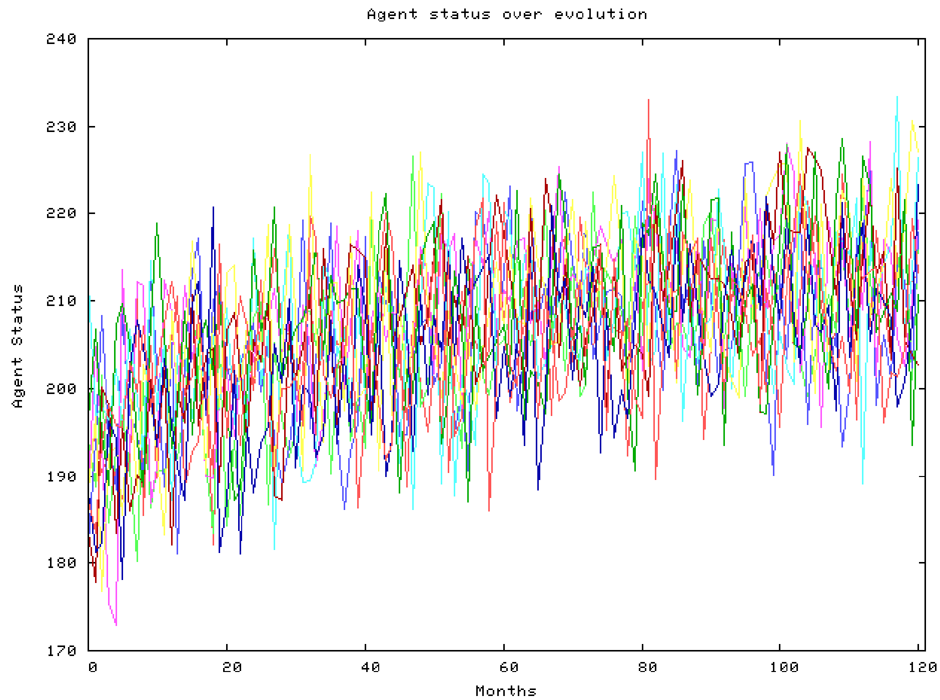


Figure 6.1: Productivity plot of the String representation with initial distribution of rest, base task, and special project as 0.33, 0.33, and 0.33, respectively. This test was conducted with all other parameters matching those in Table 3.1 with the inclusion of 5 covert drug users with a stress tolerance of $\alpha_d$=0.005.

variance than was expected when drug users were added. In Chapter 3 it was noted that when drug users were added into the simulation, the productivity of the agents diverged, as seen in Figure 3.5. In contrast to this Figure 6.1 demonstrates that the productivity no longer diverges; however, the overall movement of productivity is more noisy and oscillatory.

## 6.3 Conclusions

This study aimed to further explore the complexities of the String representation and the near average performance of the agents. While several of the hypotheses were affirmed, the agents exhibited much more diverse behaviour than originally expected. In particular, the noisy oscillations seen in Figure 6.4 and the ceiling of productivity in Figure 6.3

| Min. Prod | Ptask | Special Task Average |
|---|---|---|
| 60 | 0.25 | 59.97±1.22 |
| 60 | 0.33 | 79.45±0.99 |
| 60 | 0.5 | 55.87±1.12 |
| 75 | 0.25 | 60.43±1.25 |
| 75 | 0.33 | 79.05±1.23 |
| 75 | 0.5 | 56.37±1.15 |
| 90 | 0.25 | 79.05±1.23 |
| 90 | 0.33 | 55.87±1.12 |
| 90 | 0.5 | 63.19±1.10 |

Table 6.2: Average of special task results.

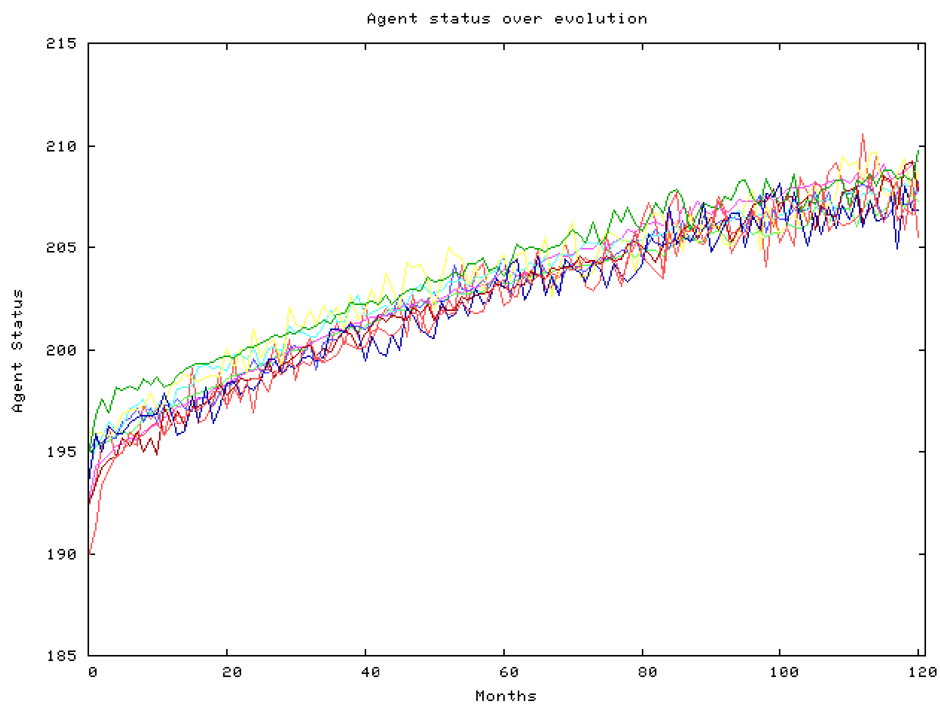


Figure 6.2: Productivity plot of the String representation with initial distribution of rest, base task, and special project as 0.33, 0.33, and 0.33, respectively. This test was conducted with all other parameters matching those in Table 3.1.

demonstrate that further investigation into the behaviour of the String representation should be conducted in the future if used to model stress in the workplace. It would also appear that the mismatching of stress profiles due to training with a drug mentor can

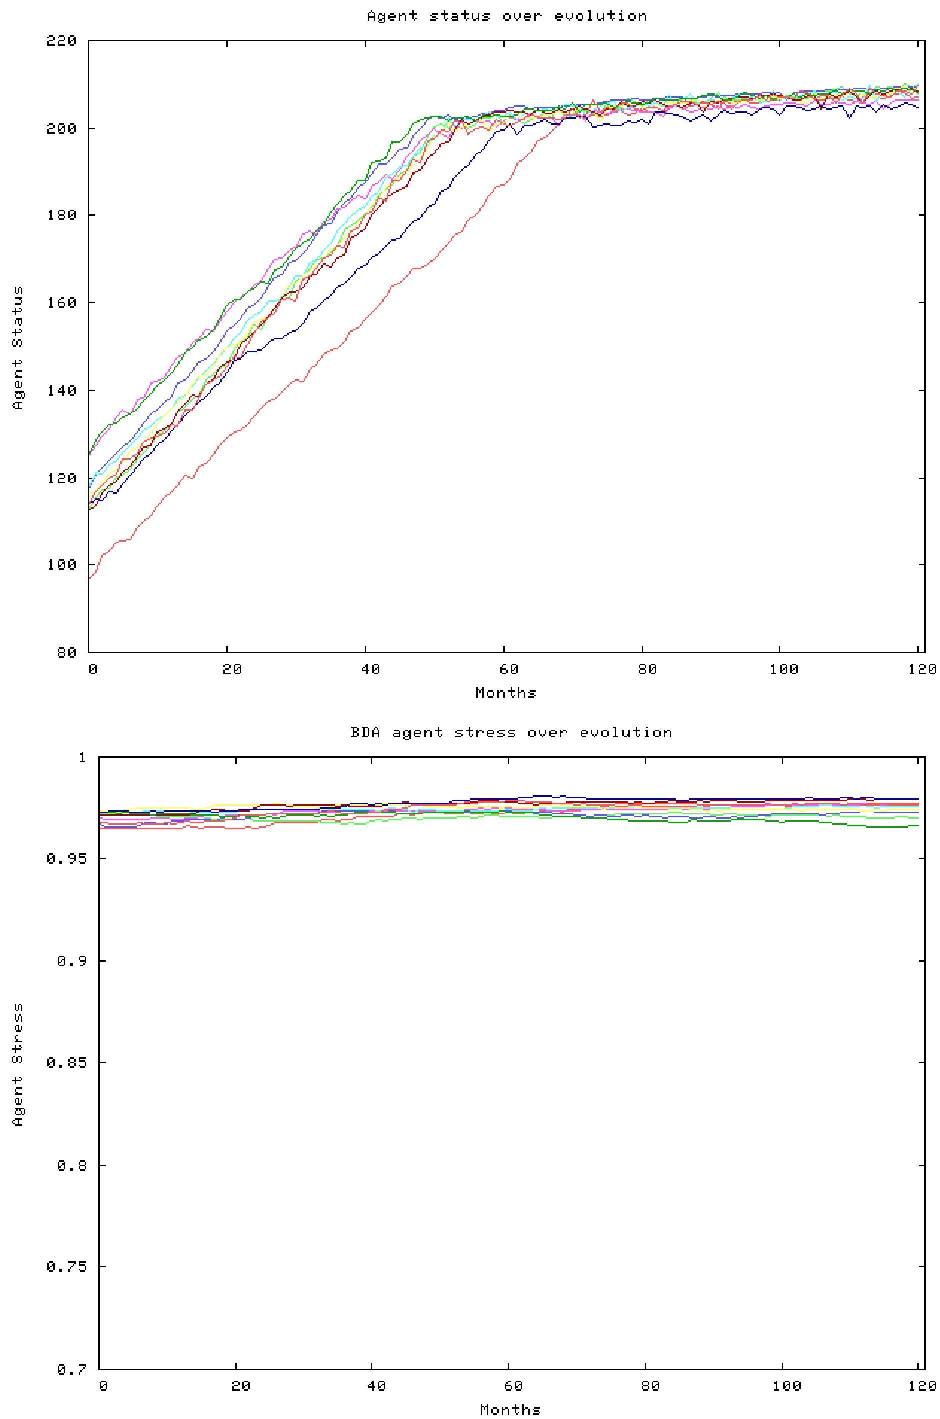


Figure 6.3: Productivity and stress plot of the String representation with initial distribution of rest, base task, and special project as 0.25, 0.50, and 0.25, respectively. This test was conducted with all other parameters matching those in Table 3.1.

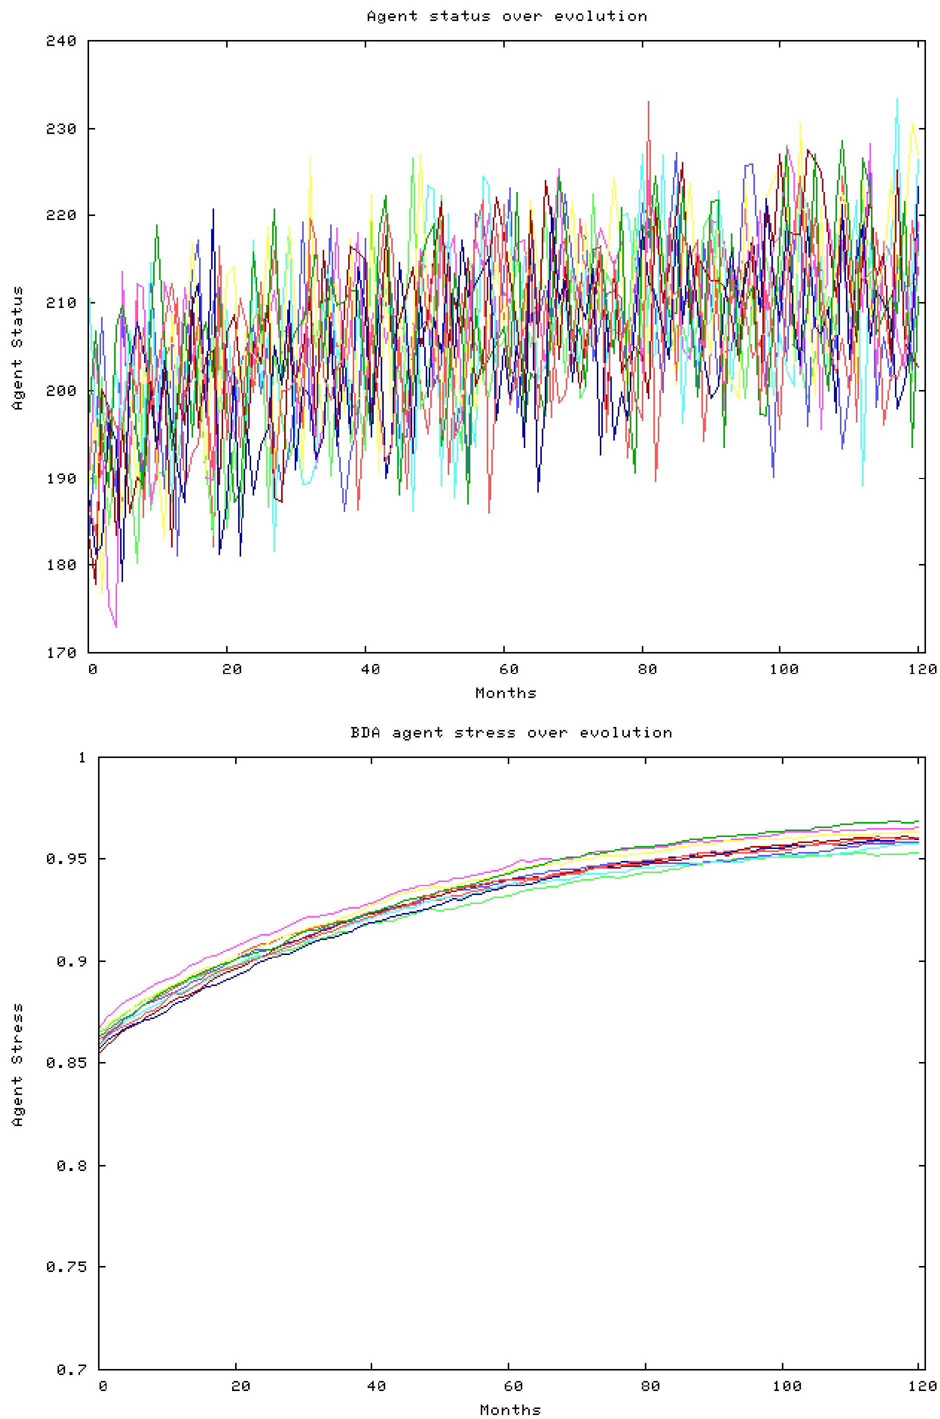


Figure 6.4: Productivity and stress plots of the String representation with initial distribution of rest, base task, and special project as 0.5, 0.25, and 0.25, respectively. This test was conducted with all other parameters matching those in Table 3.1.

be overcome by setting the initial *Ptask* distribution higher than the equally distributed values as was done in Chapter 3.

# Chapter 7

# Conclusions and Future Research

This study demonstrates that Computational Intelligence can be effectively utilized to model complex social issues such as stress in the workplace. Agents encoded through String representations yield solutions to the model however the complexities due to fixed distribution of tasks needs to be researched further to ascertain the full potential of the representation. The BDA representation was a more effective predictor for modelling the nuances of the real world in comparison to the String representation. While corporate drug use increases individual productivity, this increase leads to a lower overall productivity in the organization due to an increased number of firings. Adding a weekend stress reduction parameter resulted in the overall productivity of the organization being lower, but is a more realistic depiction of reality. Varying the multinomial probabilities of tasks for the String agents affects agent performance and leads to more diverse performances. Overall this model, albeit simplified, is a useful foundation for further research and is of importance in attempting to establish a healthy work-life balance for individuals and maximizing corporate productivity.

## 7.1 Future Work

The rationale behind the technology used in this study was to contribute to research in the field of work-life balance. However throughout the research, many additional novel applications of this model were discovered. It is hoped that once the simplifications have been reduced in the model that actuaries and insurance companies can use the model to tailor their health insurance policy premiums to different job types for individuals. Further research into the mitigation of the ill effects of stress on the body under certain controlled circumstances could establish a greater understanding of how the body reacts in a biological sense to stressors [18, 42, 49, 56]. If the agent's representation was increased to measure biological components such as interleukin-6 and epinephrine, the relationships between cognitive and physical performance could also be better understood. Through the maximization of eustress in individuals a maximally productive state from employees can be achieved and organizations can benefit from increased corporate productivity and ultimately profit. Additionally, employees could have a greater appreciation for work and enjoy their lives on a daily basis leading to more career fulfilment [20]. If the predictive model of stress were enhanced and tested against real world data sets, governments and corporations could use the model as a framework to test further policy changes to establish the social and economic impact before implementing the change. Moreover, the framework could be used by organizations to study possible means of maximizing eustress through initiatives. Finally, should the biological component be further established, the model could be used as a predictive test for the increased risk of developing certain medical and mental health conditions amongst employees. For example, an employee that logs a substantial number of work hours for prolonged periods of time has a greater propensity to develop medical conditions such as obesity or heart disease and could also experience the effects of depression. An advanced version of this model could potentially identify individuals at risk and management could intervene to reduce the stress loads of the individual to healthier levels.

### 7.1.1 Modifications and Simplifications

As there are many possible areas of enhancement to the model, some changes can be made to more accurately model stress in the workplace. Some of the interesting possibilities include:

- Stress decays by a fixed parameter over the weekends in the model however it does not model the possibility of agents having to work on weekends. It would be more accurate to allow agents to work on the weekend and still allow for rest.
- Allowing the model parameters to vary during the simulation would enable more complex interactions to take place. The value of individuals could be adjusted during the simulation to model new work requirements or life events.
- Currently covert drug use is tolerated within the organization. Allowing for agents to be fired if found to be using illicit compounds would more accurately depict reality.
- Incorporating different drug profiles would enable the study of effects of harmless stimulants compared to harmful illicit compounds. This would involve adding in the long term ill-effects of drug use possibly by increasing $\alpha$ values for the short term but substantially reducing them in the long term. Secondly, the productivity loss associated with long term drug use would need to be addressed.
- Having social and non-social rest periods would allow for the effects of spending time with colleagues, lunches out at work and meetings with one's boss as a possible way of advancing one's career. Furthermore, the more complicated effects of modelling rest at home could be modelled as not all time at home is restful.
- Introducing more than one special project that an agent could undertake would more accurately portray reality as there are many options for additional work one

could choose. To do this though there would need to be a minimum investment of time from an individual into a project to have the possibility of a pay-off.

- Using contact networks from epidemiology would allow for the modelling of networks within an organization. Each team would be a network and thus the effects of stress diffusing through a network could be established.

- Incorporating the notion of teams would allow individuals to share credit for work and allow for additional outlets for the diffusion of stress. Furthermore, agents in the group could additionally train one another to help further their careers together.

- Introducing possible policy changes in the model could allow for effects such as time off for being overworked thereby reducing the alpha value for individuals.

- Modelling the effect of social networks within a given organization would allow for the study of different standards for individuals. For example, an employee who is a personal friend of someone in management might have a different base productivity requirement and might also be assigned as a mentor in spite of their actual productivity. Consequently this would allow for the modelling of nepotism.

- Modelling prisoner's dilemma through cooperating (choosing to work) or defecting (resting all the time).

# Bibliography


[1] T. Akerstedt. Psychosocial stress and impaired sleep. *Scandinavian Journal of Work, Environment and Health*, 32(6):493–501, 2006.

[2] J. Arifovic. Evolutionary algorithms in macroeconomic models. *Macroeconomic Dynamics*, 4(3):373–414, 2000.

[3] D. Ashlock. *EVOLUTIONARY COMPUTATION FOR MODELING AND OPTIMIZATION*. Springer, New York, 2006.

[4] D. Ashlock and E. Knowles. Deck-based prisoner's dilemma. In *Computational Intelligence and Games (CIG), 2012 IEEE Conference on*, pages 17–24. IEEE, 2012.

[5] D. Ashlock and C. Lee. Agent-case embeddings for the analysis of evolved systems. *IEEE Transactions on Evolutionary Computation*, 17(2):227–240, 2013.

[6] D. Ashlock and M. Page. An agent based model of stress in the workplace. Accepted to SSCI 2013, 2013.

[7] D. Ashlock and M. Page. Binary decision automata modelling stress in the workplace. Accepted to CEC 2012, 2013.

[8] D. Ashlock and N. Rogers. A model of emotion in the prisoners dilemma. In *Computational Intelligence in Bioinformatics and Computational Biology, 2008. CIBCB'08. IEEE Symposium on*, pages 272–279. IEEE, 2008.

[9] W. Ashlock. Using very small population sizes in genetic programming. In *Evolutionary Computation, 2006. CEC 2006. IEEE Congress on*, pages 319–326. IEEE, 2006.

[10] K. Svrdsudd L. Welin-G. Tibblin L. Wilhelmsen P. Bjrntorp B. Larsson, J. Seidell. Obesity, adipose tissue distribution and health in menthe study of men born in 1913. *Elsevier Appetite*, 13(1):37–44, 1989.

[11] T. Back. *EVOLUTIONARY ALGORITHMS IN THEORY AND PRACTICE.* Oxford University Press, New York, 1996.

[12] J. Barsh and S. Cranston G. Lewis. *HOW REMARKABLE WOMEN LEAD: THE BREAKTHROUGH MODEL FOR WORK AND LIFE.* Crown Publishing Group, New York, NY, 2011.

[13] J. Biethahn and V. Nissen. Combinations of simulation and evolutionary algorithms in management science and economics. *Annals of Operations Research*, 52(4):181–208, 1994.

[14] JP. Bouchaud and M. Potters. *THEORY OF FINANCIAL RISK AND DERIVATIVE PRICING: FROM STATISTICAL PHYSICS TO RISK MANAGEMENT.* Cambridge University Press, 2003.

[15] G.E.P. Box. Evolutionary operation: A method for increasing industrial productivity. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 6(2):81–101, 1957.

[16] H.J. Bremermann and M.C.Yovits. *OPTIMIZATION THROUGH EVOLUTION AND RECOMBINATION, IN SELF-ORGANIZING SYSTEMS.* Spartan Books, 1962.

[17] C. Artemio C. Coello and G.B. Lamont. *APPLICATIONS OF MULTI-OBJECTIVE EVOLUTIONARY ALGORITHMS*, volume 1. World Scientific Publishing Company.

[18] S. Cartwright C. Cooper. An intervention strategy for workplace stress. *Journal of Psychosomatic Research*, 43(1):7–16, 1997.

[19] F.L. Cooke. Human resource strategy to improve organizational performance: a route for firms in britain? *International Journal of Management Reviews*, 3(4):321–339.

[20] M. Csikszentmihalyi and J. LeFevre. Optimal experience in work and leisure. *Journal of Personailty and Social Psychology*, 56(5):815–822, 1989.

[21] C. Kuusela D. Ashlock and N. Rogers. Hormonal systems for prisoners dilemma agents. In *Computational Intelligence and Games (CIG), 2011 IEEE Conference on*, pages 63–70. IEEE, 2011.

[22] C. McGuinness D. Ashlock and W. Ashlock. Representation in evolutionary computation. *Advances in Computational Intelligence Lecture Notes in Computer Science*, 7311:77–97, 2012.

[23] E.Y. Kim D. Ashlock and N.Leahy. Understanding representational sensitivity in the iterated prisoner's dilemma with fingerprints. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 36(4):464–475, 2006.

[24] D. Dasgupta and Z. Michalewicz. *EVOLUTIONARY ALGORITHMS IN ENGINEERING APPLICATIONS*. Springer Verlag.

[25] J.M. Deutsch. Evolutionary algorithms for finding optimal gene sets in microarray prediction. *Bioinformatics*, 19(1):45–52, 2003.

[26] A.E. Eiben and J.E. Smith. *AN INTRODUCTION TO EVOLUTIONARY COMPUTATION*. Springer, New York, 2003.

[27] J.M. Epstein. Agent-based computational models and generative social science. *Generative Social Science: Studies in Agent-Based Computational Modeling*, pages 4–46, 1999.

[28] P.J. Fleming and R.C. Purshouse. Evolutionary algorithms in control systems engineering: a survey. *Control engineering practice*, 10(11):1223–1241, 2002.

[29] D.B. Fogel. *EVOLUTIONARY COMPUTATION*. IEEE Press, 1995.

[30] G.B. Fogel and D.W. Corne. *EVOLUTIONARY COMPUTATION IN BIOINFORMATICS*. Morgan Kaufmann, San Francisco.

[31] L.J. Fogel. Autonomous automata. *Industrial Research*, 4(2):14–19, 1962.

[32] C.M. Fonseca and P.J. Fleming. An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary computation*, 3(1):1–16, 1995.

[33] Organisation for Economic Co-operation and Development. *BETTER LIFE INDEX*. OECD Publishing, 2011.

[34] Organisation for Economic Co-operation and Development. *THE OECD MENTAL HEALTH AND WORK PROJECT*. OECD Publishing, 2012.

[35] R.M. Friedberg. A learning machine: Part i. *IBM Journal of Research and Development*, 2(1):2–13, 1958.

[36] R.M. Friedberg. A learning machine: Part ii. *IBM Journal of Research and Development*, 3(7):282–287, 1959.

[37] C. Greeno and R. Wing. Stress-induced eating. *Psychological Bulletin*, 115(3):444–464, 1994.

[38] D. Guest. Perspectives on the study of work-life balance. *Social Science Information*, 41(2):255–279, 2002.

[39] D. Gusfield. *ALGORITHMS ON STRINGS, TREES AND SEQUENCES: COMPUTER SCIENCE AND COMPUTATIONAL BIOLOGY*. Cambridge University Press, 1997.

[40] H.W. Hethcote. The mathematics of infectious diseases. *SIAM review*, 42(4):599–653, 2000.

[41] J.H. Holland. Outline for a logical theory of adaptive systems. *Journal of the ACM*, 9(3):297–314, 1962.

[42] B. Hale K.L. Speckman H.R. Lieberman, W.J. Tharion and R. Tulley. Effects of caffeine, sleep loss, and stress on cognitive performance and mood during u.s. navy seal training. *Psychopharmacology*, 164(3):250–261, 2002.

[43] I. Rechenberg. *SIMULATIONSMETHODEN IN DER MEDIZIN UND BIOLOGIE*. Springer Berlin Heidelberg, 1978.

[44] S. Humphries V. Mohamed-Ali J. Yudkin, M. Kumari. Inflammation, obesity, stress and coronary heart disease: is interleukin-6 the link? *Elsevier Atherosclerosis*, 148(2):209–214, 2000.

[45] W. Kirsten. Health and productivity management in europe. *International Journal of Workplace Health Management*, 1(2):136–144, 2008.

[46] R. Leardi. Genetic algorithms in chemometrics and chemistry: a review. *Journal of chemometrics*, 15(7):559–569, 2001.

[47] K. Lindgren and M.G. Nordahl. Evolutionary dynamics of spatial games. *Physica D: Nonlinear Phenomena*, 75(1):292–309, 1994.

[48] F.G. Lobo and D.E. Goldberg. A review of adaptive population sizing schemes in genetic algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 228–234, 2005.

[49] G. Johansson M. Frankenhaeuser. Stress at work: psychobiological and psychosocial aspects. *Applied Psychology*, 35(3):287–299, 1986.

[50] S. Gregory M. LeFevre, G. Kolt and J. Matheny. Eustress, distress and their interpretation in primary and secondary occupational stress management interventions: which way first? *Journal of Managerial Psychology*, 21(6):547–565, 2006.

[51] D.E. Goldberg M. Pelikan and E. Canty-Paz. *BAYESIAN OPTIMIZATION ALGORITHM, POPULATION SIZING, AND TIME TO CONVERGENCE*. Illinois Genetic Algorithms Laboratory, 2000.

[52] P. McGovern C. Mills M. White, S. Hill and D. Smeaton. High-performance management practices, working hours and worklife balance. *British Journal of Industrial Relations*, 41(2):175–195, 2003.

[53] S. Martínez and C. Coello. A proposal to hybridize multi-objective evolutionary algorithms with non-gradient mathematical programming techniques. In *Parallel Problem Solving from Nature–PPSN X*, pages 837–846. Springer, 2008.

[54] K. Miettinen. *EVOLUTIONARY ALGORITHMS IN ENGINEERING AND COMPUTER SCIENCE: RECENT ADVANCES IN GENETIC ALGORITHMS, EVOLUTION STRATEGIES, EVOLUTIONARY PROGRAMMING, GE.* John Wiley & Sons, Inc., 1999.

[55] B. Kevin M.A. Summer-B. Seal M.R. Rosekind, K. Gregory and D. Lerner. The cost of poor sleep: Workplace productivity loss and associated costs. *Journal of Occupational and Environmental Medicine*, 52(1):91–98, 2010.

[56] R. Kessler E. Schilling N. Bolger, A. DeLongis. Effects of daily stress on negative mood. *Journal of Personality and Social Psychology*, 57(5):808–818, 1989.

[57] A. Narayanan and M. Moore. Quantum-inspired genetic algorithms. In *Evolutionary Computation, 1996., Proceedings of IEEE International Conference on*, pages 61–66. IEEE, 1996.

[58] Mental Health Commission of Canada. *CHANGING DIRECTIONS, CHANGING LIVES: THE MENTAL HEALTH STRATEGY FOR CANADA.* Mental Health Commission of Canada, 2012.

[59] World Health Organization. *THE WORLD HEALTH REPORT 2001 - MENTAL HEALTH: NEW UNDERSTANDING, NEW HOPE.* WHO Publishing, 2001.

[60] H. Pierreval and L. Tautou. Using evolutionary algorithms and simulation for the optimization of manufacturing systems. *IIE transactions*, 29(3):181–189, 1997.

[61] H. Timmerman R. Mannhold, H. Kubinyi and D.E. Clark. *EVOLUTIONARY ALGORITHMS IN MOLECULAR DESIGN.* Wiley-VCH, 2008.

[62] J. Robinson. The business case for wellbeing. *Gallup Business Journal*, 1(1):1–3, 2012.

[63] F. Rothlauf. *REPRESENTATIONS FOR GENETIC AND EVOLUTIONARY ALGORITHMS*. Springer, Netherlands, 2006.

[64] H. Friedman S. Booth-Kewley. Psychological predictors of heart disease: A quantitative review. *Psychological Bulletin*, 101(3):343–362, 1987.

[65] S. Ahmad S. Imtiaz. Impact of stress on employee productivity, performance and turnover; an important managerial issue. *International Review of Business Research Papers*, 5(4):468–477, 2009.

[66] M. Lipson S. Preble and H.Lipson. Two-dimensional photonic crystals designed by evolutionary algorithms. *Applied Physics Letters*, 86(6):061111–061111, 2005.

[67] A. Sanders. Towards a model of stress and human performance. *Elsevier Acta Psychologica*, 53(1):61–97, 1983.

[68] H.P. Schwefel. *EVOLUTIONSSTRATEGIE UND NUMERISCHE OPTIMIERUNG*. Technische Universitt Berlin, 1975.

[69] H. Selye. *STRESS WITHOUT DISTRESS*. J.B. Lippincott Company, Philadelphia, 1974.

[70] M.R. Gunnar S.J. Lupien, B.S. McEwen and C. Heim. Effects of stress throughout the lifespan on the brain, behaviour and cognition. *Nature Reviews Neuroscience*, 10:434–445, 2009.

[71] W. Staiano S.M. Marcora and V. Manning. Mental fatigue impairs physical performance in humans. *Journal of Applied Physiology*, 106:857–864, 2009.

[72] V. Seyfried M. Strehle T. Baumert, T. Brixner and G. Gerber. Femtosecond pulse shaping by an evolutionary algorithm with feedback. *Applied Physics B: Lasers and Optics*, 65(6):779–782, 1997.

[73] J. Harter T. Rath and K. James. *WELL-BEING: THE FIVE ESSENTIAL ELEMENTS*. Gallup Press.

[74] M. Tausig and R. Fenwick. Unbinding time: Alternate work schedules and work-life balance. *Journal of Family and Economic Issues*, 22(2):101–119, 2001.

[75] D. Fogel T.Back and Z. Michalewicz. *EVOLUTIONARY COMPUTATION 1: BASIC ALGORITHMS AND OPERATIONS*. Institute of Physics Publishing, Bristol, United Kingdom, 2000.

[76] KY. Tsai and FS. Wang. Evolutionary optimization with data co-location for reverse engineering of biological networks. *Bioinformatics*, 21(7):1180–1188, 2005.

[77] H. Weigand V. Dignum and L. Xu. Agent societies: towards frameworks-based design. In *Agent-Oriented Software Engineering II*, pages 33–49. Springer, 2002.

[78] C.P. Katsaras V.K. Koumousis. A saw-tooth genetic algorithm combining the effects of variable population size and reinitialization to enhance performance. *IEEE Transactions on Evolutionary Computation*, 10(1):19–28, 2006.

[79] R.E. Keller W. Banzhaf, P. Nordin and F.D. Francone. *GENETIC PROGRAMMING: AN INTRODUCTION*. Morgan Kaufmann Publishers.

[80] R.M. Yerkes and J.D.Dodson. The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology*, 18:459–482, 1908.

[81] R. LeRiche Z. Michalewicz, D. Dasgupta and M. Schoenauer. Evolutionary algorithms for constrained engineering problems. *Computers & Industrial Engineering*, 30(4):851–870, 1996.