MODELING AFFECT-BASED RELATIONSHIPS AND SOCIAL ALLOSTASIS AS MECHANISMS FOR ADAPTATION IN SOCIAL GROUPS

By

IMRAN M. KHAN

A dissertation submitted in partial fulfillment of
the requirements for the degree of

DOCTOR OF PHILOSOPHY

UNIVERSITY OF HERTFORDSHIRE
School of Engineering and Computer Science

SEPTEMBER 2020

ACKNOWLEDGMENT

I would like to express my gratitude to my principal supervisor, Dr. Lola Cañamero, for allowing me to undertake my Ph.D. research at the EECAiA lab under her supervision. Thank you for the years of advice, guidance and patience; and for investing the time and energy into training me to become a better researcher.

I would also like to thank my second supervisor, Dr. Matthew Lewis, for his supervision and willingness to always assist with a host of seemingly-random issues over the years.

Finally, I would like to thank the University of Hertfordshire, for awarding me a Ph.D. studentship that granted me the opportunity to complete my research programme full-time.

Abstract

Artificial social agents are becoming increasingly present in the real world. For these agents, their success is tied to their ability to remain truly autonomous—to appropriately adapt their decisions and behaviours across changing physical and social contexts—and still remains a key area of research in the field of embodied autonomous agents. Numerous approaches towards models of adaptation have been proposed, including biologically-inspired models based on the regulation of an agent's internal environment. A more recent concept from biological systems, called "(social) allostasis", proposes a mechanism for the anticipatory adaptation of the internal environment through (affect-based) interactions with the (social) environment. Despite the long-term adaptive benefits proposed by social allostasis, current approaches towards models of adaptation for artificial social agents are yet to consider its principles in their approaches. In this manuscript, we address this shortcoming and investigate how mechanisms of social allostasis through affect-based relationships can be used for long-term adaptation of an embodied agent model. Using a biologically-inspired artificial life approach, we systematically investigate numerous hypothesised, hormonal mechanisms that underpin social allostasis and the survival-related benefits associated with affect-based social relationships. We conduct these investigations using a small society of simulated agents across numerous dynamic physical and social contexts. Our results find that adaptation of an embodied agent model through social allostasis mechanisms can provide survival-related advantages across several dynamic conditions in some contexts through physiological and behavioural adaptation. Throughout the manuscript, we also find numerous social interactions and dynamics that mirror biological systems. Our work addresses limitations in current approaches to adaptive models for embodied artificial agents, and presents a novel framework to guide future work in this field. The work also contributes a biologically-inspired artificial life model as a scientific tool to address some methodological challenges in natural systems.

TABLE OF CONTENTS

Page

APPENDIX

LIST OF TABLES

LIST OF FIGURES

To my parents, who have sacrificed more for me than I will ever know.

To my family and friends, for their endless support throughout the years.

Finally, to my brother, Mebz, for the lifelong memories and making me the person I am today.

Rest in peace.

Chapter One

Introduction

1.1 Introduction

Recent years have seen a rise of social agents in the real world [1]—from applications in health care [2]–[4], education [5]–[9], communication and use in public spaces [10], [11]—and are only likely to become more widespread in the near future. One aspect of consideration for the long-term success of these agents in the real world is their ability to remain autonomous—to use their own (cognitive) processes in order to regulate their interactions with their (changing) environment appropriately [12]—in a way that allows them to remain viable, long-term solutions. Given the dynamic nature of the real world, autonomous regulation requires the ability to adapt to, and even anticipate, constantly changing conditions. How to approach this adaptation for social agents is a non-trivial challenge that still remains a key area of research in the field of embodied artificial intelligence.

Looking to the natural world, autonomous regulation—and therefore long-term survival—has been underpinned by continuous adaptation to ongoing changes in the physical and social environments [13], [14]. For social species (including humans), this adaptation is facilitated, in part, through the formation and maintenance of socially-affective relationships [15]–[17]. Individuals with affect-based social support have consistently been found to benefit from improvements in health and longevity of life [18]–[20]. When developing approaches to long-term adaptation in social environments, the adaptive potential proposed by these socially-affective relationships should be considered.

One long-standing concept suggested to be responsible for long-term adaptation is that of homeostasis [21]: that stability (i.e. survival) of an organism is driven by the maintenance of its essential parameters within an ideal range or "safe zone". The "viability" of an organism or system, according to Ashby

[22] describes the ongoing maintenance of these essential parameters—within a range that promotes system stability—through behavioural and physiological regulation and adaptation. Numerous approaches towards adaptation have been proposed using homeostatically-controlled models (Section 1.3). In recent years, however, the notion of homeostatic processes being solely responsible for long-term survival has been reconsidered [14], [23], [24]. Rather than the depletion-replenishment approach of Cannon's definition of homeostasis [21], the concept of allostasis [14], [25], [26] proposes a more anticipatory, adaptive mechanism of regulation.

Allostasis proposes that physiological processes—including the set points or ranges of homeostatically-controlled parameters—can be adjusted when faced with, and in anticipation of, changing internal and external environmental conditions, by accounting for prior information and the current internal and external stimuli. Unlike traditional homeostatic systems, where stability is achieved by correcting deviations of a physiological parameter from a static set point, allostasis describes a mechanism that anticipates environmental changes and dynamically adjusts the homeostatically-controlled set point before the need to correct (large) deviations occurs. We illustrate the difference between homeostasis and allostasis with an example. Consider a biological agent A (i.e. an animal) that is sleeping and must maintain its blood pressure (a physiological parameter) within a given range of values. When A wakes, her blood pressure rises accordingly. In a strictly-homeostatically-controlled system, this change would result in an error signal (the deviation of the present blood pressure value from the static range of "safe" values) which threatens the stability of the system and requires correction through physiological or behavioural adjustments. Allostasis anticipates the changing environment (i.e. from sleeping to not sleeping) and adjusts the "safe range" of blood pressure accordingly (to relatively higher levels) in order to maintain stability of the system. Therefore, this anticipatory adaptation of physiological processes, driven by allostasis, can be seen as an adaptive mechanism that adapts the adaptive mechanism (of homeostasis) and proposes long-term stability across dynamic environments.

By extension, social allostasis [14] proposes that this anticipatory adaptation of physiological processes is driven by interactions with the social environment. Such a re-conceptualisation of homeostatic regulation challenges classic homeostatic theory—though not without criticism [23], [27]—and provides a biologically-plausible mechanism for how individual in a society adapt their physiology and behaviours to constantly changing physiological, physical and social environments. Despite recent criticisms that the definition of homeostasis sufficiently describes the physiological adaptation proposed by allostasis [27], [28], we use the term "(social) allostasis" to describe a mechanism that is complementary to homeostasis [29] and which adapts the homeostatically-controlled system through (several) anticipatory physiological and behavioural changes.

Underpinning this socially-driven adaptation is the formation and maintenance of affect-based social relationships, such as parent-offspring [14], [30] and close-kin relationships [31]. These affect-based relationships may provide mechanisms through which individuals are able to adapt their physiology and behaviours, permitting long-term survival across a range of conditions. The mechanisms underpinning the formation and maintenence of these social relationships, along with their adaptive functions on behaviour and physiology, however, are still unclear.

Despite its biological plausibility, its relevance to the Ashbyan framework of "ultrastability"—the ability for a system to maximise stability by adjusting its internal environment in response to on-going changes in environmental conditions [32]—and suggestions that models of adaptation should consider allostasis as part of their approach [33], approaches explicitly using (social) allostasis principles in their models of adaptation still appear to be largely unexplored. Where such agents are situated in the world that they are (socially) interacting with in real time, making context-relevant, socially-affective decisions can be seen as an essential component of successful decision-making, which may contribute to their long-term success. These mechanisms of social allostasis, facilitated by social interactions and relationships, may provide a framework to guide approaches towards developing models for long-term adaptation in a social environment.

The work presented within this manuscript aims to address this limitation in the present literature by approaching adaptation of embodied agent models using these principles of affect-based relationships and social allostasis in the contexts of small social groups. We systematically investigate some of the underlying mechanisms that are hypothesised to play a role in the physiological and behavioural adaptation through affect-based social relationships, inspired by the principles of social allostasis. We focus our attention on two specific biological hormones—oxytocin and cortisol—and investigate a number of their hypothesised, adaptive functions within social relationships and physiological and behavioural adaptation. Through the manuscript, we move from an agent model driven by strictly-homeostatic control to a socially-adaptive model underpinned by affect-based social relationships and mechanisms of social allostasis.

Despite being widely studied, the specific effects of the hormones we investigate still lack universal consensus, presenting an additional challenge to our research. As we will see throughout this manuscript, we also use our model as a tool to investigate, validate, and even challenge some of the contrasting, often-contradictory hypotheses of these hormones from biological literature.

As we find in later chapters, physiological and behavioural adaptation through social allostasis-inspired mechanisms promotes the performance of our model across a number of dynamic, challenging physical and social conditions. We suggest that such an approach towards adaptation provides significant advantages with respect to the long-term viability of embodied agent models, and we suggest that future work builds on our framework and continues to explore socially-affective relationships and social allostasis as mechanisms for adaptation in their approaches.

1.2 Aims & Hypothesis

In this manuscript, we aim to investigate how the adaptation of a homeostatically-controlled embodied agent model through mechanisms of affect-based relationships—and inspired by social allostasis principles—affects the performance (in terms of "viability", Section 2.4.1) and social dynamics of a small society of artificial agents.

We will systematically investigate several hypotheses motivated by biological mechanisms that underpin affect-based social interactions and relationships, and evaluate their effects on the viability of our agent model across several dynamic physical and social environments.

Our research is therefore guided by the following research question:

How does the adaptation of a homeostatically-controlled embodied agent model through affect-based social relationships and social allostasis mechanisms affect the viability and social dynamics of a society of artificial agents?

To address this research question and focus our investigation throughout the manuscript, we ask the following questions:

- How does oxytocin's effects on increasing or decreasing the attention given to social cues affect group wellbeing and interactions? (Chapter 3)
- Does oxytocin-mediated adaptation of internal satisfaction signals associated with tactile interaction affect group viability and interactions? (Chapter 4)
- How does oxytocin-mediated social bond partner preference for social interactions affect the wellbeing of agents with social bonds in a rank-based society? (Chapter 5)
- How do different, rank-related tolerances to stress affect the social dynamics and agent viability in a rank-based society? (Chapter 6)
- How do the effects of social support on reducing the release of a stress hormone affect social dynamics and the viability of agents with and without social support? (Chapter 7)

- How does the adaptation of tolerances to stress, through social support and oxytocin mechanisms, affect the viability and social interactions of agents? (Chapter 8)
- Does the adaptation of energy metabolism, through the stress-reducing effects of social support, improve agent wellbeing and reduce socio-negative interactions for agents with social support? (Chapter 9)

We aim to use our results to develop and contribute a novel framework for a socially-adaptive model for embodied artificial agents built on the mechanisms that we systematically investigate and evaluate.

We aim to investigate several existing hypotheses from biology, with respect to hormonal effects and social dynamics, by using our biologically-inspired model as a tool to address some methodological challenges of studies in the natural world. We therefore aim to contribute the results of our investigations to inspire future work for those fields.

We also aim to contribute our artificial life model as an experimental tool for hypothesis generation and validation, and promote future cross-disciplinary research.

Based on existing literature from natural systems and current models of adaptation for artificial agents, we hypothesise that mechanisms of adaptation, underpinned by affect-based relationships and the principles of social allostasis, will result in significant improvements to the performance—with respect to viability—of a homeostatically-controlled agent model across a number of challenging social and physical conditions.

1.3 State of the Art

We take a cross-disciplinary approach to our investigation and are therefore motivated by the work and limitations from a number of different fields. Here, we discuss related and current approaches of adaptive models using (hormonally-modulated) homeostatic models. We also discuss current work related to models of social interaction and adaptation in simulation models, and approaches towards modelling affect in multi-agent social systems. We also highlight some related approaches towards adaptation using principles of allostasis.

1.3.1 Hormonal Modulation of Homeostatically-Controlled Agent Models

The use of "hormones" as a form of modulation in embodied agent models is not a new area of research. On the contrary, a significant amount of work has been undertaken in this area to investigate how such mechanisms can provide more adaptive decision-making and behaviours in autonomous agents. Inspired by cybernetics, earlier work [34] proposed an action-selection architecture that used hormonal mechanisms, secreted as a function of “affective” states, to modulate the intensity of motivations and internal physiological signals which then influenced behaviour selection. Building on this, [35] used hormonal modulation in their homeostatically-driven model to modulate the perception of external cues, and [36] investigated the neuromodulation of a neural-network-based Braitenberg vehicle, modulating internal thresholds related to a robot’s “energy”. Later studies built upon this to modulate perceived internal errors [37] and internal physiological variables [38]. In those investigations, the researchers demonstrated adaptive behaviours and viability benefits through hormonal modulation of their decision-making models. More recently, [39] investigated the effects of ‘pleasure-like” hormonal mechanisms in a homeostatically-controlled agent model, and [40], [41] investigated the effects of several hormones to adapt the strength of sensory input and the speed of agent movement—as well as the sensitivity of “hormone receptors”—in an epigenetic-like model of hormonal modulation. Additional investigations [42]–[45] have also been inspired by similar neuromodulatory functions for adaptation of their homeostatically-driven controllers, showing adaptive benefits of these mechanisms in physical agents.

Additional models of neural network-based hormonal adaptation have been proposed by [46], [47], with [48], [49] using these mechanisms to create a model of “homeostatic plasticity” in their models. More recent approaches have also considered the social environment as part of their investigations. Models of homeostatic hormonal control have been used in social learning in robot-animal colonies [50]

1.3.2 Models of Adaptation through (Affect-Based) Social Interactions

Models of adaptation through social interaction have also been proposed, though largely concerning themselves with dyadic interactions and not group-level interaction. The study in [51] used a probabilistic model in a simulated human-robot simulation to compare the performance of strictly-reactive and proactive models during a cooperation task. The authors evaluated how a proactive model could adapt to learning behavioural intent of the (simulated) human, and showed significant performance benefits of using a proactive approach in their agent’s decision-making model. Similar performance benefits resulting from anticipatory social action have been reported by [52]–[54] who modelled mechanisms of adaptation—albeit using a computationally-intensive approach—to determine optimal behavioural strategies in physical agents. Behaviour and attitudes of humans towards virtual agents [55], [56] have also been seen to improve when the agent would adapt its behaviour through affective social interaction. The maintenance of social habits, facilitated by embodied interactions with other social individuals, has also been demonstrated to improve the performance of homeostatically-controlled agent models [57].

Models of social learning have also been proposed as a mechanism for adaptation for social agents. Using neural network-based simulation models, [58], [59] showed how behavioural adaptation through social learning resulted in improved task-related performance for their agent model. More recent approaches towards social learning have shown improvements in agent performance compared to individual learning in both artificial [60]–[63] and physical agents [64]. Current approaches find how social learning is maximised particularly when multiple social partners are used to learn from, and when “teachers” determine what is being learned [65], as well as the emergence of group “leadership” in a small collective of social agents [66].

In sum, researchers have shown task-specific improvements in the performance of their models when artificial agents adapt behaviours through their social environment and social interactions using a range of approaches.

Investigations using large multi-agent systems have considered the relationship between small-scale interactions and wider group dynamics [67]–[71]. Biologically-inspired embodied models, particularly those modelled on primate societies [72]–[77], have been used to study emergent group dynamics primarily through dominance-based interactions. Conversely, recent biologically-inspired models [78]–[81] have sought to understand the role of cooperative social interactions on group dynamics.

Social dynamics have also been investigated through affect-based decision-making models. For instance, [82], [83] used discrete models of emotion, independent of a cognitive decision-making architecture, to dynamically adapt agent behaviour in their evacuation-related models. An alternative approach towards emotion modelling, based on cognitive appraisal theory, has been used by [84], [85], whose modelling of emotion also accounts for internal cognition, drives, and personality traits of individual agents. A similar approach was used by [86], [87]: using different intensities of predetermined emotional states to determine an agent’s “mood” to adapt decision-making. More recent work has included approaches using multi-dimensional models of personality and emotion. This includes both agent–agent, [88]–[91] and multi–agent [92] interactions. Additional studies have investigated models of (socially–)affective decision–making, focusing on human–agent interactions [93]–[95]. More recently, [96] demonstrated how a triadic awareness of resource-sharing (altruism) shaped future group interactions and dynamics through indirect reciprocity of behaviours. However, many of these approaches suffer from several limitations which we aim to address in our investigation. We discuss these in Section 1.4.

1.3.3 Current Approaches of Allostasis

Despite the increasing discussion of allostasis in the field of autonomous agents [12], [97]–[99] and the notion that allostasis aligns with Ashby's [22] discussion of "ultrastable" systems [100], the development of allostasis-inspired models still appear to be somewhat limited. Nevertheless, some investigations have been undertaken, using allostatic mechanisms to adapt homeostatically-driven systems, both with and without hormonal mechanisms.

Using a probabilistic approach in their physical agents, [101] used an allostatic controller to control numerous homeostatic subsystems which concerned themselves with foraging behaviour. They found that such an approach created behavioural stability, demonstrating behavioural similarities between their physical agent and a rodent model. More recently, [102] used an allostatic control layer as a second-layer above a homeostatic system for a robot during social human–robot interactions. Similarly, the model proposed by [103] proposes an allostatic-type model by using a second-order adaptive layer above predefined sensorimotor control loops. Finally, the work undertaken by [40] also constitutes a process of allostasis. In those investigations, the epigenetic mechanisms adapted the internal physiology of the agents over time, resulting in improved performance during challenging conditions.

In sum, models that either use, or are explicitly inspired by, allostatic mechanisms have been limited in the current literature. As a result, models using social allostasis principles are yet to be proposed.

1.4 Motivations for the Research

Overall, numerous approaches towards models of adaptation, including hormonally-modulated homeostatic systems, have shown performance benefits across a range of dynamic world conditions. Despite some early work using allostatic mechanisms for adaptive models [40], [101], [102], much more work needs to be undertaken. By extension, models related to social allostatic principles—that of the adaptation of the internal environment via the (external) social environment—are yet to be systematically investigated. Despite researchers discussing the relevance of their adaptive models for (physical) social agents, the relationship between micro-level social interaction and wider social dynamics have rarely been the focus of investigation.

As we mentioned in Section 1.1, the re-evaluation of homeostatic mechanisms in the form of social allostasis [14], [26] provides researchers of embodied agent models with a biologically-plausible framework that can be used to better conceptualise and develop potential models of long-term adaptation for social agents. Any model that concerns itself with adaptation in a social context must also consider their social milieu as part of their adaptation [14]. We are therefore motivated by this limitation in the current literature.

Potentially as a result of suffering from both the "Scaling" and "Dynamics" problems [104], a large focus of current affect-based social models of adaptation—including social simulation models—have concerned themselves either with dyadic interactions, which makes it difficult to extrapolate findings to larger social groups, or the modelling of large-scale group behaviours, that overlook the micro-level phenomena that underpins aggregated behaviours (Section 1.3.2). Though this appears to be a significant challenge to overcome in the field of agent-based modelling [104], social models concerning themselves with understanding the emergence of social dynamics should also allow for analysis at the micro-level of individual behaviours. Our methodology is motivated by the limitations of these current models to adequately investigate both the micro and macro-level phenomena: to understand how group dynamics can potentially emerge as a result of smaller-scale interactions.

Many of the approaches towards adaptation of homeostatically-controlled models [40], [46]–[49], [103], [105] as well as mechanisms for social adaptation [53], [54], [64], [65] have been developed using artificial neural networks for both artificial and physical agents. In our view, such approaches are computationally-costly with application that is limited by resource availability. As part of our motivation in developing a socially-adaptive model inspired by social allostasis principles, we are also motivated to simplify the computational complexities of current adaptive models: to not let such a model be restricted by computational or hardware limitations, in order to facilitate potential future investigations and application.

Finally, though not immediately relevant to the field of socially-affective adaptive models, we are also motivated by the non-conclusive findings of hormonal mechanisms presented in the ethological and neuroscience literature [106]–[111]. There are a number of methodological challenges that exist when attributing causality to hormones (including oxytocin and cortisol) in the real-world: including temporal delays of hormone release [109], [112], the relative stress associated with subject analysis [113], and inconsistencies between exogenous treatments and endogenous effects [111]. Therefore, abstracting some of these proposed effects and investigating them with a model removed from the natural-world challenges may allow us to contribute to this divided field of research, and to allow our model to be used as a tool to validate and generate hypotheses for biological systems.

1.5 Overview of Dissertation

Our research concerns itself with systematically investigating the adaptive effects of some biologically-inspired mechanisms underpinning affect-based relationships and social allostasis, and to develop an embodied agent model of adaptation based on such mechanisms for social agents. To this end, we have focused our research on two hormones—oxytocin and cortisol—and, through a combination of our own results and observations from the natural world, have tested several hypotheses of these hormonal mechanisms in our model while systematically developing our model using these mechanisms. At the start of each investigation chapter (Chapters 3 to 9), we provide the relevant background literature to justify the mechanisms and hypotheses that have been investigated.

Our investigation can be considered in two parts. In the first part (chapters 3 to 5) we investigate a number of hypotheses related to the effects of oxytocin on the attention given to social stimuli (the "social salience" hypothesis). In the second part of this dissertation (chapters 6 to 9), our focus shifts towards the stress-reducing effects of oxytocin and social support (the "social buffering" hypothesis).

Throughout the chapters, we incrementally build upon our agent model, using the findings from each chapter to generate further hypotheses while developing our model of adaptation using mechanisms inspired by social allostasis. Our final model is presented in Chapters 9.

This first chapter (**Chapter 1**) provides the relevant background and related work, along with our motivations and aims for the undertaken research.

Chapter 2 gives an overview of the basic agent model and experimental methodology used throughout this investigation. We describe our initial agent model, our experimental design and methodology, and provide justification of methods where appropriate.

Our investigations start at **Chapter 3**, which describes our initial study into the effects of oxytocin on the processing of social cues. In this chapter, we investigate the effects of two contrasting hypotheses (the "prosocial" and "interactionist" approaches, respectively) regarding oxytocin's effects on the attention given to social stimuli, and test these contrasting approaches on the performance of our agent model.

Chapter 4 builds upon our results from Chapter 3, which showed some support for the "prosocial" approach of oxytocin mechanisms. From the results, we concluded that further contextual information needed to be considered as part of our investigation. Continuing with the investigation into oxytocin's contrasting effects on social salience, we further test these hypothesised effect when one additional hypothesised role of the hormone—related to the internal physiological signal associated with prosocial behaviours—was considered as part of our investigation.

Contradicting the initial results, the results from Chapter 4 found that the prosocial approach did not always apply across all contexts. Instead, oxytocin's effects on social salience may be context-dependent: supporting the "interactionist" hypothesis. Based on our observations, we suggested that these contextual effects may be dependent upon the social environment, though our investigations lacked an appropriate social context. In **Chapter 5**, we consider this social context in the form of a social hierarchy and the presence of bond partners. In this social environment, we investigate a hypothesised contextual effect of oxytocin on social cues: that it specifically enhances the valence of social bond partners for social interactions.

These results showed further support for the interactionist approach of oxytocin, and we observed how different types of social bonds resulted in the emergence of a number of different affective behaviours. Based on both our qualitative and quantitative findings, we hypothesised that the wider social dynamics of a small society may be determined by the interactions and behaviours of dominant individuals. We considered that a rank-based society can also be a source of stress for some individuals, and that social support can create a "buffer" against the potentially-negative consequences of stressful conditions. **Chapter 6** introduces our initial work into investigating this "social buffering" hypothesis. Based on our results from the previous chapter and additional biological literature, this first study investigated the role of individual, rank-related "personality" differences—related to stress tolerance—on social interactions and group dynamics.

The investigation in Chapter 6 showed how strictly-rank-based social interactions and the wider group dynamics were affected by a combination of individual stress tolerances and the physical challenges presented by the environment. These results showed some contextual support for our hypothesis. In **Chapter 7**, we continue the investigation into the "social buffering" hypothesis by accounting for the presence of affective social bonds in the social environment. We investigate the effects of one hypothesised stress-reducing role of social support on both individual and group dynamics and wellbeing, across a number of different physical and social contexts.

Contrary to our hypothesis, we found that the stress-reducing effect of social support was dependent upon the environmental and social context, and that the adaptive benefits of social support may be limited by these contexts when only considering one of the "social buffering" effects. Building on earlier results from Chapter 6, **Chapter 8** investigates a second hypothesised effect of social buffering: that affect-based social relationships also "buffer" the tolerance to stress in individuals. We investigate how accounting for two of these hypothesised effects of social support affect social interactions and group viability.

The findings from Chapter 8 found how accounting for two of the stress-reducing effects of social support resulted in different types of physiological and behavioural adaptation, dependent upon the degree of the environmental challenge. In **Chapter 9**, we consider a third adaptive effect of these affect-based social relationships: that they not only regulate the stress response, but that they also provide a more direct physiological effect by reducing the "cost" associated with stress for socially-supported individuals. Using our previous results and addressing some methodological challenges from the natural world, we propose a potential mechanism through which social support directly adapts one type of physiological cost, related to energy metabolism. We investigate the effects of this adaptive mechanism on both individual and group viability, and the wider social interactions.

Chapter 10 summarises all of the work completed and discusses several areas of consideration for future work.

1.6 Contributions of the Research

The research undertaken makes the following contributions:

- An embodied agent model of adaptation through mechanisms of affect-based relationships, inspired by social allostasis principles. We systematically investigate a number of adaptive, hormonal mechanisms that underpin affect-based social interactions, and build upon a homeostatically-controlled action-selection model to develop a model of social adaptation inspired by social allostasis. We demonstrate improved performance of our final model across a number of social and physical conditions (Chapter 9).
- A framework for modelling and quantifying the quality of affective and dyadic relationships for artificial agents. Capitalising on the capabilities of artificial systems, we have proposed a dynamic, real-time measure of (affective) bond strength between dyads or small groups, called the Dyadic Strength Index (DSI). This addresses a number of limitations from previously-used measures of social relationships in agent models (Chapter 5) and contributes a quantifiable model and measure of affect-based relationships for future investigations.
- A mathematical model of the hypothesised "trade-off" between social partner selection in biological agents, responsible for the context-dependent adaptation of social behaviours (Social Assessment Component, Chapter 5). The modular implementation of the Social Assessment Component allows it to be modified and integrated within future socially-adaptive embodied agent models.
- The contribution of our observations and generation of a number of hypotheses to inspire future cross-disciplinary work. Throughout the research, our biologically-inspired agent model has been used as a tool to both support and challenge existing hypotheses from biological literature. We summarise these findings in Chapter 10.

- An open-source, biologically-inspired social simulation model, documented to the Overview, Design concepts, and Details (ODD) Protocol, to provide ease of replicability and to be used as a scientific tool to conduct future cross-disciplinary investigations.

This Ph. D. has contributed several journal articles and conference proceedings:

- Khan, I., Lewis, M. and Cañamero, L., 2018, April. Adaptation and the social salience hypothesis of oxytocin: Early experiments in a simulated agent environment. *In Proceedings of the 2nd Symposium on Social Interactions in Complex Intelligent Systems (SICIS)*, Liverpool, UK (pp. 4-6). (Chapter 3)

- Khan, I. and Cañamero, L., 2018. Modelling adaptation through social allostasis: Modulating the effects of social touch with oxytocin in embodied agents. *Emotions in Robots: Embodied Interaction in Social and Non-Social Environments, MTI Special Issue*, 2(4), p.67. (Chapter 4)

- Khan, I., Lewis, M. and Cañamero, L., 2019, September. The Effects of Affective Social Bonds on the Interactions and Survival of Simulated Agents. In *2019 8th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)* (pp. 374-380). IEEE. (Chapter 5)

- Khan, I., Lewis, M. and Cañamero, L., 2020. The Effects of Rank-Related Differences in Stress Tolerance on the Social Dynamics and Group Viability of Artificial Affective Agents. *IEEE Transactions on Affective Computing* (Chapter 6, submitted).

- Khan, I., Lewis, M. and Cañamero, L., 2020, July. Modelling the Social Buffering Hypothesis in an Artificial Life Environment. *In Artificial Life Conference Proceedings (pp. 393-401).* MIT Press. (Chapter 7)

- Khan, I., and Cañamero, L., 2021. The Long-Term Efficacy of “Social Buffering” in Artificial Social Agents: Contextual Affective Perception Matters. *Frontiers in Human-Robot Interactions (Special Issue: Affective Shared Perception)* (Chapter 8)

Chapter Two

Simulation & Agent Model

2.1 A Biologically-Inspired Artificial Life Model

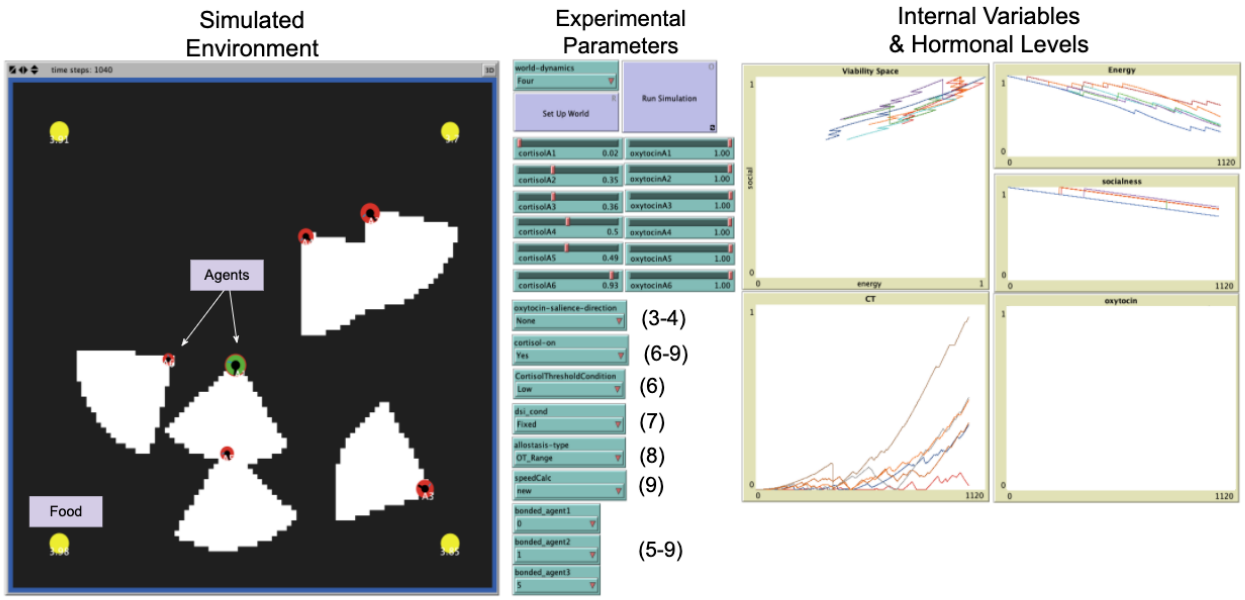

world-dynamics: Four
cortisolA1 0.02
oxytocinA1 1.00
cortisolA2 0.35
oxytocinA2 1.00
cortisolA3 0.36
oxytocinA3 1.00
cortisolA4 0.5
oxytocinA4 1.00
cortisolA5 0.49
oxytocinA5 1.00
cortisolA6 0.93
oxytocinA6 1.00
oxytocin-salience-direction: None
cortisol-on: Yes
CortisolThresholdCondition: Low
dsi_cond: Fixed
allostasis-type: OT_Range
speedCalc: new
bonded_agent1: 0
bonded_agent2: 1
bonded_agent3: 5

Figure 2.1 Screenshot of the graphical user interface (GUI) of the simulated model developed in NetLogo. Left: The simulated world environment, showing agents (doughnut-shapes) and their field-of-vision (white), and food resources (yellow discs). Middle: Experimental controls and parameter settings. Numbers alongside each parameter denotes the chapter each parameter was introduced. Right: Plots of internal states, variables and hormone levels, updated in real-time as experiments are being conducted.

In this chapter, we introduce the basic overview of our simulated environment and agent model used throughout our investigation, along with a brief description of the basic experimental design. We also provide justification for our methodology choices where appropriate.

We use a simulated, artificial life approach throughout the course of our research, developed using the NetLogo [114] platform, version 5.3.1. We propose advantages to using an artificial life approach over physical agents for reasons related to scalability, cost, resource requirement, noise reduction during data capture, minimisation of complexity, and accessibility. For ease of reading, we extend on these points in Appendix C.

The nature of the mechanisms that we are looking to investigate—that of hormones and social adaptation—has roots in several disciplines, including neuroscience [25], [26], [29], developmental biology [13] and ethology [115]–[117]. To develop a plausible model, we must consider their literature as part of our own investigation. Researchers of biologically-inspired (henceforth bio-inspired) artificial life models are able to deduce mechanisms and behaviours from real-world observations, abstract them into simple computational models; embed them within artificial agents, and use the observation of the emergent behaviours to validate ethological behaviour or potentially propose alternative hypotheses of such behaviours in the real world.

We use this approach for our investigations, and our agent model may then be likened to that of artificial animals. This approach has traditionally been referred to as the "animats" approach of modelling, though "ethorobotics" has been coined more recently in an attempt to formalise the multidisciplinary approach of evolutionary, ecological, and ethological concepts towards the development of social (physical and artificial) agents [118], [119]. In each chapter, we aim to provide relevant background literature to justify our abstractions of mechanisms and hypotheses that we investigate.

However, despite the proposed benefits towards the study of these fields, the implications of extrapolating our findings to biological agents has limitations. Though similarities can be deduced from emergent behaviours with an explanation of causal mechanisms, we recognise that such modelling does not accurately capture the nuanced complexities of real-world behaviour or physiology, especially that of human societies. The utility of simulations for this purpose is a particularly complex issue [120] and is not a challenge that we look to specifically address in this research. In defence of our bio-inspired approach, we suggest that an artificial life model may address some of the limited conclusions drawn from real-world studies, and can be used as a scientific tool in its own right to contribute to hypotheses in those fields.

Despite the proposal that artificial agents should have biological targets [121], our modelling approach did not have any specific biological target in mind in our earlier investigations (Chapters 3 and 4). Later investigations (Chapter 5 onwards), however, were partially inspired by some physiological and behavioural observations from primate (specifically chimpanzee) studies. We discuss the reasons for this in Section 5.2.1.

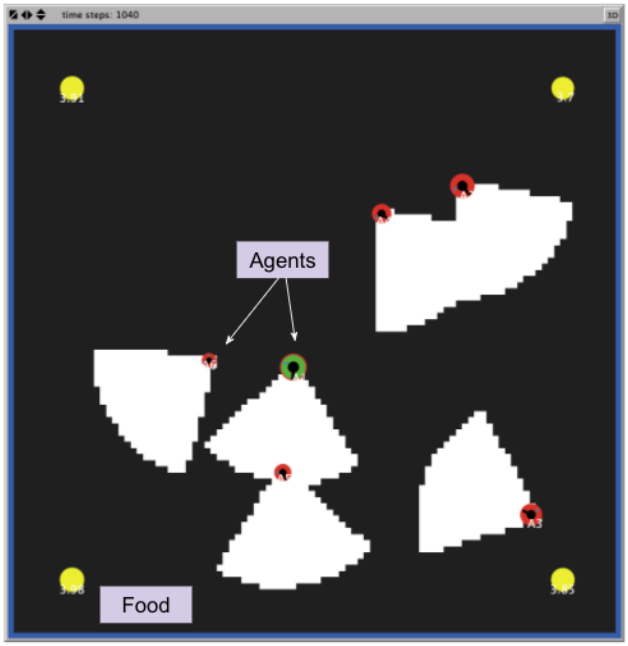

Figure 2.2 A screenshot of the two-dimensional simulation model used in this investigation (screenshot taken from experiments in Chapter 9). Both types of resources are labelled: agents are doughnut-shaped discs (red or green), and food resources are yellow spheres. Food distribution varies between world conditions and chapters. White patches show an agent's field-of-vision and remains unchanged throughout all investigations.

In those chapters, we identify and describe similarities between the behaviours and dynamics of our artificial agents and those of primate (chimpanzee) societies. As a result, we conclude (in Chapter 10) that the agents in these later chapters can be considered "artificial primates", and discuss the potential of using our findings to contribute explanations of social phenomena in (human and non-human) primate societies.

2.2 Simulation Environment

We use an artificial life approach throughout the entirety of our investigation. This model was developed using the NetLogo platform, version 5.3.1 [114]. NetLogo is a multi-agent modelling environment and was selected as an appropriate development platform based on previous work from similar investigations [72], [73]. All investigations undertaken used the same version of the NetLogo platform. Throughout the research, the complexity of the environment (with respect to resource availability) was varied. Changes to these base conditions are discussed in each respective chapter.

The simulated environment was set up as a two-dimensional square environment of size 99 by 99 patches. In the NetLogo environment, a patch constitutes one square or unit on the grid structure. Each environment contains two different types of objects or resources: autonomous **Agents** whose behaviour is driven by a motivation-based controller (Section 2.3.1) and which represent a "social" resource that can be interacted with by other agents, and **Food**, representing a "physical" resource that can be consumed by agents. This environment is surrounded by blue walls and agents cannot transgress beyond them. Agents detect a wall when it falls into their field-of-vision, and will randomly turn either left or right at an angle of 20 degrees to avoid collision if a wall is detected. Across all of the investigations, agents interacted with the environment across a two-dimensional plane. Several screenshots in this dissertation show this environment as a three-dimensional world, and this is simply for visualisation purposes.

Agents are represented as a doughnut-type shapes (Figure 2.2), initially of size 2 patches (occupying a space of 2 units in the simulation). This size was determined by social rank in later chapters (Chapters 5–9). When agents "die" (as a result of failing to remain "viable" which we discuss in Section 2.3.1), they disappear from the simulated world and can no longer be interacted with.

All investigations used a society of six agents. For all agents, behaviour was driven by a homeostatically-controlled decision-making model (called the Action-Selection Architecture, Section 2.3.1), based on the maintenance of two internal needs. Agents are endowed with two consummatory (goal-achieving) behaviours, *Eat* and *Groom* which are performed on either food or other agents respectively (described further in Subsection 2.3.3). When a resource is not available, agents have a goal-seeking, appetitive behaviour of "wandering". In biological terms, agents could be viewed as conspecifics. Initially, all agents were homogeneous (Chapter 3–4), though this changed in later chapters (Chapters 5–9).

The choice to use six agents was justified by the nature of the investigations and numerous validation experiments before our first investigation (Chapter 3). As there is limited literature studying artificial agents on both an individual and aggregated group level (Section 1.3, but see [66]), we found no precedence on which to build on or consider. Prior bio-inspired simulation models [72], [73], [84], [85], [92] used comparatively large societies, were concerned with aggregated group performance and dynamics, and not investigating micro-level agent performance and interactions. The size of the society had to be manageable enough to analyse individual agent behaviour, and flexible enough to create "groups" of agents of significant size. Furthermore, to overcome some of the limitations of previous agent-based models that we had identified, an appropriate number of agents would (a) allow us to study the dyadic interactions between agents and, (b) allow these dyadic relationships to be investigated in the context of wider group dynamics and interactions.

Food resources are represented by yellow spheres, with a size proportional to its nutritional availability. For instance, resources with a nutritional availability of 4 units would occupy space equivalent to 4 patches

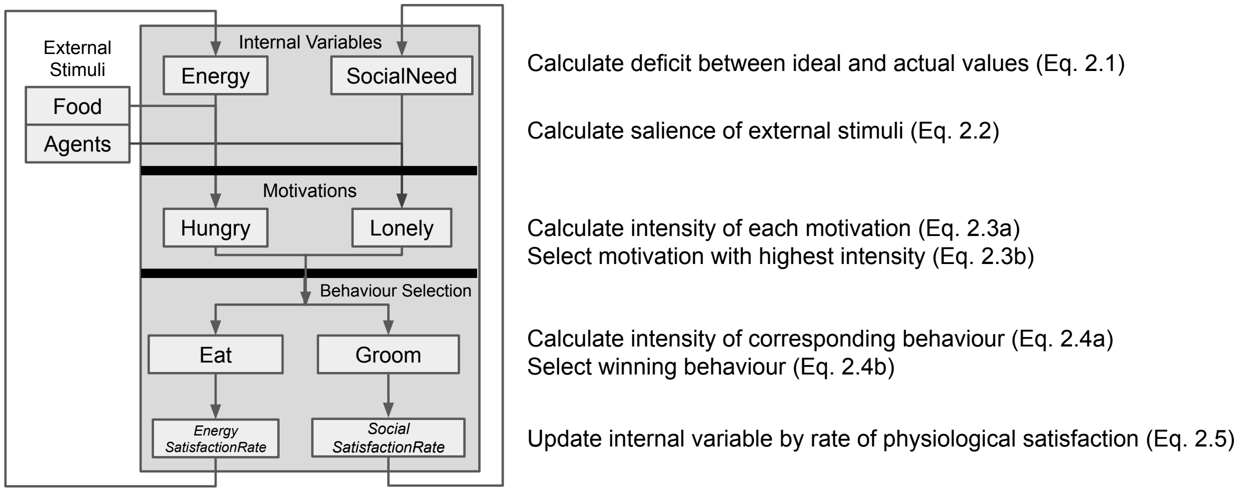

Figure 2.3 Illustration of the basic homeostatically-controlled Action-Selection Architecture used within our investigations, with a brief description of each stage and corresponding equation number. Each chapter incrementally builds upon this basic model. Details of changes are discussed as they are introduced.

(as seen in Figure 2.2). The nutritional value corresponded to how much "nutrition" is available within that particular resource in order to satisfy the agents' *Energy* deficits. Food resources reduces at a rate of $1 \times 10^{-}2$ per time step when agents take "bites" of each food resource. As the nutritional availability is reduced in this way, so does the size of the resource. When resources are not being consumed, they grow back at a rate of $1 \times 10^{-}3$ per time step.

This approach was inspired by the real-world: food grows over time at a much slower rate than it is typically consumed. It also allowed for a semi-dynamic environment with no intervention from researchers to manipulate the environment. In all experiments, these resources can only grow back to a maximum of their initial starting value. Therefore, it was impossible for any environmental condition to have more food available than the starting availability. As food resources can store their own variables, each food resource was able to store its own nutritional availability.

2.3 Agent Model

2.3.1 The Action-Selection Architecture

Building on the work conducted by [34], [35], [39], [40], each agent in our investigation is endowed with an Action-Selection Architecture (ASA): a decision-making model that drives behaviour selection through the maintenance of two homeostatically-controlled variables (Figure 2.3), by keeping their values above its lower bound so as to remain stable (or viable [22]). These variables relate to two competing needs: a physical need, denoted as *Energy*, as well as a social need, denoted simply as *SocialNeed*:

Energy represents a physical need and is analogous to a basic biological requirement to consume nutrition to stay "alive". As such, it is a survival-related variable: when it reaches its lower bound, the agent will "die".

SocialNeed, on the other hand, is non-critical, though still drives a motivation and goal-directed behaviours through the ASA. As the name suggests, it is analogous to a need for social interaction with another agent. As a non-critical variable, agents do not "die" when this value reaches its lower bound.

The goal of the agents is to maintain the two internal homeostatically-controlled variables as close to the ideal points as possible, by correcting any deficits through a relevant goal-directed behaviour: to *Eat* food resources to correct *Energy* deficits, and *Touch* other agents to satisfy their *SocialNeed* deficits. This comprises a two-resource problem of action-selection. The ASA described here selects which of these behaviours to execute at any given time.

Each of these two variables have a lower and upper bound between 0 and 1 respectively, with their ideal value set at 1. Both internal variables start at their maximum value and experience a small deficit at each time step, at an independent rate, when their relevant motivations are not being satisfied. We describe these values fully in Table 2.1. Upon initialisation, the ASA then works as follows:

Step 1: At each time step, the difference between the ideal value of an internal variable ($ideal_i$) and its current value ($current_i$) is calculated, producing an error (deficit) value for that internal variable (d_i).

$$d_i = ideal_i - current_i \tag{2.1}$$

Step 2: The salience of each of the two external stimuli (Food and Agents) is calculated as the perceived amount of each resource:

$$c_i = S_i \tag{2.2}$$

Step 3: The error value (Equation 2.1) is combined with the amount of each relevant stimuli (*Food* and *Agents*) to calculate the intensity, or "urgency", of each of the two motivations ($Hungry$, $Lonely$):

$$m_i = d_i + (d_i \times c_i) \tag{2.3a}$$

$$m_{winner} = \max(m_{Hungry}, m_{Lonely}) \tag{2.3b}$$

Step 4: The motivation with the largest intensity is then selected as the winning motivation. In our two-resource problem, each motivation has a single corresponding behaviour: *Eat* when an agent is *Hungry*, and *Groom* when an agent is *Lonely* (Table 2.1). This constitutes a "Winner-Takes-All" architecture.

$$b_i = m_i \tag{2.4a}$$

$$b_{winner} = \max(b_{Eat}, b_{Groom}) \tag{2.4b}$$

Once a winning behaviour has been selected, it drives an agent to satisfy the relevant motivation by heading towards an available resource related to that behaviour, or wandering around the environment until a relevant resource is perceived. When the winning behaviour has been executed, the relevant internal variable v (*Energy* when agents *Eat*, *SocialNeed* when agents *Groom*) is satisfied at a pre-determined rate p:

$$v_i = v_i + p_i \tag{2.5}$$

Table 2.1 describes the relationship between each homeostatically-controlled variable, the relevant motivation, the behaviour and resource required to satisfy that motivation, and the effect of the behaviour on the agent's physiology. Internal values are updated at each time step and motivations recalculated. A visual representation of the ASA can be seen in Fig. 4.1.

The basis of the architecture presented here remains unchanged throughout all of the experiments. Throughout the course of this dissertation, each chapter incrementally builds upon this framework, moving from a model driven strictly by homeostatic control to a homeostatically-controlled model that is adapted through mechanisms of social allostasis and affect-based social relationships. These changes are described both mathematically and diagrammatically within each relevant chapter.

2.3.2 Hormonal Modulation of the Action-Selection Architecture

Several biological hormones can be considered to underpin affect-based social interactions and social allostasis [29], [122]. Our investigation focuses on two: oxytocin, a hormone associated with social interaction, and cortisol, a hormone associated with stress and adaptation. To justify our selection of these hormones, we highlight some of their hypothesised roles and functions in biological agents.

Oxytocin is a neuropeptide hormone that is released during social behaviours. Given its correlations with prosocial behaviours, it has historically been dubbed the "love hormone', having been associated with the promotion of maternal behaviours [123], [124], trust [125], generosity and reward-sharing [126], and empathy

[127]. It has also been observed to have a number of effects on partner preference [128], [129], including "in-group" bias [130]. These effects are not only restricted to human studies [107], and are likely facilitated in part by oxytocin's effect on the dopaminergic or reward circuitry in the brain [131]. Its association with social behaviours is not always seen to be prosocial in nature. For instance, oxytocin has also been found to have effects on envy [106], decreased trust [132] and out-group aggression [130]. These findings have made researchers reconsider the precise effects of oxytocin, and the contexts in which these effects differ.

Cortisol, on the other hand, is a steroid hormone released as a result of (internal or external) stress. It is associated with energy production [133], heightened vigilance [134] and improved memory formation [135]: all of which are adaptive, survival-related responses. Cortisol is misunderstood as a stress hormone that only has maladaptive consequences, and should instead be considered a hormone of physiological adaptation [29]. However, long-term exposure to chronic levels of cortisol results in physiological detriments [136]: and even long-term medical disorders such as Cushing's syndrome [137] or Addison's disease [138] in humans. However, to characterise cortisol as only a hormone of "stress" overlooks its survival-related adaptation [14]. Individual [115], [116] and social [108], [117] contexts also determine responses to elevated levels of cortisol and stress.

In our model, we define an (artificial) hormone as an internal value that is "secreted" as a function of the internal and/or external environments. These hormones are used to modulate different components within the Action-Selection Architecture calculations. This is line with similar work using modulatory hormones within their action-selection models [34]–[36], [40], [41], [139]. As the specific hormonal mechanisms and modulatory effects used changed throughout this research, each mechanism will be introduced and discussed at the start of each relevant chapter.

2.3.3 Agent Interaction

Agents are endowed with two consummatory (goal-achieving) behaviours, *Eat* and *Groom*, which are performed on either food or other agents respectively to satisfy the current motivation. When a resource is not available, agents have a goal-seeking, appetitive behaviour of "wandering". We briefly describe the way in which agents interact with the environment here.

Vision System: Each agent has a fixed "field-of-view" which describes the maximum length and angle of its perception. This cone is of length 20 patches with an angle of 80 degrees (+/- 40 degrees). This modelling decision was justified after a number of pre-experimental runs prior to the investigation in Chapter 3 being undertaken. A length of 20 patches corresponds to approximately 1/5th of the world length. This cone of

Internal Variable	**Decay Rate**	**Motivation** m_i	**Behaviour** b_i	**Stimuli** S_i	**Physiological Effect** p_{b_i}
Energy	-1×10^{-3}	Hungry	Eat	Food	+ Energy $(+1 \times 10^{-2})$
Social Need	-1×10^{-3}	Lonely	Groom	Agent	+ SocialNeed $(+3 \times 10^{-2})$ +Oxytocin (+1)

Table 2.1 Description of the two internal variables, and the initial decay rate of each of these variables. Deficits of each of these variables are associated with a corresponding motivation and behaviour, with one type of stimuli with which that behaviour can be executed. The final column denotes the effect that performing behaviour b_i has on the corresponding internal variable.

vision remains unchanged for all agents across all experimental iterations. Agents are able to perceive a resource only if it falls within this vision range. A screenshot of this cone of vision is seen in Figure 2.2.

Movement: Agents move through the environment at a default rate of 0.5 patches per time step. In absence of any relevant stimuli, agents randomly wander through the environment, constituting an appetitive behaviour. When a stimulus is perceived that affords an agent to satisfy its winning motivation, wandering becomes more directed in the direction of the resource. When an agent is in close proximity of a resource (within a one-patch radius), it performs one of two consummatory behaviours in order to satisfy each of the two motivations: *Eat* a *Food* resource, or *Touch* another *Agent*.

Eat is executed when an agent is within one patch radius of a food resource. Agents stop at the resource that they want to consume and, at each time step, take "bites" of the food resource. As agents consume the resource, their *Energy* is replenished at a small rate (Table 2.1) for each bite taken. At the same time, the nutritional availability of food resources is reduced.

Groom is executed when agents are within a radius of two patches of another agent. This distance corresponds to being within "touching distance". Unlike *Eat*, the consummatory behaviour of *Groom* is executed in a single time step. Initially, this behaviour satisfies an agent's internal *SocialNeed* by a fixed value (Table 2.1). In later chapters, this rate becomes more dynamic and context-dependent. Agents can perform this behaviour on another agent or, in some scenarios, agents can perform this behaviour on each other at the same time.

Table 2.1 summarises the default conditions of each internal variable, their decay rate, the motivation and behaviours it drives, along with the rate of satisfaction.

2.4 Metrics

2.4.1 Viability Indicator Metrics

The term “viability”, as proposed by Ashby [22], describes how the “survival” of an individual or system (insofar as it can remain stable) is represented by the state of its essential variables: its internal milieu [140]. In biological creatures, some examples of these essential variables include heart rate, blood pressure and blood oxygen levels. These essential variables exist within upper and lower limits, called a “viability zone” (Figure 2.4). When one or several of these essential variables approaches either of these limits, the stability (and therefore survival) of a system or individual is challenged. If these variables are not corrected and transgress these limits, an individual will “die”.

As we described earlier, regulation of these variables can be achieved through homeostatic control, and the selection of appropriate viability measures provides researchers with a more holistic understanding of the stability and performance of a system. In the world of autonomous agents, reports of “viability” should concern itself with both quality and quantity of system stability. We use the measures described below to assess the performance and suitability of the different versions of our model. We therefore employ the approach proposed by [35] who have suggested and described the use of multiple metrics (which they refer to as “viability indicators”) to assess model performance. We describe these viability indicator metrics below.

Life Length

Life Length describes the amount of time an agent has survived (by keeping its survival-related internal variables above 0), as a percentage of the overall simulation runtime:

$$LifeLength_A = \frac{t_{life_A}}{t_{run}} \tag{2.6}$$

Where t_{life} is the time, in time steps, that an agent survived, and t_{run} is the maximum run time of the experimental run.

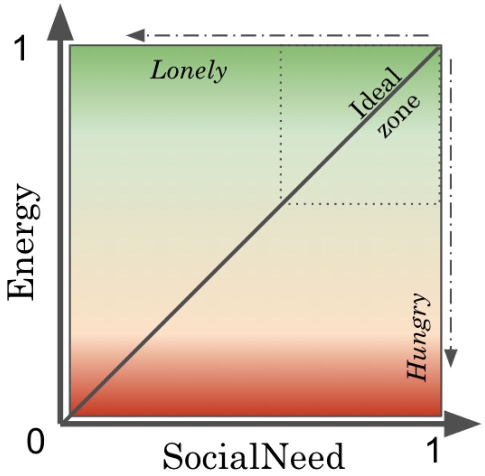

Figure 2.4 A visualisation of the internal "viability zone" for each agent. Each internal variable (*Energy* and *SocialNeed* can exist in the range 0 to 1, with the goal for each agent to keep both variables as close to the green zone as possible (i.e. *Energy* as close to 1 as possible, as *SocialNeed* is non-critical). The closer these internal variables are to the red zone, the closer an agent is to "death". The value of each internal variable is used along with the presence of external stimuli to calculate the intensity of each motivation (*Lonely* or *Hungry*).

Average Comfort

The Average Comfort metric describes the mean satisfaction of the internal homeostatically-controlled variables:

$$AverageComfort_A := \sum_{i=1}^{t_{life_A}} \frac{(1 - \bar{d}_i)}{t_{life_A}} \tag{2.7}$$

Physiological Wellbeing

Finally, Physiological Wellbeing describes the homogenity with which both internal variables are satisfied during an agent's life.

$$PhysiologicalWellbeing_A = \sum_{i=1}^{t_{life_A}} \frac{(1 - \sigma^2(d_i))}{t_{life_A}} \tag{2.8}$$

2.4.2 Additional Metrics

While the viability indicators were the main metrics with which we quantified the model performance, we also reported a number of additional metrics throughout the course of the investigation where we deemed it appropriate. These additional metrics reported quantified measures of individual and group behaviours, and are defined and discussed as they are used in different chapters.

Finally, with the view that the simulated can be viewed as "artificial animals", we supplemented these quantified measures by observing experimental runs and providing qualitative analysis where appropriate in order to provide a holistic analysis of emergent dynamics and agent behaviour. Despite related work analysing agent behaviour using dynamical systems theory [141], we employed an agent-based approach as we were interested in understanding the contexts of agent interactions and behaviours, and then to understand social dynamics as a bottom-up, emergent property of non-homogeneous agent interactions. Given our interest in individual-level, and not purely aggregated, dynamics, an agent-based approach was better suited than a dynamical systems approach for this research. However, we highlight that any subjective discussions regarding agent behaviours and dynamics remain an interpretation of the researcher.

Statistical significance testing was performed using one-way ANOVA between comparison groups, with significance reported at $p \leq 0.05$. In later chapters, correlations between variables was assessed using Pearson's Correlation Coefficient.

Chapter Three

The Social Salience Hypothesis of Oxytocin: An Initial Investigation

3.1 Introduction

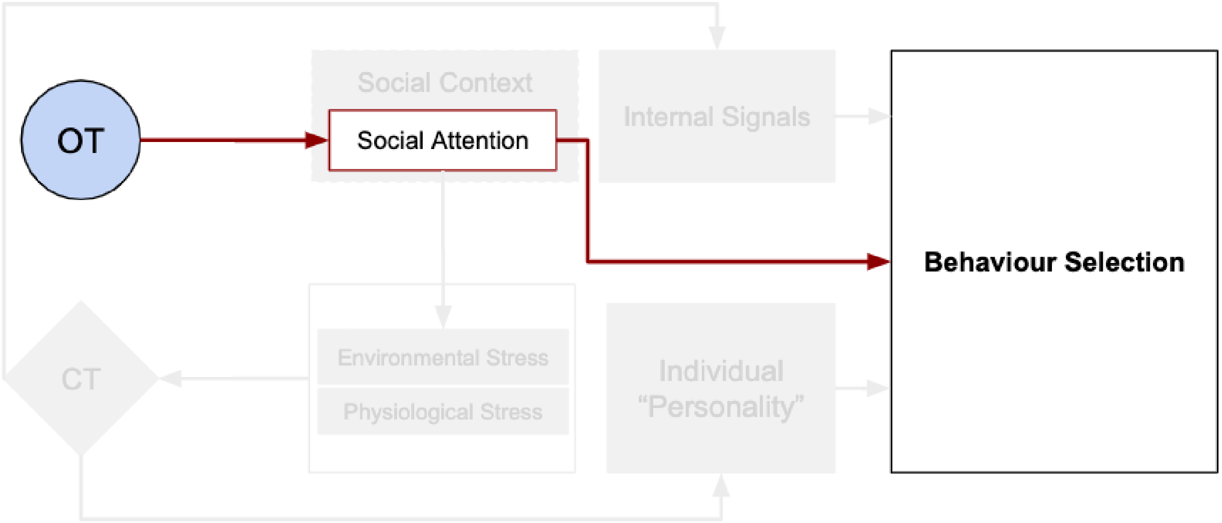

Figure 3.1 A high-level view of the effects investigated in this chapter. Red highlights extensions to the model. OT = Oxytocin. A more granular view of the Action-Selection Architecture can be seen in Figure 3.2

In Chapter 2, we described our simulated agent model as well as the foundation of our decision-making model (called the Action-Selection Architecture) and how each of its steps are calculated. We described the base environmental conditions in which the experiments are conducted. This chapter describes our first investigation into one of the potential mechanisms to use for an adaptive model of social allostasis.

As we have discussed (Subsection 2.3.2), adaptation via social allostasis is driven, in part, by adaptation to an individual's social environment. The ability to regulate, anticipate and adapt social behaviours (that is, actions performed on other social entities) based on internal and external contexts constitutes one type of adaptation. One way in which these social behaviours are adapted may be facilitated by the different levels of attention given to social cues.

Oxytocin has been seen to affect the attention given to social cues (the "social salience" hypothesis)—which may underpin the formation of social relationships and social behaviours—though the precise effects are still not clear. A number of hypotheses have been proposed, and we focus on two contrasting ones in this chapter. One hypothesis is that oxytocin promotes these behaviours by strictly increasing the attention given to social cues (the "prosocial" hypothesis): another that its effects are context-dependent and, in some cases, oxytocin decreases social salience for survival-related outcomes (the "interactionist" approach).

In this chapter, we test the effects of these two contrasting hypotheses on agent viability under three different environmental conditions related to increased challenges in food availability. Our results showed that, in the conditions observed, increased social salience via oxytocin presents significant behavioural and viability advantages in challenging world environments, whereas decreased social salience through oxytocin mechanisms significantly hindered performance. We also observed the emergence of a number of social behaviours related to trust, imitation, and manic-type behaviours, and considered oxytocin's role in these behaviours in real-life individuals.

3.1.1 Oxytocin & Social Salience

For social individuals, social interaction can provide survival-related advantages [107], but also comes with attention, temporal and energy cost [142]. Attention given to social cues may come at the cost of taking away attention given to other survival-related resources. When individuals have multiple needs that must be managed (as is the case in natural systems), or live in socially-challenging environments, operating at an "optimal" level of social vigilance can be a case of life-or-death. Freeing up the attention given to less-relevant stimuli can provide an individual with "attentional resource" that can be focused towards the pursuit of other survival-related goals [107], [126]. In a social context, adaptation—and ultimately survival—is therefore dependent upon how and when attention should be given to cues in the social environment [25], [107], [142].

For social agents, including humans, oxytocin is seen to be the primary hormone that mediates such social interactions. The social salience hypothesis of oxytocin states that oxytocin's effect on social behaviours is facilitated by its effect on attention given to social cues [143]. A number of theories regarding oxytocin's effects on social salience exist, and we focused on two particular ones. The first, described as the "prosocial" hypothesis of oxytocin [125], considers that oxytocin promotes social behaviours, and one of the ways in which this is achieved is through increased attention to social cues. A competing theory, called the "interactionist" approach [106], [107], finds that oxytocin's effects on social behaviours, including its effect on social attention, are diverse and dependent upon individual, social and environmental context. We highlight some of these contextual effects in Subsection 2.3.2. Overall, the precise effects of oxytocin are incredibly diverse and not unanimously agreed upon.

In this chapter, we use the capabilities of our artificial life model to investigate two of these proposed effects in isolation. We created an abstraction of two hypothesised effects of oxytocin on social salience—that it increases social salience (the "prosocial" approach), and that it can decrease social salience (the "interactionist" approach)—and investigated their effects on survival-related (viability) metrics and behaviours in a small society of artificial agents.

3.1.2 Aims & Hypothesis

In this chapter, we aim to investigate two contrasting effects of oxytocin—that it increases, or that it decreases the attention given to social cues—on the survival-related benefits (viability) and behaviour selection of a small group of agents across three different physical environments related to food availability. We hypothesised that group viability would be significantly improved across all conditions when oxytocin increases social salience. We aim to use the results of this investigation to assess one potential mechanism for a model of social adaptation.

3.2 Extensions to the Agent Model

Each of the agents were endowed with the Action-Selection Architecture as described in Chapter 2. In this initial iteration, we introduced the hormone of "oxytocin" with a single mechanism: to modulate the salience of external social cues in the Action-Selection Architecture in one of two ways. We describe these effects below.

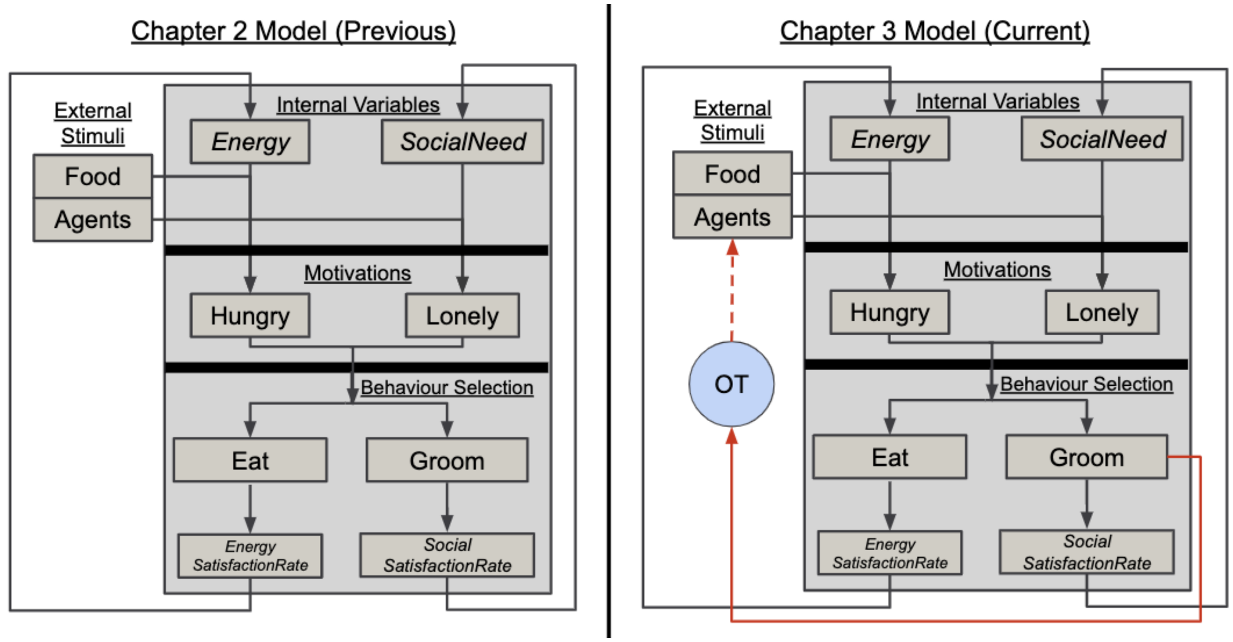

Figure 3.2 A granular view of the Action-Selection Architecture (ASA) used in this investigation (right), compared to the initial model presented in Section 2.3.1. Extensions to the agent model are highlighted in red. In the current model, oxytocin (OT) modulated the salience of social cues in one of two ways, as described in Section 3.2.1. Dashed lines indicate modulatory effects.

3.2.1 Modulation of Social Salience via Oxytocin

We developed a computational model to capture some of the hypothesised effects of real-world oxytocin (OT) mechanisms on social cues, as discussed in Section 3.1. This hormone OT modulated the salience value of social cues (i.e. other agents) within the Action-Selection Architecture, in one of two opposing ways (Figure 3.2). In the first condition, salience of social cues was increased as a function of an agent's internal level of OT. In the second, salience of social cues was decreased as a function of OT.

Increased Social Salience was calculated as:

$$S_{\text{agents}} = C_{\text{agents}} \times (1 + OT) \tag{3.1}$$

Decreased Social Salience was calculated as:

$$S_{\text{agents}} = C_{\text{agents}} \times (1 - OT) \tag{3.2}$$

where C_{agents} is the total number of other perceived agents, and S_{agents} is the final salience calculation. OT values take the range 0–1. Agents are initialised with the maximum OT value of 1. This value of OT decreases at a rate of -0.005 per time step in absence of all behaviours, and increased to its maximum value

	SalienceEffect		
World Condition	No Effect	Increase Social Salience	Decrease Social Salience
Easy	E0	E+	E−
Challenging	C0	C+	C−
Super Challenging	SC0	SC+	SC−

Table 3.1 Table of each experimental condition in this investigation with the corresponding label, using the world condition (E, C, SC) and the modulatory effect of oxytocin on social cues (0 (None), + (Increase), − (Decrease)). Results are reported using these labels.

of 1 in a single time step when agents perform the *Groom* behaviour. The possible range of the modulation functions described above was therefore 0–2.

3.3 Experimental Set Up

Experiments were conducted across three different environmental conditions corresponding to increasing physical challenges related to food availability. These worlds were simply called *Easy*, *Challenging* and *Super Challenging*. Figure 3.3 shows screenshots of each of the respective conditions. Below, we give a brief description of each of the world conditions.

Easy conditions were world conditions in which food resources were considered to be abundant. 18 yellow food resources were available throughout the environment grouped together in small clusters. We proposed that this would serve as the baseline environmental condition, and hypothesised that the abundance of food posed little to no physical challenges for our agents (Figure 3.3, left).

Challenging conditions were world conditions where food resources were scarce. This condition had 6 yellow food resources in two small clusters placed in two corners of the environment (Figure 3.3, centre).

Super Challenging conditions had the same food distribution as Challenging conditions, though the resources had half of the available nutritional value as the Challenging conditions. This constituted an extremely challenging world condition, and we hypothesised that agent performance would be comparatively poorer in these conditions (Figure 3.3, right).

This created a total of 9 experimental groups: 3 corresponding to the direction of modulation of social cues (No Effect, Increase, Decrease) which we call the *SalienceEffect*, and 3 related to world condition (Easy, Challenging, Super Challenging). The labels for each of these experimental groups can be seen in Table 3.1. For simplicity, we have referred to each experimental group using these labels henceforth.

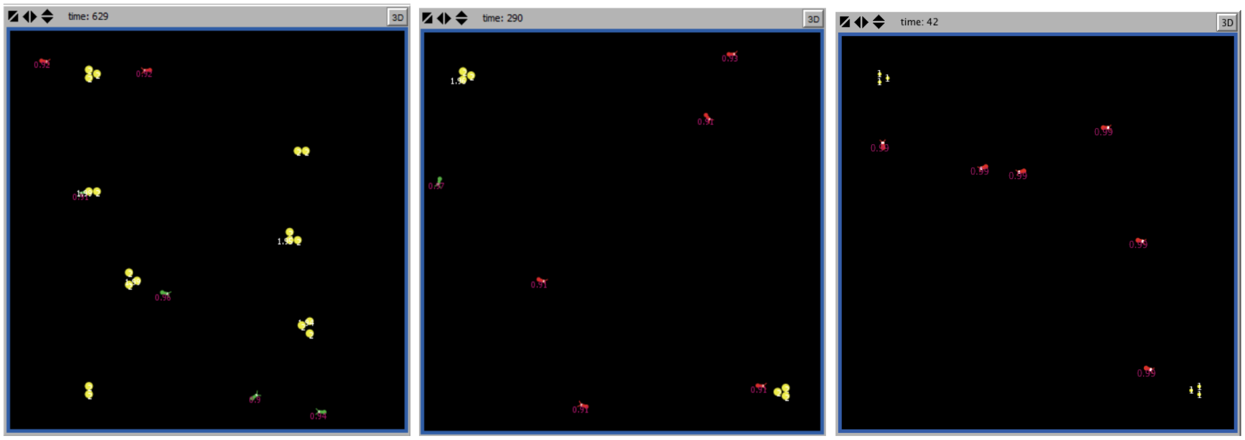

Figure 3.3 Screenshots of the three world conditions: Easy (left), Challenging (centre), and Super Challenging (right). Yellow circles represent food resources. In this chapter, agents were represented using small bug shapes. Agent colour denotes their motivational state (green = Lonely, red = Hungry).

In the Easy and Challenging world conditions, each food resource had a starting nutritional value of 2 units. In the Super Challenging condition, each food resource had a nutritional value of 1 unit. The effects of food resources are discussed in Section 1.3.2.

Each of the six agents in this investigation were homogeneous, and were endowed with the same Action-Selection Architecture, hormonal modulation effects and behaviours across each experimental condition. For the purpose of visualisation, agents would turn one of two colours depending on their motivational state: red when their current motivation was to *Eat*, and green when their current motivation was to *Groom*. These colours were for the benefit of observers and were not detected by the other agents.

Each experimental group consisted of 20 simulation runs, resulting in a total of 180 simulation runs in total. Simulations were run for 25,000 time steps. A time step is defined as one complete iteration of the code. This equated to approximately 45 hours of simulation run time.

Results for each condition were reported using aggregated measures of the viability indicators, Life Length and Average Comfort (Section 2.4.1), for the overall group. To assess the effects of social cue modulation on decision-making, we also reported two additional measures, Motivation Distribution and Behaviour Time. Motivation Distribution described the time, as a percentage of an agent's life, that the agent spent in each of its two motivational states (*Hungry* or *Lonely*). Since agents always experience one of two motivational states, the sum of these percentages is 100%. Behaviour Time described how much time was spent executing each of the behaviours, as a total percentage of Life Length.

Statistical testing was performed on the two viability indicators. Between-group statistical testing was

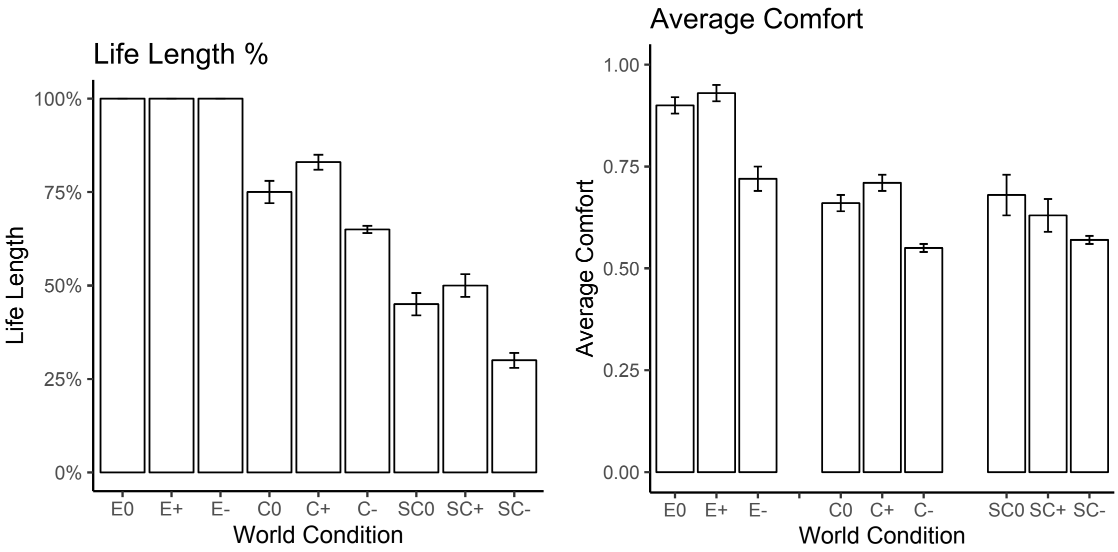

Figure 3.4 Mean Life Length (left) and Average Comfort levels across all world conditions. These results are aggregated values for all agents across all simulation runs. Error bars show the standard error of the mean (SEM).

conducted using one-way ANOVA. Statistical significance was determined using Pearson's correlation coefficient and declared at $p \leq 0.05$.

3.4 Results

3.4.1 Viability Indicators

Easy World

All agents in Easy conditions maintained a 100% survival rate across all experimental runs, and we observed no relationship between the different *SalienceEffect* conditions (E0, E+, E−) and mean Life Length (Figure 3.4, left).

We observed a variance in Average Comfort results between the E0 and the E+ group (0.90 and 0.93 respectively, Figure 3.4, right) groups, though this did not achieve statistical significance ($p = 0.07$). However, we find a statistically significant difference in Average Comfort between the E− group (0.72) and control group (E0: $p < 0.01$), and E− and the E+ group ($p < 0.01$).

Due to the near-perfect viability performance in these world conditions, we considered it to be an adequate baseline to compare the results of the more physically-challenging conditions (Challenging and Super

Challenging) against.

Challenging Worlds

In Challenging conditions, we observed differences in viability performance across the different *SalienceEffect* groups. Agents in the non-oxytocin group (C0) experienced a mean Life Length of 75%: a statistically significant decrease when compared to the same group in the Easy conditions (E0: $p < 0.01$). Agents in the C0 group achieved an Average Comfort level of 0.66: a statistically-significant reduction when compared to the E0 group ($p < 0.01$).

The mean Life Length in the C+ group was 83% (Figure 3.4): a statistically significant improvement when compared to the C0 group (75%, $p = 0.022$), but lower than the mean Life Length reported in the Easy world conditions (E+: 100%),

In the C− group, we observed a reduction in both Life Length performance (65%) and Average Comfort (0.55). For Life Length, this was a statistically significant result compared to both C0 ($p = 0.027$) and C+ ($p < 0.01$). Average Comfort saw a significant decrease compared to the C0 and C+ conditions ($p < 0.01$ for both, respectively).

Super Challenging Worlds

There was a significant reduction in Life Length performance in the Super Challenging world conditions across all *SalienceEffect* groups when compared to the other two world conditions. In non-oxytocin conditions (SC0), agents experienced a mean Life Length of 45%: a statistically significant decrease when compared to both the Easy and Challenging world conditions ($p < 0.01$ for both, respectively).

When oxytocin increased social salience (SC+), mean Life Length increased to 50% ($p = 0.042$ compared to SC0). For SC− groups, we observed a statistically-significant reduced performance in Life Length (30%) , compared to both SC0 and SC+ groups ($p < 0.01$ for both, respectively).

Average Comfort levels were highest in this world condition in the non-oxytocin groups (SC0: 0.68, Figure 3.4, right): higher than the control condition of the Challenging world environment (C0, 0.66), though with a larger variance (± 0.04 vs. ± 0.02). Finally, we found a non-significant reduction in Average Comfort levels in the SC+ group (0.63) compared to SC0 ($p = 0.43$), with a further reduced performance in SC− group (0.57, $p = 0.02$ vs. SC0).

	Experimental Condition								
Behaviour	**E0**	**E+**	**E−**	**C0**	**C+**	**C−**	**SC0**	**SC+**	**SC−**
Eating	3%	3%	1%	4%	4%	1%	11%	10%	8%
Grooming	15%	15%	14%	13%	12%	11%	7%	6%	2%
Wandering	82%	82%	85%	83%	84%	88%	82%	84%	90%

Table 3.2 Time spent executing each of the behaviours, Eating and Grooming. Time spent not executing one of these two consummatory behaviours is spent *Wandering*.

3.4.2 Motivational State Distribution and Behaviour Time

Across the control conditions (E0, C0, SC0), Motivational State Distribution was in line with the relative physical challenge of the environment. In Easy world conditions, we observed a 10%pts difference in time spent Hungry (45%) and Lonely (55%). In the Challenging world conditions, this difference reduced to a more even distribution, with only a 4%pts difference between Hungry (52%) and Lonely (48%) motivational states. In the Super Challenging world condition, this difference increased to 20%pts, with agents spending more time Hungry (60%) than Lonely (40%).

In conditions where oxytocin increased social salience (E+, C+, SC+), we expected Motivational State Distribution to favour the Lonely motivational state. This was seen in only the Easy condition (E+: 59%). The C+ group saw a slight reduction in the time in the Lonely motivational state (-3%pts) compared to C0, while the distribution in the Super Challenging world (SC+) remained unchanged.

In the Easy world, when oxytocin decreased social salience, agents spent 60% of their total life in the Hungry motivational state. In more challenging conditions, we saw the inverse, with agents spending 83% and 78% of their time in the Lonely motivational state, in the Challenging and Super Challenging worlds respectively.

Table 3.2 shows the Behaviour Time of the society across each of the conditions. In the Easy and Challenging environmental conditions, groups spent significantly more time *Grooming* (11−15%) than *Eating* (1 − 4%). In Super Challenging conditions, this trend is reversed. Across all environmental conditions, we find that decreased *SalienceEffect* groups (E−, C−, SC−) show decreased time spent eating when compared to the other *SalienceEffect* groups. In the latter two conditions, this reduced time eating corresponds with reduced motivation to *Eat*. Across these same conditions (E−, C−, SC−), we observed increased time wandering and not performing any consummatory behaviour, compared to the other two *SalienceEffect* groups.

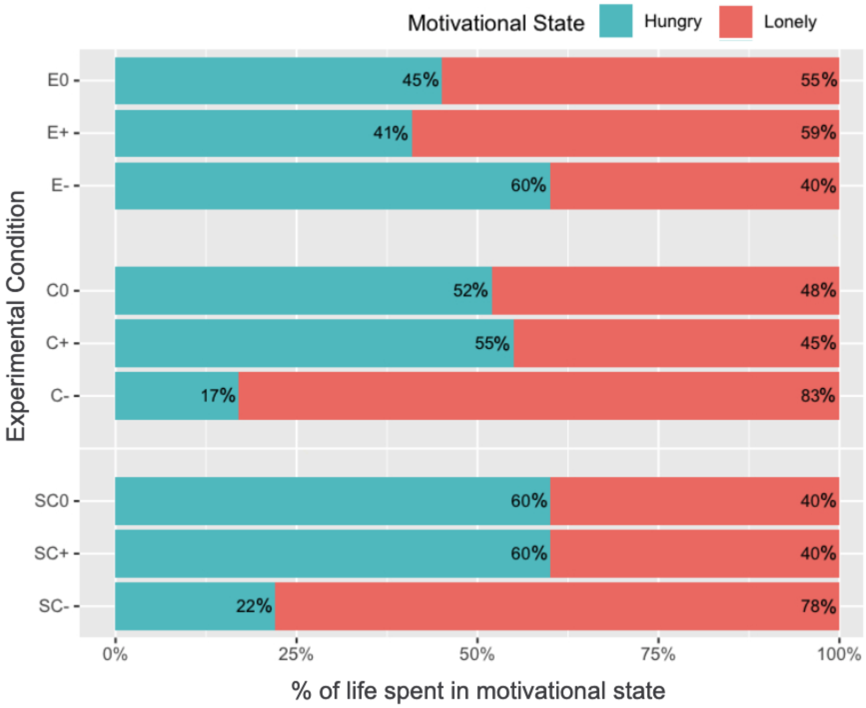

Figure 3.5 Motivational State distribution, showing the spent in each motivational state as a percentage of an agent's life, broken down by experimental condition (y-axis).

3.5 Discussion

Overall, we find that agent viability was significantly improved when oxytocin increased the salience of social cues across all three world conditions. Where Life Length remained unchanged in the Easy condition, viability improvements were still reported with increased Average Comfort. We observed both increased Life Length and Average Comfort levels in the Challenging world condition when oxytocin increased social salience. We observed a non-significant decrease Average Comfort when oxytocin increased social salience in the Super Challenging group, though we consider that this may ne an effect of the relative difficulty of the environment.

In addition to improving agent viability performance by increasing social salience through oxytocin, we found decreased viability across all environments when oxytocin decreased social salience. Furthermore, the time spent performing the prosocial behaviour (*Grooming*) was also significantly reduced in these conditions. Here, we suggest a causal link between the prosocial behaviour of *Grooming*, facilitated by oxytocin mechanisms, and agent viability. We discuss a number of observations of our experiments below.

3.5.1 Attention to Social Cues Affects Eating Behaviours

In our model, oxytocin's effect on social salience affected the motivation to perform a prosocial behaviour (Equation 3.1), and we had predicted that increased social salience groups (E+, C+, SC+) would show increased motivations to *Groom* as a result of the increased attention given to social cues.

Contrary to this prediction, in Challenging and Super Challenging world conditions, we found that agents spent more time in the *Hungry* motivational state rather than the *Lonely* (Figure 3.5). We propose that this is due to the relative physical challenge of the world condition. Additionally, agents in these conditions were able to perform a similar amount of *Grooming* with others compared to the C0 and SC0 groups (Table 3.2).

In Super Challenging conditions, we observed a significant difference in the time spent eating compared to the other two environmental conditions (Table 3.2). Watching the simulation runs, we found that, due to the relative sparsity of food resources, agents would maximise the opportunity to eat once they had located a food resource. As a result, agents in these conditions were found to spend longer periods of time eating for longer periods to satisfy larger internal deficits—akin to a "feasting-type" behaviour. In contrast, where the internal *Energy* deficit was smaller in the other two world conditions (Easy and Challenging), eating behaviours in those groups were seen to be "little-and-often", allowing for internal deficits to be satisfied regularly and freeing up attention to satisfy their other internal need (*SocialNeed*).

3.5.2 Increased Social Salience through Oxytocin Results in Emergent "Imitation"

We highlight that the increased viability in the Increase Social Salience conditions was, in part, due to the emergence of agent "imitation" resulting in indirect cooperation towards the acquisition of food resources.

As we mentioned in Section 2.2, agents wandered through the environment and head towards salient resources in their vision if that particular resource satisfies their current motivational state. As another agent represents a dynamic resource that both affects and satisfies their motivation to *Groom*, these agents (for instance, A1 in Figure 3.6) would follow the other agent (A2) order to *Groom* with it. In some cases, the agent being "followed" would locate a food resource and—as a result of this following behaviour—both agents (A1 and A2) are able to locate the food resource. Agents then satisfy their *Energy* deficit through this opportunistic finding of resources. The proximity afforded by both agents eating together also increased their likelihood to *Groom*, therefore satisfying both internal needs in a short space of time. An illustration of this behaviour can be seen in Figure 3.6.

This behaviour was also seen to extend to three agents in some simulation runs and, in some rare

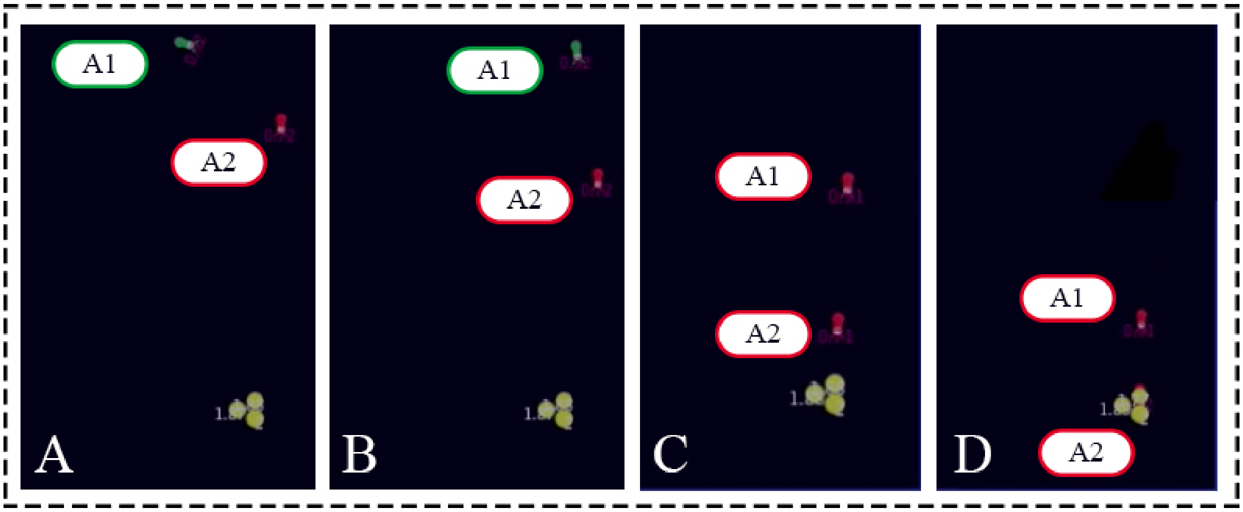

Figure 3.6 Screenshots showing how agents followed or "imitated" others towards food resources. A: Agent A1 is motivated to groom, and sees agent A2 who is searching for food. B: A1 begins to follow A2. C: A2 sees food to satisfy its Hunger motivation, and by following A2, A1 also sees this food. D: Both A1 and A2 successfully find and acquire food. This satisfies their Energy variable and, by virtue of being close to each other, are more likely to Groom with each other (not shown). Agent colours represent motivational states: Green when they are Lonely and looking for agents, and red when they are Hungry and looking for food. Picture C shows how A1's motivation changes as a result of seeing food.

circumstances, extended to four agents organising in such a way. Therefore, increasing salience to social cues not only allowed agents to exploit the presence of other agents, but also to utilise the other agents' range of vision in the location and consumption of food resources.

Such behaviour did not emerge in groups where oxytocin decreased social salience (E−, C−, SC−). In those conditions, we observed more of a solitary approach towards food acquisition, compared to this cooperative behaviour towards food acquisition described above. In conditions where food resources were difficult to find (Challenging and Super Challenging), this behaviour provided survival-related benefits for agents.

We highlight that this emergent cooperation driven by increased attention to social cues—the effect related to the prosocial hypothesis of oxytocin— was a social behaviour that resulted in improved agent viability in these conditions. We propose that this behaviour, underpinned by oxytocin's effect on social salience, may constitute a form of behavioural adaptation through the social environment for these agents [144].

3.5.3 Decreased Social Salience Results in Social "Ignorance"

As we saw no survival-related benefits for decreasing social salience through oxytocin, we performed several observations on simulation runs in these conditions to investigate the reasons further. From these observations, we highlight a number of emergent behaviours seen in these conditions that may have contributed to reduced viability performance.

We found that when oxytocin levels were high, agents in the C$-$ and SC$-$ groups (i.e. when oxytocin decreased social salience) would not perform *Grooming* with other agents to satisfy their *SocialNeed* deficit. In some cases, this *SocialNeed* deficit would exceed the *Energy* deficit. Instead, we found that these agents "ignored" others and continue to wander through the environment, in favour of food acquisition.

As internal motivational states were driven as a function of both internal deficits and external cues (Equation 2.3a), oxytocin's inhibitory effects on attention given to other cues (Equation 3.2) resulted in agents not being motivated to *Groom* in this scenario, despite an available social resource and an internal deficit.

Such social ignorance related to oxytocin levels may underpin states of distrust, and has also been observed in the real world. For instance, [145] discussed how elevated oxytocin can decrease trust and prosocial behaviours in individuals with borderline personality disorders. We suggest that the social withdrawal and ignorance of other individuals that we have described may have emerged as a result of oxytocin's effect on attention to social cues. These observation may provide an insight into one potential mechanism underpinning these behaviours in biological agents, and provides a potential avenue for further investigation.

3.5.4 Oxytocin's Secondary Effects as "Social Memory"

After watching the simulations, one unexpected consequence of our model was how oxytocin levels were indicative of recency of a *Grooming* behaviour taking place. Specifically, higher levels of oxytocin indicated a recent social interaction, and lower levels of oxytocin indicated that such an interaction had not taken place for an extended period of time. This was most notable in groups where oxytocin decreased social salience (E-, C-, SC$-$): where, after *Grooming* with another, agents would be seen to "ignore" other agents as a result of elevated oxytocin levels (Subsection 3.5.3).

Here, we suggest that oxytocin may serve a secondary function related to "social memory" in our model: an effect also seen in biological systems [146]. This effect may also play a role in oxytocin's diverse, context-dependent effects [145]: where a recent social interaction can "free up" attentional resource in pursuit of other goals, which may therefore facilitate long-term survival.

Currently, our agent model does not consider the recency of a social behaviour (via oxytocin levels) when satisfying the internal *SocialNeed* variable. Instead, *SocialNeed* is corrected by the same value (+0.5) after all *Grooming* interactions. Taking the view that oxytocin serves as a measure for how recently a social behaviour has been performed, we propose that adaptation of this internal physiological signal (*SocialNeed*) as a function of oxytocin may provide further insight into our initial findings. We consider this as an avenue for further investigation.

In terms of a low-cost computational model for social interactions, such a social memory model may provide some computational advantages over other approaches [147]–[150].

3.6 Summary

This chapter presented our initial work investigating one potential mechanism that underpins social behaviours and the formation of social relationships. Amongst other effects on social behaviour, oxytocin has been hypothesised to play a role in social behaviours by affecting the attention given to social cues, though views on its precise effects are still not unanimous.

We investigated two contrasting effects related to two hypotheses of oxytocin's effects—by increasing and decreasing the attention given to social cues, respectively—and accounted for these effects in our agent model. We assessed the effects on the viability and behaviours of a small society of homogeneous agents across three world conditions related to food availability. We had hypothesised that agent viability would be improved when oxytocin increased social salience across all conditions.

Based on the quantitative results and qualitative observations, we observed viability benefits for all agents when oxytocin increased the salience of social cues. From our results, we also found some support for the prosocial hypothesis of oxytocin (which posits that oxytocin's effects specifically enhance prosocial choices and behaviour) in our experimental conditions. Conversely, we found that groups where oxytocin decreased the attention given to social cues caused a detrimental effect to viability performance.

Despite the low-level hormonal mechanisms that we have modelled, the emergence of higher-order behaviours related to trust or imitation may provide some insight into the underlying mechanisms of such behaviours in natural systems. We highlight that these emergent behaviours would not have been identified by the quantitative data analysis alone. This initial investigation therefore provided insight into the value of using observational analysis to understand the behaviour of (artificial) agents, and we continue this approach in future chapters.

We found that, as oxytocin is released through social interaction, internal levels of oxytocin provided a proxy measure of recency of social interaction: where high levels were indicative of recent social interaction, and vice-versa. These hormonal mechanisms may also help shape internal representations of social conditions—such as the availability of support, or a recent social interaction—which may provide contextual information for the adaptation of oxytocin's effects. We believe that such a mechanism may overcome the computational requirements of previously-proposed social memory models.

Though we had found support for the prosocial hypothesis, these findings are limited as we only considered one potential effect of the hormone. Though we do not aim to account for all of the potential contexts regarding oxytocin's proposed diverse effects, we identified that one additional area that could be investigated is through the internal signals related to social behaviours. As a result of the different frequencies of grooming behaviours between our experimental conditions, along with our observation that oxytocin levels serve as an internal representation of recency of social behaviours (akin to social memory), we considered additional context in the form of internal signals of physiological satisfaction of social behaviour. In the next chapter, we consider oxytocin's potential effects on these internal signals in the context of "affective touch".

Chapter Four

Investigating the Role of Affective Touch on the Social Salience Hypothesis of Oxytocin

4.1 Introduction

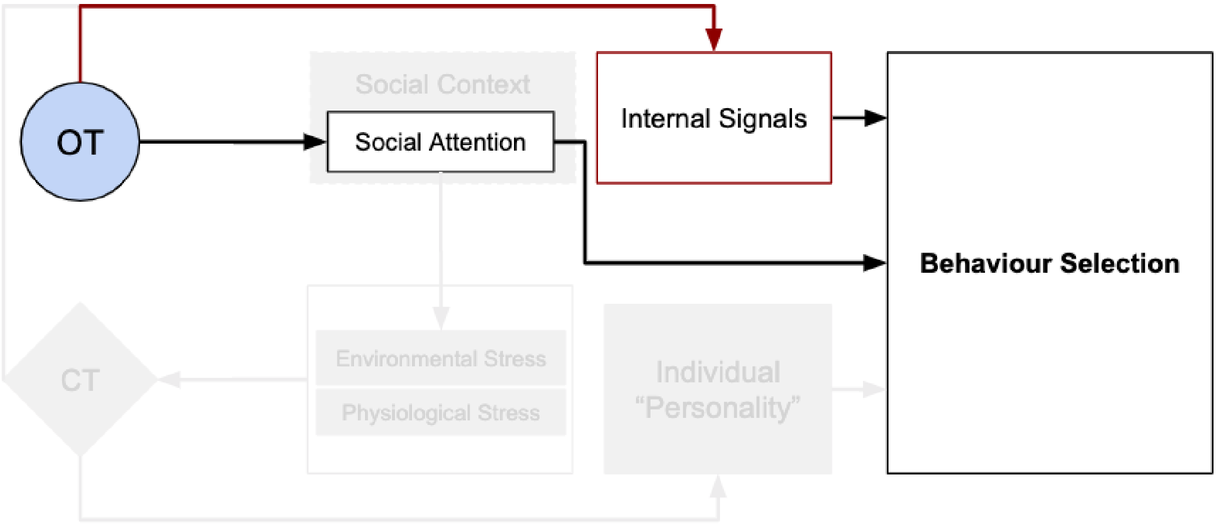

Figure 4.1 A high-level view of the effects investigated in this chapter. Red denotes extensions to the model from the previous chapter. OT = Oxytocin. A more granular view can be seen in Figure 4.2

In the previous chapter, we investigated two contrasting effects related to oxytocin's role in the attention given to social cues—by both increasing and decreasing attention to social cues as a function of oxytocin levels—and tested their effects on group viability and behaviours. Though we found some viability benefits through effects related to the "prosocial" approach of oxytocin—that oxytocin promotes prosocial choices—these findings required further investigation as the model lacked any of the additional context associated with oxytocin's diverse effects on social behaviours. As we continued our investigation into the social salience hypothesis of oxytocin, we aimed to address this limitation in our model.

Observations of simulations in Chapter 3 found that oxytocin, released as a function of social interaction, can also serve as a measure of how recently a positive social interaction has been performed. From this observation, we considered that this may also provide an additional mechanism through which individuals adapt social behaviours, which may provide some further insight into oxytocin's contextual effects.

One way in which this additional context can be accounted for is through the internal signals associated with performing prosocial behaviours. One example of a prosocial behaviour is that of tactile contact with others, which can express or invoke an emotional state and can be considered "affective" touching.

However, the perceived satisfaction associated with tactile contact can be dependent upon the social or environmental contexts, as well as the affective state of interacting individuals. Oxytocin has been seen to facilitate impressions and responses of such tactile contact, though its precise effects on these internal signals are still not fully understood. The diverse effects of oxytocin on these internal signals, along with its context-dependent effects on external social cues, have made it difficult to study these effects at the same time in real-world individuals.

We extend on the initial investigation into the social salience hypothesis of oxytocin (Chapter 3) by taking into consideration a potential secondary effect of oxytocin. In this chapter, we accounted for oxytocin's proposed effects on modulating internal signals related to the pleasantness associated with performing a prosocial tactile behaviour.

We find that, in contrast to our initial results in Chapter 3, the adaptive effects of oxytocin on social salience may be context-dependent. In terms of the current investigation, when oxytocin's effects on internal "satisfaction" signals of social interaction were accounted for, its survival-related benefits may extend beyond a strictly-prosocial approach. Unlike our results from Chapter 3, our results from this chapter found survival-related benefits in some conditions when oxytocin decreased the attention to social cues. As a result, we considered our findings to be supportive of the "interactionist" approach towards oxytocin's effects on social behaviours.

In terms of our wider research investigating potential mechanisms underpinning social relationships and long-term adaptation, the results from this investigation suggest that there is no one-size-fits-all approach to oxytocin's effects with respect to agent viability. Instead, the effects of these hormonal should account for additional contextual information.

4.1.1 Oxytocin & Affective Touch

Along with auditory and visual stimuli, interpersonal tactile contact is one modality through which individuals interact and communicate socially, and plays a key role in the formation of social bonds [151]. "Affective touch" can be considered to be any physical tactile contact between individuals which communicates, express or evokes an emotional state [152]–[154]. In humans, tactile contact such as touching can amplify the intensity of emotional displays from facial and auditory (e.g., vocal) cues [155] and can also communicate the positive or negative hedonic values of emotional states [156].

In social mammals (including humans), oxytocin has been found to be released after positive social tactile behaviours [157], [158]. Oxytocin has been seen to increase the perceived pleasantness of tactile interaction [159] as well as social evaluations after receiving it [160]. Importantly, social touch can act as an important, adaptive behaviour for stress-reduction, promoting longer-term survival and wellbeing for social individuals [108], [161].

However, the perceived pleasantness of this behaviour is shown to have inverse behaviours in individuals with traits related to Autism Spectrum Disorder [159]. Additional mixed findings of oxytocin on the internal signals and behavioural responses to "affective" touch [162]–[164] suggest that these effects are more diverse than previously thought. More recent work has found that oxytocin can modulate attention switching between internal and external cues [165], and that it may facilitate integration between internal signals and external cues [166].

Potentially as a result of oxytocin's diverse, context-dependent effects, studies investigating the concurrent effects of the hormone on both internal and external signals have been difficult to study in the real-world [165], [166]. Nevertheless, oxytocin appears to play a role in the processing of internal social signals. We are also motivated by the limitations of studies integrating oxytocin's effects on internal and external signals in biological systems, and look to address this limitation in this chapter.

4.1.2 Aims & Hypothesis

We continue our investigation into the two contrasting hypotheses of oxytocin's effects on social cues—the "prosocial" and "interactionist" approaches—by accounting for its effects on the perceived pleasantness (the physiological satisfaction) associated with performing a social behaviour. We aim to investigate how accounting for this second modulatory effect of oxytocin affects the viability of a group of agents.

We investigate two contrasting effects of oxytocin on the physiological satisfaction associated with tactile behaviours—both increasing and decreasing internal signals of satisfaction—across three world conditions related to food availability.

Based on our findings from Chapter 3, we hypothesised that group viability would be improved when oxytocin both increased social salience, and increased the internal satisfaction ("pleasantness") associated with performing social (tactile) behaviours.

In terms of the development of our adaptive model, we aim to use the findings from this investigation to understand the viability effects of one type of physiological adaptation through social interaction.

4.2 Extension to the Agent Model

In this investigation, we retain oxytocin's effects on modulating social salience from Chapter 3, and introduce a second effect: modulating the rate at which the internal variable *SocialNeed* is corrected after a *Grooming* behaviour takes place. In the previous model, this rate was a static value of +0.05 (Table 2.1). In this investigation, this rate becomes dynamic, modulated by internal oxytocin levels in one of two ways.

The rate of satisfaction of the internal *SocialNeed*, called the *SocialSatisfactionRate*, was dynamically modulated through internal oxytocin levels in one of two ways. In the first instance, *SocialSatisfactionRate* was directly modulated by the amount of oxytocin in an agent's physiology (referred to as the Direct modulation groups). Specifically:

$$\text{SocialSatistfactionRate} = p_{\text{touch}} \times \text{OT} \tag{4.1}$$

In the second, it was inversely modulated by the amount of oxytocin (referred to as the Inverse modulation groups):

$$\text{SocialSatistfactionRate} = p_{\text{touch}} \times (1 - \text{OT}) \tag{4.2}$$

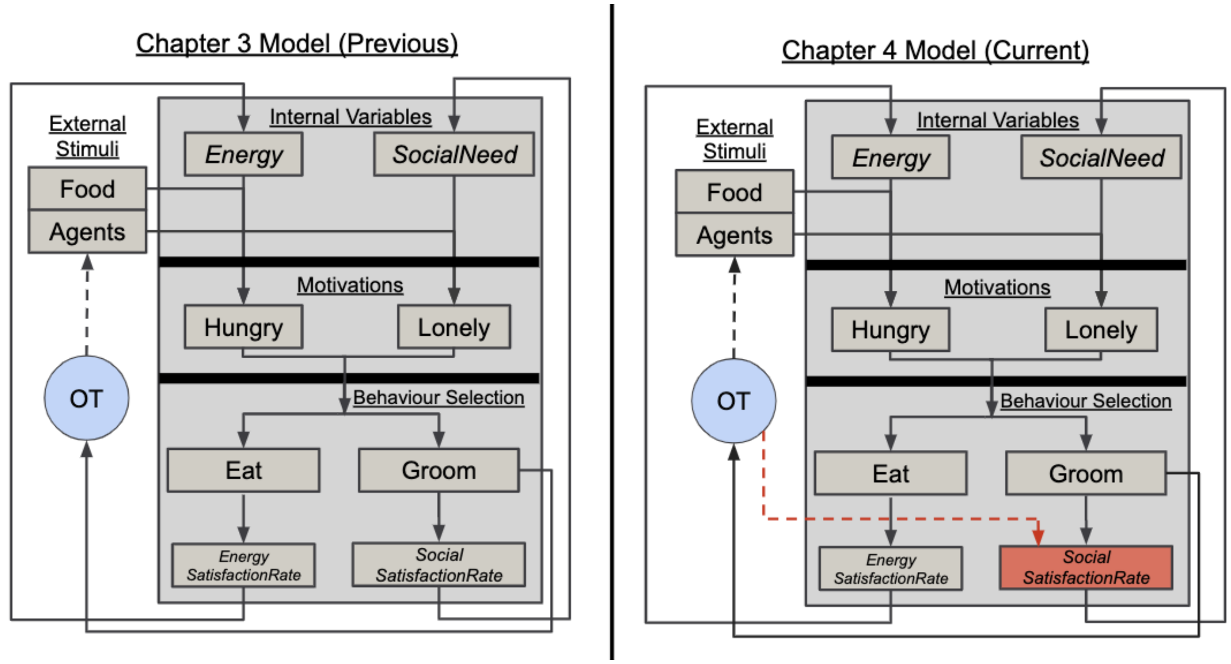

Figure 4.2 Illustration of the agent model used in this investigation (right) compared to the previous chapter (left). Extensions to the previous model are denoted in red. In the current model, oxytocin modulates both the salience of social cues (Equations 3.1 and 3.2) and the physiological *SocialSatisfactionRate* of executing the *Groom* behaviour (Equations 4.1 and 4.2). Dotted line indicates modulation of values.

where p_{touch} is the default value of physiological satisfaction of *Grooming*. The internal *SocialNeed* variable is then updated by the *SocialSatisfactionRate* value:

$$\text{SocialNeed}_t = \text{SocialNeed}_{t-1} + \text{SocialSatisfactionRate} \tag{4.3}$$

In Chapter 3, the rate at which *Groom* satisfied the *SocialNeed* (p_{touch}) was a fixed value of +0.05. Using the modulatory effects of oxytocin described above, the internal rate of satisfaction was a dynamic value between 0 and 1 in the current investigation.

4.3 Experimental Set Up

Experiments were conducted across several parameters: three different conditions related to oxytocin's effects on modulating social salience, two effects of oxytocin's modulation on internal satisfaction signals, and three world conditions. This created a total of 18 different combinations of experimental conditions. For concision, results of each condition will be referred to using their labels in Table 4.1.

Firstly, as per the previous chapter, this investigation used three different effects of oxytocin on social

Oxytocin Effects		**Environmental Conditions**		
Social Salience Direction	**SocialSatisfactionRate Modulation Direction**	**Easy**	**Challenging**	**Super Challenging**
None	Direct	E0/D	C0/D	S0/D
None	Indirect	E0/I	C0/I	S0/I
Increase	Direct	E+/D	C+/D	S+/D
Increase	Indirect	E+/I	C+/I	S+/I
Decrease	Direct	E−/D	C−/D	S-/D
Decrease	Indirect	E−/I	C−/I	S-/I

Table 4.1 Table of each experimental condition with the corresponding label. Experimental conditions are labelled with the environmental condition (E = Easy, C = Challenging, S = Super Challenging), the modulation effect of oxytocin on social cues (0 = No Effect, + = Increase Social Salience, − = Decrease Social Salience) and modulation effect of oxytocin on the *SocialSatisfactionRate* (D = Direct, I = Inverse).

salience: None (0), Increasing Social Salience (+), and Decreasing Social Salience (−). The modulatory effects remained unchanged from Chapter 3 (Section 3.2.1).

Secondly, we investigated two different effects related to oxytocin's modulation of the physiological satisfaction associated with performing *Grooming* (called the *SocialSatisfactionRate*): as both a direct function of oxytocin (Equation 4.1), and an inverse function of oxytocin (Equation 4.2). These were referred to as the "Direct" (D) and "Inverse" (I) modulation groups, respectively.

In terms of world conditions, we investigated these effects across three different world conditions related to food availability: Easy (E), Challenging (C), and Super Challenging (S). These conditions remained unchanged from Chapter 3. Descriptions of these world environments are found in Section 3.3. Figure 3.3 shows screenshots of each of these conditions.

To create an appropriate control group, experiments were conducted in two stages. In the control experiment (Experiment 1), we used static values of *SocialSatisfactionRate* from 0.05–1.0 (i.e. 5% to 100% of the maximum value of *SocialNeed*), in 0.05 (5%) increments. Our previous model had validated the model using values between 0.01–0.10, but further investigation was limited due to computational limitations. We address those limitations in this investigation.

In the second stage (Experiment 2), we dynamically modulate all values of *SocialSatisfactionRate* of between 0.05–1.0 by internal values of oxytocin using the *Direct* and *Inverse* modulation functions described in Section 4.2.

Each experimental condition was run for 20 simulation runs. Unlike the previous study which had a maximum run-time of 25,000 time steps, each simulation in this investigation was run for 30,000 time steps. This decision to increase simulation run time was due to both an increase in computational resource, as well as to further validate whether the results from the Easy environments in Chapter 3 (where we observed agents surviving 100% of the simulation run) was due to the simulation run time or the environmental conditions.

In total, the combination of different parameters across all experimental conditions resulted in a total of 7,200 experimental runs. This equated to approximately 122—125 hours of real time.

We reported results using the same viability indicators as Chapter 3: Life Length and Average Comfort (Section 2.4.1). Unless specified, results were aggregated for all values of *SocialSatisfactionRate* in that condition. Statistical significance testing was performed using one-way ANOVA tests between groups. Statistical significance was reported at a $p = 0.05$.

4.4 Results

We first report results of our control condition which used static values of *SocialSatisfactionRate* across all values between 0.05 and 1 (Experiment 1) in 0.05 increments. In those conditions, oxytocin modulated the salience of social cues, but had no effect on the *SocialSatisfactionRate*. Our results from this experiment found little change between the values of *SocialSatisfactionRate*, and, unless specified, results are reported as aggregated for all *SocialSatisfactionRate* values. We then report report results for experiments where oxytocin modulates both social salience and *SocialSatisfactionRate* (Experiment 2)

4.4.1 Experiment 1: Static Rate of Physiological Satisfaction

Easy Environment

In terms of Life Length, all agents in the Easy environment survived for the entirety of the simulation length (100% survival rate) regardless of oxytocin's effects on social salience (Table 4.3). This was in line with the results observed in those same environment from Chapter 3.

Mean Comfort levels of the group were similar between each of the *SocialSalienceDirection* and *SocialSatisfactionRate* conditions. In control conditions (E0), Comfort levels remained consistent at 0.93 across all *SocialSatisfactionRate* values, 0.94 for Increased Social Salience groups, and 0.92 for Decreased Social Salience groups. Differences were non-significant between groups. (Appendix D).

Despite an increased simulation run time, all agents in these Easy environments experienced both a long

Experimental Condition	**World Condition**	**SocialSalience Direction**	**Life Length**	**vs. 0**	**vs. +/-**	**Comfort**	**vs. 0**	**vs. +/-**
E0	Easy	None	100%	n/a	NS	0.93	n/a	NS
E+	Easy	Increase	100%	NS	NS	0.94	NS	NS
E-	Easy	Decrease	100%	NS	NS	0.92	NS	NS
C0	Challenging	None	68%	n/a	NS	0.71	n/a	NS
C+	Challenging	Increase	68%	NS	*	0.72	NS	*
C-	Challenging	Decrease	67%	NS	*	0.70	NS	*
S0	Super Challenging	None	43%	n/a	*	0.66	n/a	*
S+	Super Challenging	Increase	43%	NS	*	0.67	NS	*
S-	Super Challenging	Decrease	40%	*	*	0.64	*	*

Table 4.2 Experiment 1 results for Life Length and Comfort levels, aggregated across all simulation runs, *SocialSatisfactionRate* values and agents. "vs. 0" indicates statistical result compared to the control group of each world condition. "vs. $+/-$" indicates statistical result compared to the opposing *SocialSalienceDirection* group. Statistical results are denoted with a * for $p < 0.05$, ** for $p < 0.01$ and NS for no statistical significance. "n/a" indicates that a statistical test was not performed.

and comfortable life across all simulation runs, regardless of oxytocin's effect on social salience or the value of *SocialSatisfactionRate*. Trends across each of the *SocialSatisfactionRate* values can be seen in Appendix D.

Challenging Environment

As in Chapter 3, we saw a slightly decreased Life Length (68% vs. 100%, $p < 0.01$) and Average Comfort (0.71 vs. 0.93, $p < 0.01$) in the control conditions in the Challenging world (C0) when compared to the Easy environments (E0). Groups where oxytocin increased social salience (C+) saw no changes in Life Length and non-significant improvements in mean Comfort (0.72 vs. 0.71, $p < 0.01$) when compared to the control group (C0). For the C$-$ groups, differences for both Life Length (67%) and Comfort (0.70) were both non-significant vs. control.

Focusing on groups where *SocialSatisfactionRate* was less than 0.20 (Appendix D), we found some statistically-significant differences between C+ and C0 groups in terms of Life Length (70% vs. 66%, $p < 0.01$) and Mean Comfort (0.69 vs. 0.67, $p < 0.01$). For C$-$ groups, we saw reduced viability compared to control, in terms of both Life Length (57% vs. 66%, $p < 0.01$) and a non-significant change in mean Comfort (0.65 vs. 0.67, $p = 0.09$).

In sum, we found little differences between the *SocialSalienceDirection* conditions for values of *SocialSatisfactionRate* beyond 0.20, but saw some statistically-significant differences for lower values.

Super Challenging Environment

As with our previous investigation, we saw significantly reduced performance in terms of both Life Length and Comfort levels in the Super Challenging environments when compared to both the Easy ($p < 0.01$ across all metrics) and Challenging ($p < 0.01$ across all metrics) world conditions.

For the S+ groups, aggregating across all *SocialSatisfactionRate* values, Life Length performance remained unchanged vs. control (S0: 43%, S+: 43%, $p = 0.80$), and there were non-significant differences in Comfort levels (S0: 0.66, S+: 0.67, $p = 0.41$). In terms of S− groups, we saw a statistically-significant ($p < 0.01$) reduction in both Life Length (40%) and Comfort (0.64) compared to control. Results were statistically-significant differences compared to both control and the S+ groups (Table 4.2).

Looking specifically at groups where *SocialSatisfactionRate* was below 0.20 (Appendix D), results from the S0 and S+ groups remained similar, though we saw reduced performance in terms of Life Length in the S− groups, ranging from 28%–37%. This reduced performance was statistically-significant vs. the S0 and S+ groups.

4.4.2 Experiment 2: Modulating Physiological Satisfaction Rates Through Oxytocin

Easy Environments

When oxytocin Directly modulated *SocialSatisfactionRate* in Easy world conditions, agents survived for 100% of all simulation runs. Average Comfort levels remained at 0.94 for E0/D and E+/D conditions, with similar values observed in the E−/D condition (0.80 to 0.92, Appendix D). We find no statistically-significant differences between the Average Comfort levels in all of these conditions ($p = 0.81$).

In conditions where oxytocin inversely modulated *SocialSatisfactionRate*, we saw 100% Life Length for all agents.

Average Comfort levels were similar between the E+/I (0.79) and E−/I (0.80) groups. We did not observe a statistically significant difference between these two groups ($p = 0.27$). However, both groups outperformed the Average Comfort of the control group (E0/I: 0.64, $p < 0.01$). In these conditions, the modulation of *SocialSatisfactionRate* showed improvements to Average Comfort, regardless of the direction of modulation.

Experimental Condition	**Environment**	**Social Salience Direction**	**SocialSatisfactionRate Modulation Direction**	**Life Length**	**Average Comfort**
E0/D	Easy	None	Direct	100%	0.94
E+/D	Easy	Increase	Direct	100%	0.94
E−/D	Easy	Decrease	Direct	100%	0.92
E0/I	Easy	None	Inverse	100%	0.64
E+/I	Easy	Increase	Inverse	100%	0.79
E−/I	Easy	Decrease	Inverse	100%	0.80
C0/D	Challenging	None	Direct	70%	0.71
C+/D	Challenging	Increase	Direct	70%	0.72
C−/D	Challenging	Decrease	Direct	63%	0.68
C0/I	Challenging	None	Inverse	70%	0.71
C+/I	Challenging	Increase	Inverse	57%	0.63
C−/I	Challenging	Decrease	Inverse	58%	0.62
S0/D	Super Challenging	None	Direct	42%	0.67
S+/D	Super Challenging	Increase	Direct	43%	0.67
S-/D	Super Challenging	Decrease	Direct	39%	0.63
S0/I	Super Challenging	None	Inverse	22%	0.47
S+/I	Super Challenging	Increase	Inverse	34%	0.60
S-/I	Super Challenging	Decrease	Inverse	34%	0.59

Table 4.3 Results of Life Length and Average Comfort across all environments in Experiment 2. Values are aggregated across all experimental runs. Experimental conditions are named using the environmental condition (E = Easy, C = Challenging, S = Super Challenging), the modulation effect of oxytocin on social cues (0 = None, + = Increase, − = Decrease) and modulation effect of oxytocin on the *SocialSatisfactionRate* (D = Direct, I = Inverse). Green results highlight improved performance vs . control, red results highlight decreased performance vs. control. These results are discussed further in this section.

Challenging Environments

In the Challenging environments, when oxytocin directly modulated the *SocialSatisfactionRate* value, we observed similar performances of Life Length between the control group (C0/D; 70%) and groups where oxytocin increased social salience (C+/D, 70%). We observed a statistically-significant decrease in Life Length in the Decreased Social Salience group compared to control (C−/D, 64%, $p < 0.01$).

Average Comfort levels were similar between the control group (0.71) and groups of increased social salience (0.72), with a statistically-significant weaker performance in the Decreased Social Salience groups (0.68).

However, when oxytocin inversely modulated the *SocialSatisfactionRate* values, both the Increased and Decreased Social Salience groups were outperformed by the control group. Life Length for the control group was 70%; a statistically-significant ($p < 0.01$) difference compared to the groups of Increased social salience (57%) and Decreased social salience (56%).

Average Comfort was also similar between the latter two groups (C+/I: 0.63, C−/I: 0.62), which were statistically-significant differences compared to the control group (C0/I: 0.71).

In sum, we found no viability improvements for either of oxytocin's effects compared to control, with some decreases in performance observed when oxytocin decreased social salience (C−/D, C−/I) or inversely modulated *SocialSatisfactionRate* (C+/I) (Table 4.3).

Super Challenging Environments

In the Super Challenging environments where oxytocin directly modulated *SocialSatisfactionRate*, mean Life Length of agents was lowest in Decreased Social Salience groups (S/D, 39%) when compared to both the control group (S0/D, 43%, $p < 0.01$), and the Increased Social Salience group (S+/D, 42%, $p < 0.01$). Average Comfort was similar for both the S0/D (0.66) and the S+/D (0.67) groups, with a statistically-significant reduction in performance in the S-/D group (0.62) compared to both groups ($p < 0.01$).

We saw viability benefits when oxytocin inversely modulated *SocialSatisfactionRate* and decreased social salience (S-/I), both in terms of Life Length (34%) and Average Comfort (0.60) when compared to the control condition (S0/I).

Unlike previous conditions, performance increases were also observed in the Decreased Social Salience groups. We observed an average Life Length of 34% and Average Comfort level of 0.59. There was also no statistical significant difference between the performance of the S+/I and S-/I groups, and both groups outperformed the control conditions similarly across both viability metrics.

In summary, in these conditions, inverse modulation of the *SocialSatisfactionRate* resulted in improvements in group viability. Unlike previous results and those from the Challenging world conditions, decreasing social salience also improved group viability in one instance (S−/I).

4.5 Discussion

By accounting for a second effect of oxytocin—the modulation of internal signals of performing social behaviours—we saw comparable viability effects for when oxytocin both increased and decreased the attention given to social cues. This is contrast to the findings from Chapter 3, which saw significant advantages to viability when oxytocin increased social salience, and disadvantages when it decreased social salience.

We found that the effect of performance on these two modulatory effects were dependent upon the environmental challenge related to food availability (in particular, Challenging or Super Challenging worlds), and the direction in which the internal satisfaction signal was modulated (Direct or Inverse). However, we found that setting the internal satisfaction signal at higher values resulted in negligible differences between experiment groups, and so we take this finding into future investigations.

Taken together, these findings suggest that there is no one-size-fits-all approach to oxytocin and that the strictly-prosocial mechanisms previously observed may not always be appropriate for survival when additional context is considered. This is in line with the "interactionist" hypothesis, which suggests that the mechanisms of oxytocin are not fixed, but dynamic and dependant upon a wide array of contextual factors.

In terms of our model of social adaptation, we found how in conditions where oxytocin inversely modulated internal social satisfaction signals, infrequent social contact allowed agents to "free up" attentional resource for survival-critical behaviours, which contributed to viability performance in these conditions. As with the findings from Chapter 3, we consider that this was due to oxytocin serving as a type of "social memory" related to the frequency and recency of social interaction.

We also highlight that the physiological adaptation of the internal physiological signal (*SocialSatisfactionRate*) through social interactions also adapted the satisfaction of the second physiological variable in our model (*Energy*). This adaptation of a second physiological variable is a mechanism akin to social allostatic principles. We discuss some of these findings further below.

4.5.1 Physiological and Behavioural Adaptation via Infrequent Social Contact

Adaptation of the "Intensity" of Social Interaction

In conditions where oxytocin inversely modulated the *SocialSatisfactionRate*, we saw comparative performance between both the increased and decreased social salience groups, contrasting the findings from Chapter 3 which saw significant differences between the two groups. In these groups (C−/I and S−/I), internal oxytocin levels—released after (tactile) social interaction—played an adaptive role for agents, who could remain motivated to satisfy their survival-related *Energy* variable, and maximise their physiological satisfaction when they found an opportunity to *Groom*.

Figure 4.3 illustrates the temporal distribution of behaviours for one agent, across two simulation runs in the C+/I and C−/I groups, respectively. In conditions where oxytocin increased social salience and inversely modulated *SocialSalienceRate* (C+/I), agents performed *Grooming* frequently. This kept oxytocin at high levels, resulting in agents continually being "distracted" by the presence of other agents (as oxytocin increases

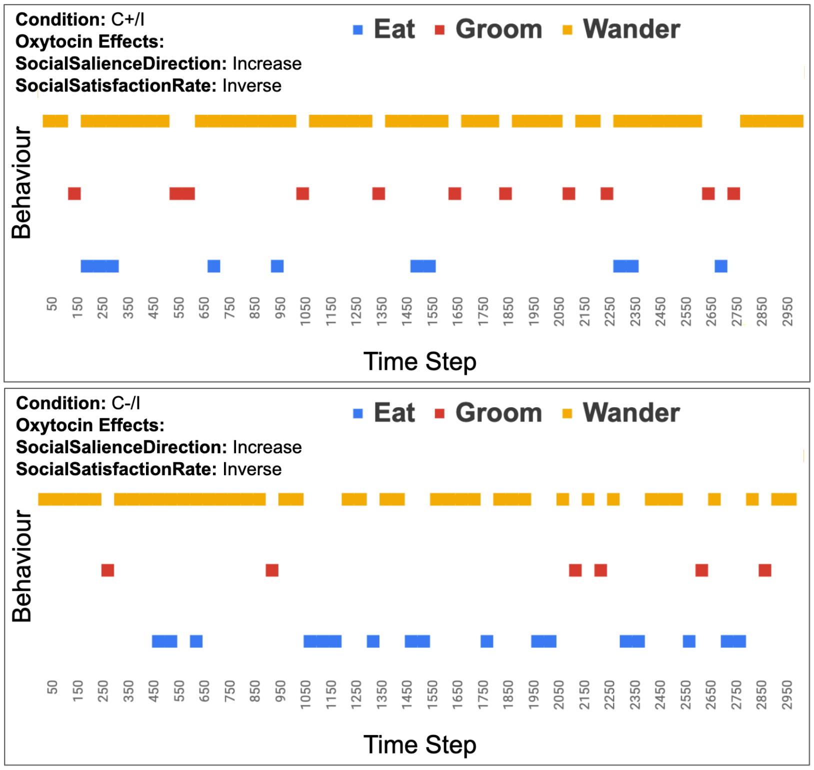

Figure 4.3 An example of the temporal distribution of behaviours of one agent when oxytocin inversely modulates *SocialSatisfactionRate*, between conditions where oxytocin increases (top) and decreases (bottom) attention to social cues. Top: When oxytocin increases social salience (+ groups), agents *Groom* more frequently, despite it not being a survival-critical behaviour. Bottom: When oxytocin decreases social salience (- groups), *Grooming* is more infrequent, but satisfies internal needs at a higher rate (Figure 4.4), allowing agents to *Eat* more frequently.

salience to social cues, Equation 4.2). As a result of these elevated oxytocin levels, the frequent *Grooming* resulted in smaller *SocialSatisfactionRates* (Equation 4.2) and therefore smaller physiological improvements (Figure 4.5, left).

In the C$-$/I groups, agents performed *Grooming* less frequently (Figure 4.3, bottom), but were able to maximise the satisfaction of the internal *SocialNeed* variable (Equation 4.2). This increased internal satisfaction then "frees up" the attention of agents, permitting them to spend longer periods of time looking for survival-related food resources (Figure 4.4, right).

Figure 4.4 shows an example of the evolution of these internal variables for one agent in a simulation run in the C$-$/I groups. When agents did not *Groom* for extended periods of time, oxytocin levels decreased, increasing the intensity of the *SocialSatisfactionRate* (Figure 4.4, A). When *Grooming* is performed (Figure 4.4, B), the internal *SocialNeed* is then satisfied by a larger value (Figure 4.4, C). The differences in the *Grooming* frequency and resultant physiological effects between these two conditions can be seen as "little-

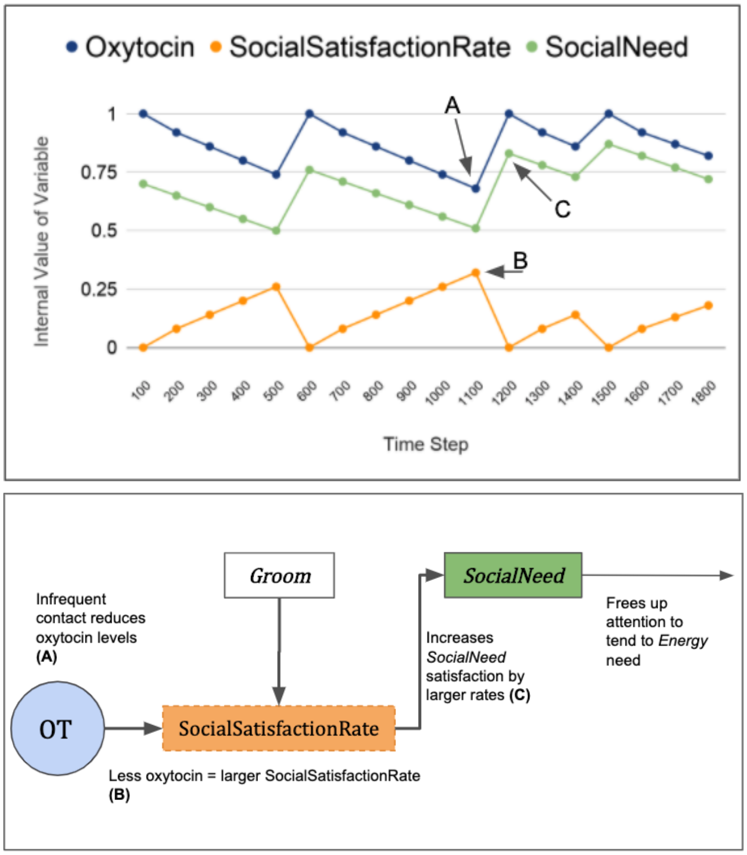

Figure 4.4 Top: An example of the evolution of internal variables over time in conditions where oxytocin decreases social salience and inversely modulates the *SocialSatisfactionRate* (C-/D). A: When the social behaviour *Groom* hasn't been performed for some time, internal levels of oxytocin decreases to reduced values. B: This lower oxytocin value inversely modulates the internal *SocialSatisfactionRate* value to a larger value. C: When *Groom* is performed, *SocialNeed* is satisfied by this higher *SocialSatisfactionRate* value. Bottom: A diagram of how the points described above are calculated in the agent model.

but-often" in conditions where oxytocin increases social salience, and "rare-but-intense" in conditions where oxytocin decreases social salience (and inversely modulates internal satisfaction signals).

Adaptation of the Satisfaction of Survival-Related Physiological Variable

The dynamic adaptation of the *SocialSatisfactionRate* also has a secondary effect, indirectly adapted the survival-critical *Energy* variable by creating a higher "upper bound" for *Energy* satisfaction (Figure 4.5, right). This allowed agents to correct their *Energy* to higher values, promoting their long-term viability, particularly in challenging environments. We note that this is a type of adaptation through social allostasis, and that the dynamic, internal satisfaction of social behaviours directly affected and adapted the survival-

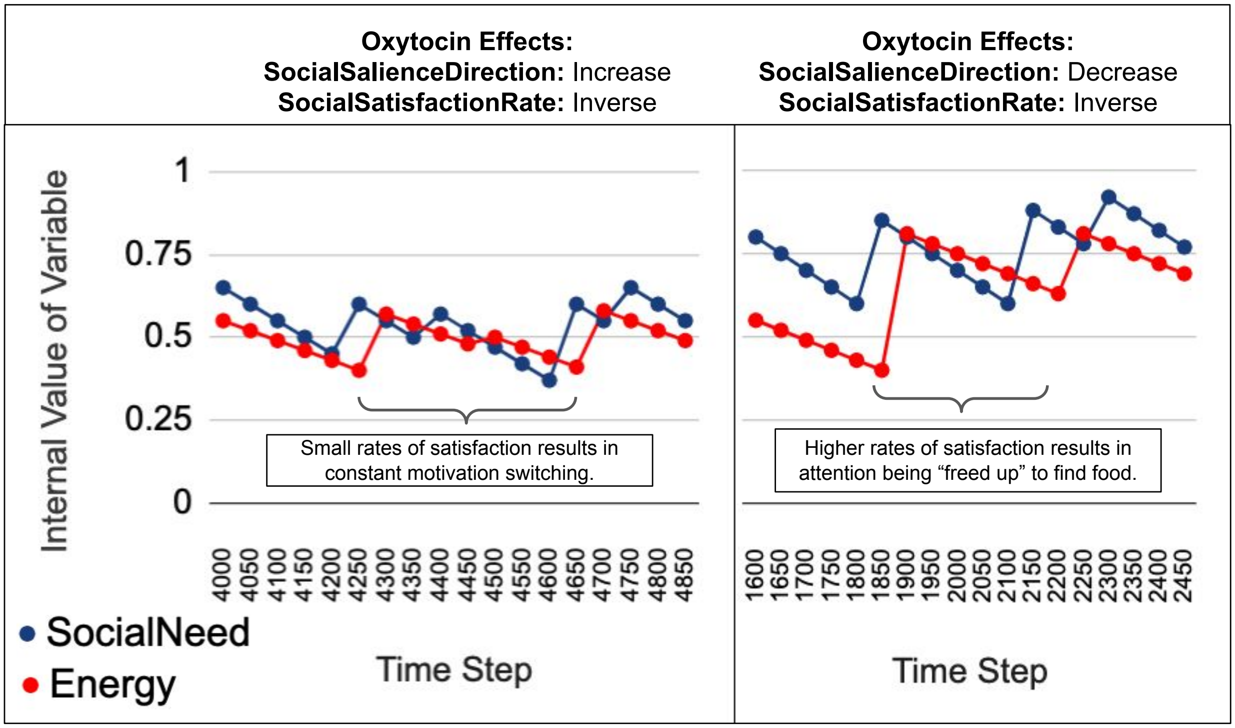

Figure 4.5 A comparison of the evolution in internal variables of one agent between conditions where oxytocin inversely modulated *SocialSatisfactionRate*, and Increased (left) and Decreased (right) *SocialSalienceDirection*. Red lines indicate an agent's internal *Energy* variable, and a blue line indicates an agent's *SocialNeed* variable. Left: Increasing Social Salience through oxytocin mechanisms results in frequent *Grooming* behaviours which satisfies *SocialNeed* by smaller values. The maximum value that *Energy* (red line) can be satisfied to is limited by the value of *SocialNeed* (dashed circles, red lines). Right: When social salience is decreased by oxytocin mechanisms, infrequent *Grooming* results in higher *SocialNeed* satisfaction, and therefore increases the value that *Energy* can be satisfied to via *Eating*.

related behaviour (*Eat*) as we describe below.

Frequent social (tactile) contact, along with agents being "distracted" by other agents due to oxytocin's effects on increasing social salience, caused agents to eat smaller amounts and resulted in a lower overall satisfaction of the *Energy* variable (Figure 4.5, left). As agents satisfy their *Energy* incrementally at each time step that they *Eat*, *Energy* was only satisfied to a value slightly above that of *SocialNeed* (Figure 4.4, left). At that point, the homeostatically-controlled agent model would determine a more urgent motivation (*Groom*) and agents would move from resources in search of other agents to satisfy their more urgent *Lonely* motivation. Conversely, when oxytocin decreased social salience, the infrequent contact resulted in higher rates of *SocialNeed* satisfaction, and therefore created a higher bound for *Energy* to be satisfied to (Figure 4.5, right).

Though counter-intuitive, in these conditions (where oxytocin increased attention given to social cues, and inversely modulated the *SocialSatisfactionRate*), frequent socio-positive interactions may have potentially limited the viability of agents, by reducing the upper limit that *Energy* could be satisfied to. In environmentally-challenging conditions, where food resources may be harder to find, this lack of adapting the internal satisfaction signal resulted in agents not appropriately maximising their *Energy* need when the opportunity arises, which may have caused a significant survival risk. These behaviours may have played a role in the reduced viability of some groups (i.e. C+/I), or contributed to the comparative performance of groups with contrasting effects (i.e. C+/I and C−/I) in this investigation.

In sum, agents adapted to their infrequent social contact, through a type of "social memory" associated with social contact recency: the less social contact an agent has, the less oxytocin is present and the more the agent maximises its opportunity to satisfy its *SocialNeed* when the opportunity arises. Oxytocin is then released in the system, meaning that any further social contact in the short-term results in a smaller satisfaction of their internal *SocialNeed* (seen in Figure 4.3 at approximately Time Step = 1400). At the same time, this maximises the rate at which the survival-critical *Energy* can be satisfied.

The comparative performance suggests that, when additional adaptive mechanisms of oxytocin are considered (in this case, its effects on internal satisfaction signals) oxytocin's effects on social salience should be more dynamic or context-dependent to maximise viability in these environments. This result may also provide further support to the interactionist perspective of oxytocin's effects on promoting social interaction and wellbeing in a social environment.

For our model, we posit that longer-term viability may be achieved through oxytocin contextual effects on social salience. For instance, oxytocin's effects on attention to social cues may depend upon current internal states—to decrease social salience when other needs are more urgent—or the perceived relationship with (or identity of) other individuals.

4.5.2 Oxytocin Levels can Serve as "Social Memory"

Like the observations from our previous chapter, we again highlight that the oxytocin mechanisms in our model provides a contextual, low-level representation of the social environment and a representation of "social memory" with respect to recency of social interaction. High levels of oxytocin corresponds to recent performance of a prosocial behaviour, which may be representative of social support in the environment.

This is further emphasised by its effects on the error-correction rate of its *SocialNeed* (Figure 4.4), where levels of oxytocin determined the intensity of internal signals, and has implications on a model of social adaptation. For instance, if social support is deemed to be high (e.g. through high levels of oxytocin over time), then social "vigilance" may be (temporarily) reduced in the priority of other needs. This mechanism was seen to emerge in our current model in some conditions (Figure 4.5), and contributed to viability performance through physiological and behavioural adaptation in those conditions.

Again, considering the low computational cost of such a "memory" effect and its potential for a mechanism of context-relevant adaptation of social behaviour, we propose that these mechanisms propose advantages in the development of a socially-adaptive model, and we take these findings into account for future investigations.

4.6 Summary

Building upon our limited findings from the previous chapter, this chapter continued the investigation into the contrasting hypotheses of oxytocin's effects on social salience, by accounting for the hormone's hypothesised role on affecting internal signals of satisfaction associated with tactile social interaction, which we frame in the context of "affective touch".

We investigated two contrasting effects related to oxytocin's modulation of this internal physiological signal—by directly and inversely modulating the intensity of this signal in our agent model, and retained oxytocin's effects on social salience from Chapter 3. We assessed the effects of these combined mechanisms on group viability, and supplemented these results with qualitative observations.

Contrary to the findings from the previous chapter, the results here suggested that oxytocin's prosocial effects on social salience does not always result in improved viability. Rather, depending upon the degree of the environmental challenge, our results found that agent performance was improved by different modulatory effects of oxytocin when its effects on internal signals were accounted for. In some conditions, viability was improved in absence of (one of) oxytocin's modulatory effects, suggesting that there may be situations in which oxytocin's effects can be "switched off" to maximise viability. Therefore, to maximise viability, the hormonal mechanisms that we had accounted for (the effects on social salience and internal signals) should be context-dependent. This view is also shared by the interactionist perspective of oxytocin's effects [145].

Based on our observations, we propose that these contextual effects of oxytocin may be impacted by the social environment. Prior relationships or familiarity with others (including affective or attachment bonds), the quality of those relationships, and the challenges associated with the social environment (such as hierarchy, social stability, or threats of predators), are examples of social contexts where oxytocin's effects are flexible and, at times, appear to be paradoxical [167]–[169].

We have identified that such a social context does not exist in our current experimental design. We address this limitation in the next chapter, by further investigating the context-dependent effects of oxytocin when additional social context—that of social hierarchy and the presence of bond partners—is taken into consideration.

Chapter Five

Contextualising the Effects of Oxytocin: Affective Bond Partner Preference in a Rank-Based Society

5.1 Introduction

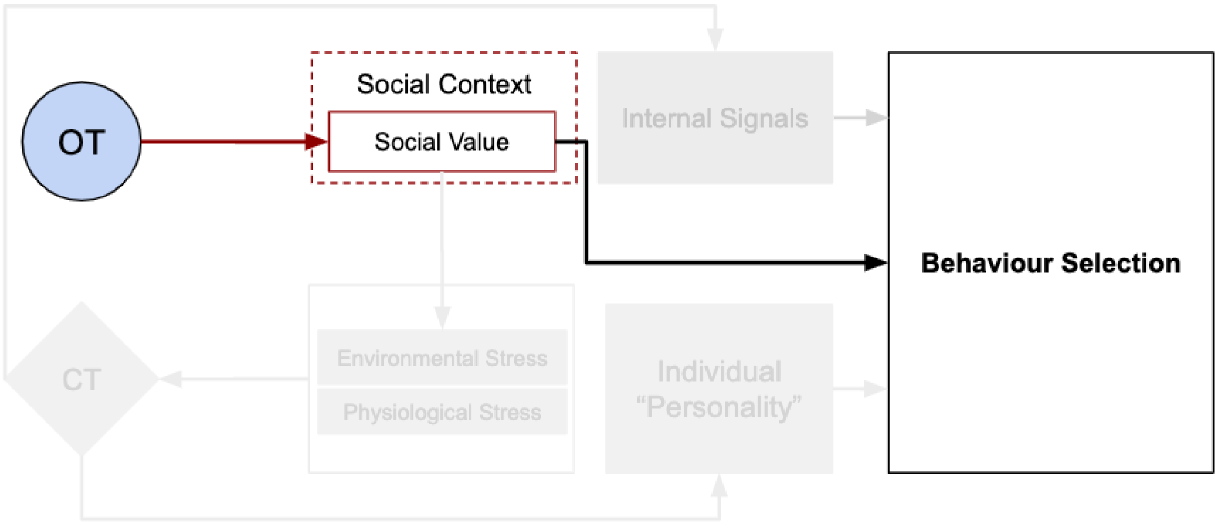

Figure 5.1 A high-level view of the effects investigated in this chapter. OT = Oxytocin. In this investigation, we remove oxytocin's effects on internal signals that we had investigated in Chapter 4. Hormonal effects on internal signals are reintroduced in Chapter 6. A more granular view of the model can be seen in Figure 5.2

In (Chapter 4), we found how accounting for a second potential effect of oxytocin—its role on modulating internal signals of social behaviours—as well as its effects on affecting the attention to social cues, resulted in more context-dependent effects on group viability. This was contrary to the investigation in Chapter 3 which found clear support for the "prosocial" approach towards oxytocin's mechanisms. We suggested that to further understand the potential adaptive role of oxytocin on social behaviours and overall wellbeing, these contexts must be investigated further. Our results also supported the interactionist approach of oxytocin mechanisms, which states that its effects on social behaviour are not universal, but dependent upon several environmental and social contexts.

In this chapter, we consider that oxytocin's effects on the attention given to social cues are not universal as we had considered in Chapter 3 and 4, but that its effects on social attention (and value) are associated with the identity of other individuals in the environment. Specifically, we focus on the hypothesis that, when presented with several potential interaction partners, oxytocin promotes preferential attention—and valence—towards familiar or affective bond partners in a social environment, and that this in turn underpins the formation of social relationships which promotes (survival-related) behavioural adaptation.

To address the lack of social context in previous investigations, we investigate this hypothesised effect in a small, rank-based society using several combinations of social bond partners related to different dominance ranks. We test how the wellbeing and social interactions of a small group of agents is affected by oxytocin's hypothesised effects on preferential value of bond partners. To minimise further experimental complexity, we perform these experiments in a single, physically-challenging, world condition.

Our results find that the oxytocin's effects on promoting bond partner preference for prosocial interactions provided a significant advantage to the wellbeing of individuals with bonds—and, in some cases, for the wider society—compared to strictly rank-based interactions. We find that the "trade-off" between rank-based and affect-based social interactions played a role in promoting the wellbeing of bonded individuals, resulting in some biologically-plausible social interactions. Our observations found behaviours akin to allo-parental and altruistic behaviours, which promoted survival of subordinate individuals, but also the emergence of distrust or social amnesia between bond partners in some cases.

We highlight two key additions to the model in this chapter. The first, the *Dyadic Strength Index* is a flexible measure of affective bond strength between individuals that addresses limitations of previously-used measures of dyadic relationships in artificial agents. The second, the *Social Assessment Component*, accounts for the perceived identity of interaction partners prior to behaviour execution, by accounting for social rank and bond partner presence. We propose this as a computational abstraction of "trade-off" seen in some biological species.

5.1.1 Social Hierarchy, Dominance Rank Advantages and Social Relationships

Social hierarchies and dominance rank are natural by-products of life in a society. Dominance rank in these societies is oftentimes associated with the most fitness benefits [170]: including priority access to resources (i.e. food) as well as greater or easier access to grooming partners, resulting in overall health and longevity of life [108], [171], [172]. For most (competitive) social species, obtaining and maintaining a higher social status is an important, survival-related type of adaptation.

However, throughout the evolution of all social species, adaptation and survival has also been based on the ability to form and maintain meaningful relationships with others [13]—to cooperate and find food, protect each other from predators and compete with other groups for resources. Much like dominance rank, the formation and maintenance of affective social bonds for social individuals plays an important role for survival in several species [18]–[20].

For individuals living in a rank-based society, the ability to control behaviours based on perception of relative rank, along with the presence of close social bonds, appears to be a crucial aspect of survival in a group setting [31], [170]. A social behaviour (such as grooming) can be costly in terms of both time and energy, taking away from other survival-critical behaviours (such as locating food). Therefore, selecting a social partner with whom to perform social behaviours is rarely random, but an active selection process that may provide future potential benefits [173].

Oxytocin is considered to play a significant role in bond formation and maintenance of socially-affective bonds, from mother-infant bonds [174] to dyadic bonds of unrelated individuals [175]. As [176] mentions, oxytocin levels may be seen as a "biomarker for the valence and strength of the relationship between interaction partners", which has been suggested to have a positive-feedback loop on partner selection and the maintenance of affective relationships.

This feedback loop, possibly via its effects on the reward system [177] or the formation of social memory [178], is hypothesised to facilitate additional partner selection in the future and therefore release further oxytocin. As [174] mentions, this affective feedback loop is potentially a key mechanism that underpins the formation of (affective) social bonds.

This effect of oxytocin on partner preference can be considered as one of the contexts where oxytocin's effects on social attention and value are not universal. As a process that drives both physiological [179] and behavioural [180] adaptation, the formation, maintenance and selection of such affect-based social bond partners for social interactions plays a fundamental role in the adaptive process of social allostasis [14].

5.1.2 Aims & Hypothesis

We continue our investigation into the context-dependent effects of oxytocin on the processing of social cues. We aim to investigate one of oxytocin's hypothesised contextual effects on social cues—that it specifically promotes the valence of, and preference towards, affective bond partners for social interaction—and study the effects of this mechanism on both individual and group wellbeing, and social interactions

Accounting for this mechanism in our agent model, we undertake this investigation in a rank-based society of agents, using several social bond combinations related to hierarchical rank, in a physically-challenging environment of limited food availability.

We hypothesise that oxytocin-mediated preference towards affective bond partners for social interactions will significantly benefit the viability of socially-bonded individuals.

We use this investigation to further understand whether oxytocin's contextual effects on social cues can provide viability advantages in a socially-challenging environment. Our results may help us further understand some of the potential mechanisms that can underpin affect-based social relationships, and can facilitate our development of an adaptive agent model.

5.2 Extensions to the Agent Model

To investigate the contextual effects of oxytocin, we higlight several changes and extensions to our agent model and modelling approach. We discuss and justify them here.

Firstly, we do not account for its modulatory effects on internal signals in this investigation, as we had investigated in Chapter 4. Secondly, we readdress the fixed rates of oxytocin secretion after behaviour execution, aligning to more contextual, biologically-inspired rates of secretion—inspired by chimpanzee societies—which we discuss in Section 5.2.1.

Extending on the homogeneous society of agents in Chapters 3 and 4, we endow each of the six agents with a different dominance rank (Subsection 5.2.4). We then introduce the notion of a "social bond": an abstraction of a mutually, socio-positive relationship between two dyads. To model this, we formulate a real-time measure of bond strength called the Dyadic Strength Index. We discuss this further in Subsection 5.2.2.

Finally, we model oxytocin's contextual effects on modulating the "valence" of bond partners through the Social Assessment Component (SAC, Subsection 5.2.5). The SAC accounts for the perceived identity of social stimuli—by accounting for difference in dominance rank and the presence and strength of an affective social bond, which is modulated by oxytocin levels. These extensions can be seen in Figure 5.2, with a detailed view of the Social Assessment Component seen in Figure 5.3.

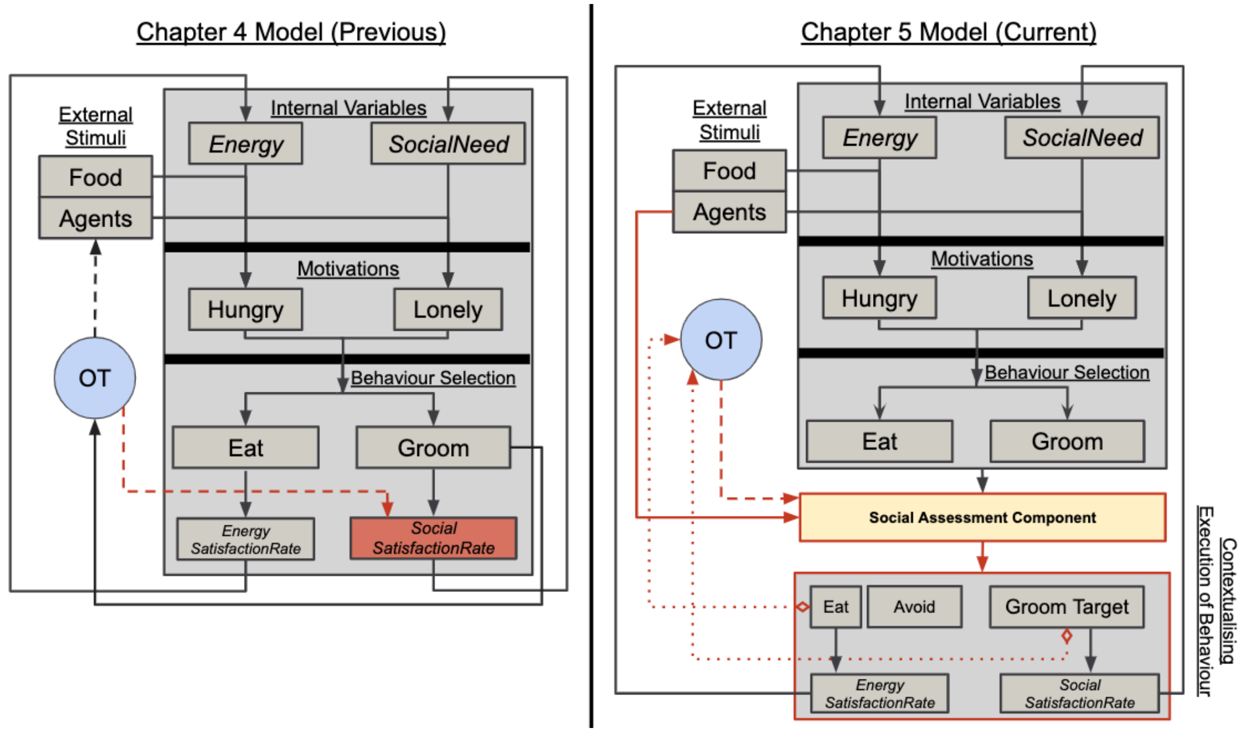

Figure 5.2 Illustration of the agent model used in this investigation (right) compared to the previous chapter (left). Extensions to the previous model are denoted in red. Oxytocin's effects on *SocialSatisfactionRate* are removed in this investigation. The current model introduces contextualisation of behaviour execution through the *Social Assessment Component* (Section 5.2.5). Oxytocin modulates partner bond preference prior to behaviour execution. Dashed line indicates modulation of values. Dotted lines to OT indicate context-dependent effects (details in Table 5.1). A more detailed view of the Social Assessment Component can be seen in Figure 5.3.

5.2.1 Chimpanzee-Inspired Hormonal Mechanisms

As we aimed to investigate oxytocin's effects in a biologically-inspired social context, we address our approach to the modelling of the hormonal mechanisms. Mechanisms in previous chapters were inspired by oxytocin effects found in species such as rats [123], sheep [124] and voles [128]. While this modelling approach was appropriate for those investigations, it limited our findings as the social behaviours and organisations of these species are not complex enough in their range of social behaviours to investigate oxytocin's diverse effects [106].

Here, we focus on non-human primate societies—in particular, chimpanzee societies—which we considered to be a suitable animal model to inspire our modelling approach. This rationale lies in their physiological similarities to humans, their diverse types of social organisation, the degree of physical and social challenges, and their flexible approach towards social interactions [112]. As [73] states, non-human primates (and their societies) provide good study objects given their their relative simplicity compared to humans, but a complexity that makes their societies interesting. Primate societies have also been used to inspire previous social simulation models [72]–[75] and so we find these models to be suitable models to abstract hormonal mechanisms from.

Here, we list some recent specific observations from chimpanzee societies that guided changes to our agent model. Further discussion of these specific studies can be found in Appendix E.

Firstly, that oxytocin levels are seen to increase after grooming between bond partners than with non-bond partners [109]. Secondly, that food sharing (independent of bond partner status) is seen to increase oxytocin levels in sharers [31]. Thirdly, that chimpanzees show flexibility in partner selection for grooming [142] and food sharing [180] based on perceived rank and researcher-defined bond "strength". These authors and others [109], [174] suggested that oxytocin systems may underpin the reinforcement of such socially-affective decision-making.

Combined, these studies suggest that for species that live in such complex societies, the release of oxytocin, as well as flexibility in social decision-making, is dependent upon several factors. We use these studies as inspiration for the changes to the agent model in this chapter.

Although a wider review of chimpanzee studies is beyond the scope of this dissertation, we remain aware of the fact that there are a significant amount of studies related to chimpanzee (and primates in general) societies and behaviours which may propose different mechanisms. We are inspired by these particular studies due to the similarities of social and environmental contexts of those investigations [109], [142], [180] and our own investigation. For a full review on studies regarding oxytocin and chimpanzee behaviour, see [176].

5.2.2 Modelling Affective Bond Strength with the Dyadic Strength Index

To model an affective bond between agents, we have developed a measure to model the strength of affective social bonds between agents, which we call the Dyadic Strength Index (DSI). The DSI is a real-time measure of social bond strength between dyads, updated as a function of intra-dyadic social interactions as well as the lack of interactions over time. In this way, the DSI is able to capture short-term dynamics of social bonds.

Our approach is motivated by recent work performed by Kulik [181] who proposed a dynamic measure of bond strength using a measure called the *Dynamic Dyadic Sociality Index* (DDSI) for primate societies, which was initially proposed to address limitations of social bond measures (such as the Composite Sociality Index (CSI) [182]) or proxy measures such as the types, frequency, duration and identity of social partners [183] used in ethology.

Though our model is motivated by biological societies, we found these measures of social relationships to be unsuitable and that they presented two key limitations which we aimed to address. Firstly, that these measures do not consider the temporal changes in dyadic relationship quality resulting from short-term interactions (or lack of interactions). Secondly, that the magnitude of effect that social interactions can have on bond strength is not a constant static value, but that it should be a contextual effect.

Despite Kulik's attempt to address some limitations from earlier measures [182], [183], the DDSI still fails to address the limitations regarding the contextual effect of social interactions on social bond strength.

One advantage for using artificial societies over biological systems is the ability to capture real-time data points non-intrusively. Observational periods do not need to be arbitrarily predetermined and dyadic bond strength can, and should, be captured and calculated in real-time. In this way, researchers can precisely measure—and model—all temporal changes in social relationships to more accurately relate emergence of behaviours with (short-term) changes in relationship quality.

Measures of social bond strength in artificial societies are limited. Recent work [184], [185] proposed measures of social relationships, but does not currently account for temporal changes in social bond strength after a social interaction has taken place. In terms of capturing temporal dynamics, measures of dyadic interactions in societies have been proposed by accounting for rank differences [72], [75], [77], and via through modifications of the Elo ranking method [186], [187]. However, these rule-based approaches towards dyadic interactions are limited to dominance-based interactions and do not currently consider affective relationships between individuals.

We build on this work to address these limitations and propose that the DSI can be used as a model and measure for future models for affect-based social relationships and interactions. The DSI between agents A and B works as follows:

Firstly each affective bond between agents A and B is initialised at time $t = 0$ with an initial value (2):

$$\mathrm{DSI}_0 = 2 \tag{5.1}$$

At each time step t, in an absence of social behaviours, the DSI experiences a small loss rate (DSI_λ).

This loss can be either a static value or an exponential decay rate. In our case, we use the latter (0.9997):

$$\begin{aligned} \mathrm{DSI}_t &= \mathrm{DSI}_{t-1} \times \mathrm{DSI}_\lambda \\ \mathrm{DSI}_t &= \mathrm{DSI}_{t-1} \times 0.9997 \end{aligned} \tag{5.3}$$

The *DSI* can be strengthened or weakened through any number of socio-positive or socio-negative behaviours in different ways, Our current approach uses two behaviours (*Eating* and *Grooming*). Each behaviour is assigned an associated weighting factor (W). The weighting for each behaviour can either be a predetermined value (as is the approach taken by [181]), or a dynamic value based on factors such as frequency, the subjective "cost" of performing a behaviour, or the direction of an interaction (i.e., whether a higher or lower-ranked individual performed the behaviour).

As initially proposed by [181], this value is multiplied by a scaling value called (max_{gain}), which sets the default "magnitude" associated with each social behaviour. Therefore, the contribution of each behaviour i is calculated using:

$$DSI_{gain} = max_{\mathrm{gain}} - (max_{\mathrm{gain}} \times \mathrm{DSI}_t \times W_i) \tag{5.4}$$

In our approach, we capitalise on the capabilities of our embodied agent model. Each of the two behaviours (*Groom* and *Eat*) from the Action-Selection Architecture returns an intensity value associated with it (calculated using Equation 2.4a). We use this value to weight the contribution towards changes in bond strength. In other words, the strengthening (or loss) of bonds is directly associated with the intensity of the behaviour being performed at any given time. In addition, the gain in bond strength is inversely associated with the current bond strength: the weaker the bond, the larger the potential gain:

$$DSI_{gain} = max_{\mathrm{gain}} - (max_{\mathrm{gain}} \times \mathrm{DSI}_t \times \alpha b_i) \tag{5.5}$$

$$\mathrm{DSI}_{t+1} = \mathrm{DSI}_t + DSI_{gain} \tag{5.6}$$

where b_i is the intensity of the winning behaviour b and α is a value to regulate b_i. In this way, the same behaviour can have different effects on dyadic bond strength at different points in time, rather than a static value for each interaction. This addresses a limitation of both the CSI [188] and the DDSI [181] by modelling bond strength changes in real-time and as a function of interactions, rather than a subjectively-weighted (albeit expertly-weighted) and retrospective approach. In later chapters (Chapters 7 to 9), we add a further dynamic element to bond strength changes, by also accounting for hormonal effects of dyadic interactions.

Unlike the multi-dimensional approach to social relationships suggested by [183], we do not consider the directionality of the interaction or take into account hierarchical rank within the calculation of our bond strength measure. Instead, we consider these features separately, in the second extension discussed in this chapter (called the *Social Assessment Component*). In this way, we aim to keep this modelling of affective bond strength flexible enough to be adapted in future work or alternative embodied social decision-making models.

Finally, in our investigation, this *DSI* value is not only a measure of bond strength during the investigation, but is also used as part of the agent's decision-making process when calculating the relative "value" of an (affective) social partner. We discuss this further in Section 5.2.5.

5.2.3 Social Dominance Rank

In our previous investigations (Chapters 3 to 4), all agents were homogeneous, representing a "flat" social hierarchy. Here, we introduce the notion of social dominance rank: an approach also taken by [72], [75], [77]. Taking inspiration from social hierarchies in the real-world, higher dominance rank grants an agent easier access to food resources and are also more desirable to be selected by other agents as a grooming target.

An individual's dominance rank in the society is represented as a value between 1 and 6. This rank can then be detected by other agents during the execution of their two behaviours. This is discussed in Subsection 5.2.5.

5.2.4 Affective Social Bonds

We consider an affective social bond between two individuals to represent a pre-existing stable, mutually-positive relationship between the two agents. In our model, this social bond is represented using a boolean flag to denote the presence of a bond between two agents (1 if a bond exists, else 0), and the strength of this bond is measured using the Dyadic Strength Index (DSI) (Subsection 5.2.2) between 0 and 2.

5.2.5 The Social Assessment Component

The Social Assessment Component (SAC) is an extension to the Action-Selection Architecture (Figure 5.2) that accounts for the perceived identity of other agents prior to (social) behaviour execution, by accounting for the difference in dominance rank, and the presence—and strength—of an affective bond between agents. In this way, social behaviour execution is contextual, based on the perceived affective relationship with another agent.

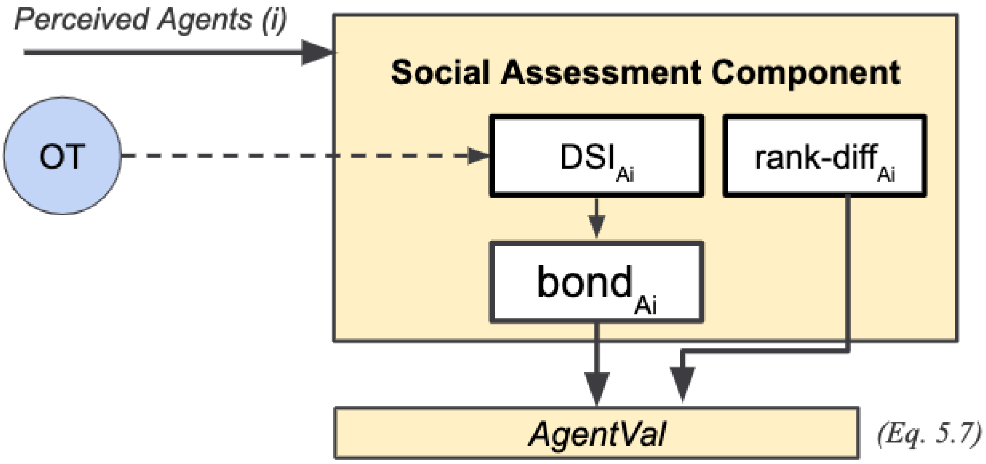

Figure 5.3 Detailed view of the Social Assessment Component used in our model, developed as an abstraction of the "trade-off" between rank and affective bond partners [31]. For each perceived agent (i), the Social Assessment Component calculates the rank difference between the current agent (A) and (i) (rank-diff_{Ai}), the presence of an affective bond between the agents (bond_{Ai}) and the strength of this affective bond (DSI_{Ai}). Oxytocin (OT) is used to modulate the valence of affective bond strength. The resulting value is the *AgentVal* used to determine behaviour execution. This calculation can be seen in Figure 5.7.

In previous iterations of our model, agents driven to execute one of their two behaviours (*Eat* or *Groom*) would move to a salient resource (any *Food* or *Agent* in its vision) and perform the winning behaviour. As all agents were homogeneous, they made no distinction as to which agent they would select for *Grooming*. Additionally, the behaviour of passive "food sharing" (i.e. multiple agents eating at a food resource) was not addressed. It was possible for multiple agents to consume the same resource at the same time.

As we moved towards a more biologically-plausible modelling of social structures and contextual hormonal release, we considered that biological social agents do not execute prosocial behaviours with random individuals, but that the choice in interaction partner is an active selection process dependent upon several social contexts [31], [170], [189], [190]. Here, we use the SAC as an abstraction of the hypothesised "trade-off" between rank and affect-based interactions suggested by [31], [170]. Other rule-based approaches have also used rank difference and affective relationships for social behaviour selection in artificial societies [75], [77]. We suggest that the SAC proposes a more dynamic modelling of this potential mechanism: modelling the relationship between rank and social bonds as a "trade-off", rather than a rule-based approach.

The SAC is used to determine two aspects of behaviour execution:

- (a) whether to approach (and share), or avoid an occupied food resource (when the winning behaviour is *Eat*)

- (b) which potential partner to select for grooming (when the winning behaviour is *Groom*).

The SAC takes into account the relative dominance rank and the presence and strength of any bond (DSI) with each agent it perceives. The SAC then modulates the valence of an affective bond by the level of oxytocin within an agent's physiology. Specifically:

$$AgentVal = \text{rank-diff}_{Ai} + (\text{bond}_{Ai} \times (\text{DSI}_{Ai} \times \text{OT}_{A})) \tag{5.7}$$

Where $bond_{Ai}$ is a boolean flag denoting the presence of a social bond between agents A and i (1 when a bond exists, else 0). DSI_{Ai} is a value denoting the strength of the social bond between these agents and takes a value between 0 and 2.

$rank\text{-}diff_{Ai}$ describes the normalised difference in rank between agents A and i respectively. Agents are assigned a rank (1–6), which is normalised between 0–1 in 0.25 increments (0, 0.25, 0.5, 0.75, 1), with 0 being the lowest-ranked agent and 1 being the highest. Agents can observe the rank of other agents in its vision. Using these normalised values, $rank\text{-}diff_{Ai}$ is calculated using:

$$\text{rank-diff}_{Ai} = \text{rank}_{A} - \text{rank}_{i} \tag{5.8}$$

A negative value indicates that the perceived agent is higher-ranked, and a positive value indicates that the perceived agent is lower-ranked. As no agents hold the same rank in our model, *rank-diff* can never equal 0. *rank-diff* takes the range -1 to 1. In absence of any social bond, *AgentVal* was determined by the rank difference between agents A and i.

Oxytocin is used to promote the valence of social bonds, by modulating the strength (DSI) of a bond between two agents. This is an abstraction of its effects on promoting partner preference [109], [128], [129], [176]. A higher level of oxytocin therefore promotes partner preference and can potentially offset difference in rank (Equation 5.7). As we look to magnify the effects of the SAC mechanisms in this investigation, the upper bound of oxytocin was increased from 1 to 2.

In sum, agents do not simply execute their winning behaviour of *Eat* or *Groom*, but account for the presence and quality of affective relationships through the SAC. We discuss how the SAC is used during each of these behaviours below. An illustration of this behaviour selection is seen in Figure 5.4.

SAC 1: Food Approach or Food Avoidance

Agents looking to eat at a food resource take into account any other agents that may be at the food resource (which we define as the "food owner"), and determines whether agents will *Approach* that food resource, or

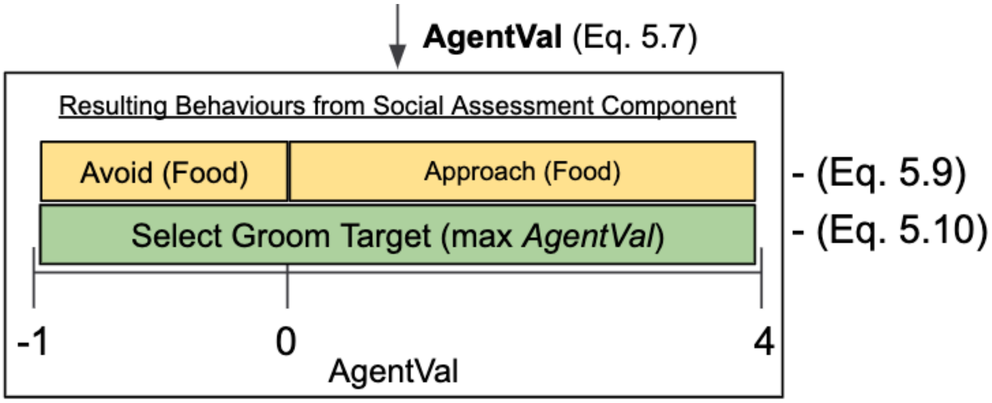

Figure 5.4 Illustration of the different behaviour selection as a result of the *AgentVal* calculation. *AgentVal* returns a value of between -1 and 4. Yellow behaviours indicates behaviours when the winning behaviour is to *Eat* (Equation 5.9). *Groom* is performed on the agent that returns the largest *AgentVal* (Equation 5.10).

Avoid it by turning around.

In this investigation, food owners maintain a passive approach towards other agents approaching that same food resource. Using the *AgentVal* calculation above (Equation 5.7), agents will approach and eat at a food resource if *AgentVal* is positive (or 0, which indicates that a food resource is not occupied). A negative *AgentVal* value—a result of the absence of a bond and the presence of a higher-ranked agent—results in agents avoiding that occupied resource. Specifically:

$$\text{ApproachFood} = \begin{cases} \text{Approach} & \text{if AgentVal} \geq 0 \\ \text{else Avoid} & \end{cases} \tag{5.9}$$

In the absence of a social bond between them, a lower-ranked agent cannot approach and access a food resource occupied by a higher-ranked agent (*AgentVal* will always be negative). In addition, the highest-ranked agent at a food resource becomes the “food owner”, regardless of whether or not they were the first to that food resource.

SAC 2: Grooming Target Selection

When presented with a number of agents, the SAC selects which other agent to execute the *Groom* behaviour with by accounting for both rank, bond partner status and oxytocin levels (Equation 5.7, the *AgentVal*). The agent returning the highest *AgentVal* is selected as the *Grooming* target:

Behaviour	**Bonded?**	**OT Effect**	**DSI Effect**	**Stimulus**
Eat	-	-	-	Food
Eat	Bond	$+0.5$	$\alpha = 0.5$	Food
Eat	Non-Bond	-	-	Food
Groom	Bond	$+0.2$	$\alpha = 0.5$	Agent
Groom	Non-Bond	-	-	Agent

Table 5.1 The different context-dependent rates of change for oxytocin (OT) and bond strength (DSI). Oxytocin is only released during interactions with bond partners at different rates for each behaviour. The weighting α in the DSI column is the weighted value used within the DSI calculation discussed in Section 5.2.2. Justification of the different α values are discussed in Subsection 5.2.1.

$$\text{Groom}_T = \max[AgentVal_i] \tag{5.10}$$

where $Groom_T$ is the target agent selected to perform *Groom* with. In absence of any social bonds, the highest-ranking agent is selected as the target by default (Figure 5.4).

5.2.6 The Role and Release of Oxytocin Mechanisms

In this investigation, oxytocin modulates the valence of bond partners within the Social Assessment Component (Figure 5.3), moving from the uniform modulation of all social cues in previous chapters. Oxytocin can be in the range 0 to 2, and experiences a decay rate of -0.005 at each time step. It is released within an agent's physiology during the execution of each of the prosocial behaviours—food "sharing" with any other agent, and grooming with a bond partner—at independent rates (Table 5.1). The dotted lines on Figure 5.2 (right) also indicate these contextual effects.

As discussed in Subsection 5.2.1, these contextual secretion rates and modulatory effects on partner preference are grounded in biological literature [31], [142], [180], [191] in order to develop a more biologically-plausible model.

5.3 Experimental Set Up

We conducted our experiments using a society of six agents, each with a different social rank, across five different conditions related to the combination of socially-bonded individuals. In each condition, three agents were socially bonded (referred to as the Bonded group) (Section 5.2), and three remained without

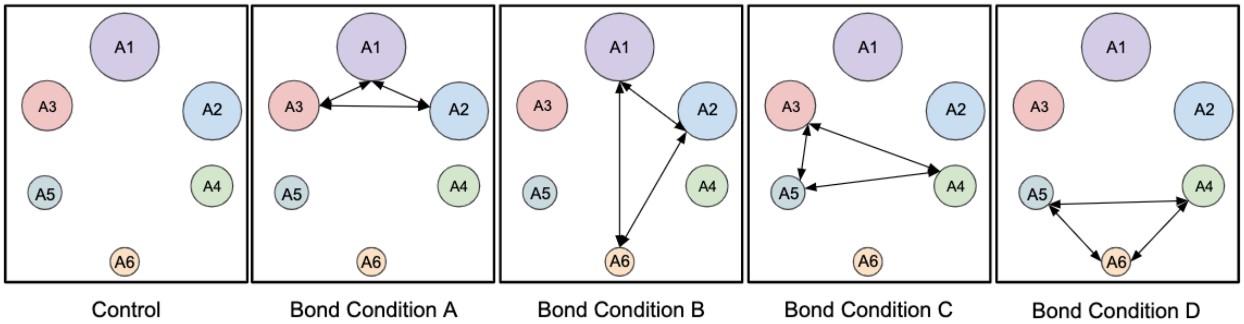

Figure 5.5 Illustration of the four different bonded conditions investigated in this experiment. Agents are ordered by rank: A1 is the highest-ranked agent; A6 is the lowest-ranked agent. Lines indicate the three agents that form a group in each condition (called the "bonded" agents), with the other three "unbonded".

social bonds (referred to as the Unbonded groups). Our control condition was set up with all six agents without social bonds. Although intuitively-selected, we took inspiration from biological societies to select combinations of affective bonds where we would expect to see significant differences. Figure 5.5 shows the different agent combinations (denoted by their rank number A1-A6).

We conducted our experiments using a single world condition consisting of two food resources, which we consider to be a physically-challenging environment (Figure 5.6).

Thirty simulation runs were performed for each of the five bond combination conditions for 20,000 time steps each. This resulted in 150 simulation runs, which equated to approximately 10 hours of experimental runs in real-time. Data was captured at each time step for each agent and aggregated across for all simulation runs.

We report results across all simulation runs using our three viability indicators (Life Length, Average Comfort, Physiological Wellbeing). Results are reported at both the individual and group (Bonded vs. Unbonded) level. Each social interaction, along with the actor and recipient of each interaction, was recorded. We used this data to calculate the frequency and direction of social behaviours between agents. We complemented all of this quantitative data with personal observations of the artificial society.

Statistical significance tests were performed on the three viability indicators. Between-group statistical testing was conducted using one-way ANOVA, with significance declared at $p = 0.05$.

Figure 5.6 Screenshot of the world condition used in this investigation. Bond Condition A shown (A1-A2-A3). Agents are doughnut-type shapes. Food resources are represented by yellow spheres. Agent size denotes rank: the larger an agent, the higher their rank. Dashed line between agents indicates a bond between them. For visualisation purposes, this line becomes more translucent as the bond strength decreases.

5.4 Results

5.4.1 Viability Indicators

Figure 5.7 shows the aggregated *Life Length* for bonded and unbonded agents. Figure 5.8 shows the results of the viability indicators for each individual agent. All results are averaged across all simulation runs.

Across all conditions, *Life Length* of Bonded agents matched or outperformed Control (Bond Condition A: +16%, $p = 0.04$; Bond Condition B: +11%, $p = 0.03$; Bond Condition C: +9%, $p = 0.13$; Bond Condition D: −1%, $p = 0.71$). We also note some detrimental impact to *Life Length* of Unbonded agents (Bond Combination A: -7%, $p = 0.02$; Bond Combination B: -6%, $p = 0.04$) compared to control.

In the Control condition, we see the two highest-ranked agents (A1 and A2) experience the highest rates of *Life Length* (A1: 55%; A2: 54%), *Average Comfort* (0.72 and 0.63 respectively) and *Physiological Wellbeing* (0.80 and 0.79 respectively) of all agents. Despite being the third-most dominant agent, A3 experienced poor viability (*Life Length*: 39%, *Average Comfort*: 0.54, Physiological Wellbeing: 0.65). For all other agents, viability performance was strongly correlated with social rank ($r = 0.742$).

In Bond Condition A (A1-A2-A3 bonded), *Life Length* of the Bonded group improved by 16% compared to when they were not bonded in Control. All agents in the bond saw improved *Life Length* (A1: +7%, A2: +4%, A3: +36%). There was a 36% difference in survival between Bonded and Unbonded groups ($p <$

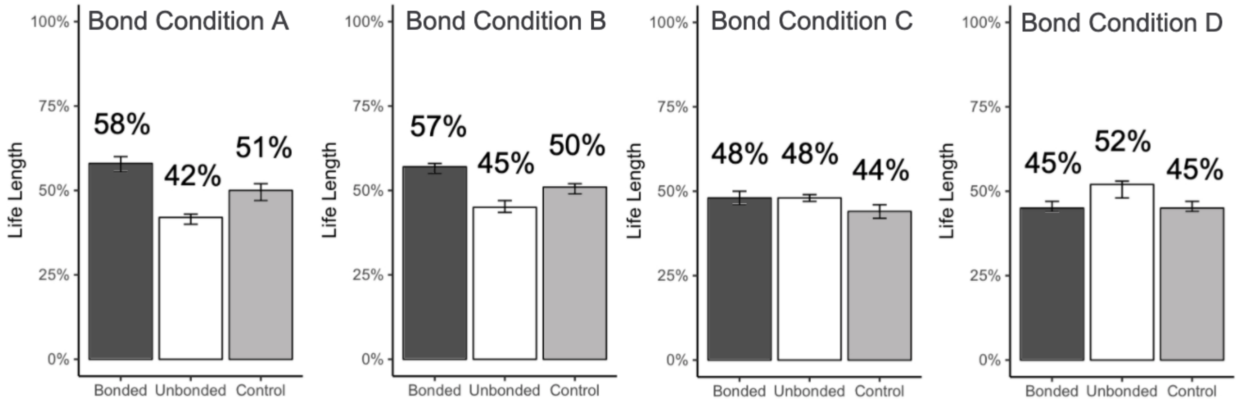

Figure 5.7 Mean Life Length of bonded and unbonded agents compared to control for each of the different bond combinations. Control results is the mean Life Length of the bonded agents when they were unbonded in the control group. Error bars indicate standard error of the mean. Life Length for individual agents can be seen in Figure 5.8.

0.01) in this condition. We note further improvements in both *Average Comfort* and *Physiological Wellbeing*, particularly for the lowest-ranking agent in this bond (A3: *Average Comfort*: +16%, *Physiological Wellbeing*: +18%).

In Bond Condition B, we note a +19% increase in *Life Length* for A6 (the least dominant agent) when bonded with the two highest-ranked agents. A1 (+8%) and A2 (+9%) also reported increases in *Life Length* in this condition. Despite being unbonded, we also see an increase in A3's life length (+20%) in this condition. Though *Average Comfort* and *Physiological Wellbeing* slightly improved for dominant individuals, the subordinate A6 experienced significant viability benefits in these conditions (*Average Comfort*: +38%, *Physiological Wellbeing*: +43%) compared to its performance in Control (Figure 5.8, Condition B). In summary, affective bonds provided significant benefits for all agents and, in particular, the lowest-ranking agent; outperforming higher-ranking agents.

Bond Condition C saw overall group survival improve by 9% vs. Control. The least dominant agent, A6, saw an 18% reduction in *Life Length* vs. the Control condition, whilst improvements were seen for the lowest-ranking bonded agent (A5: +12%). We consider A6's reduced *Life Length* to be a result of the increased cooperation between the closely-ranked agents (A3-A4-A5). Both *Average Comfort* (0.67) and *Physiological Wellbeing* (0.82) significantly improved for the Bonded group compared to control, with similar viability to higher-ranked agents (*Average Comfort*: 0.67, *Physiological Wellbeing*: 0.79).

In Bond Condition D (A4-A5-A6), aggregated bonded agent life was slightly lower than their non-bonded performance in Control: a result of a decreased *Life Length* of A4 (-4%) and A5 (-6%). The lowest-ranking

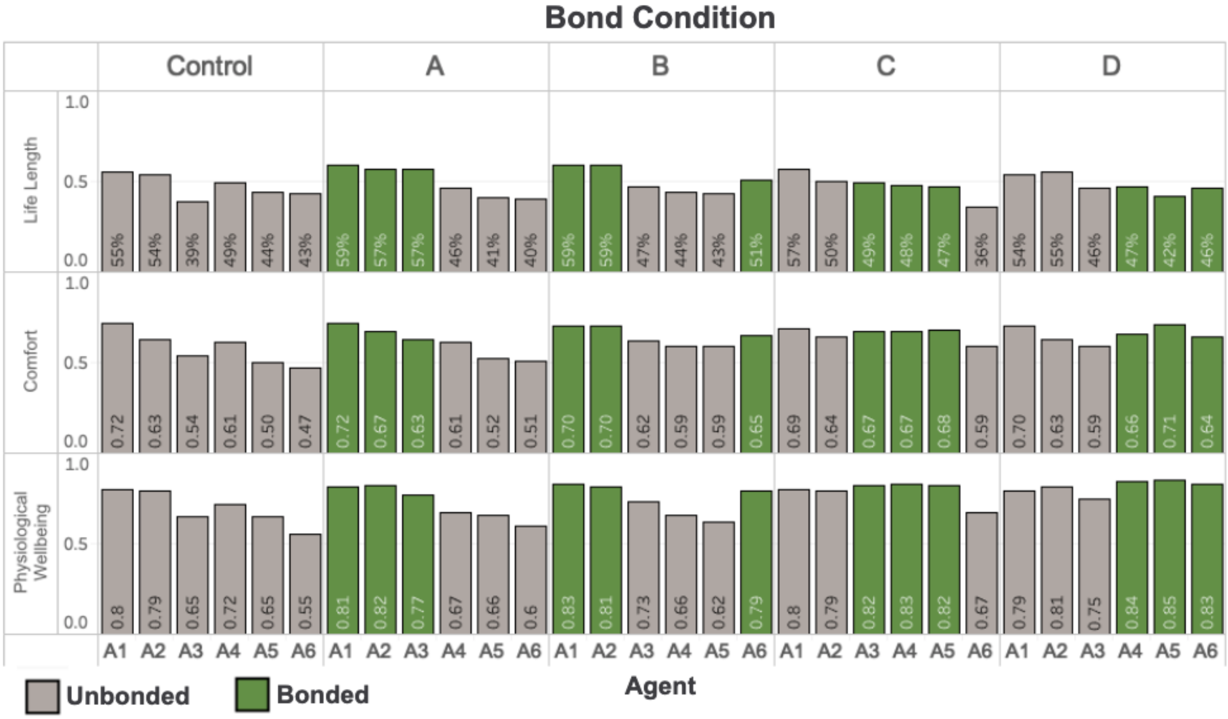

Figure 5.8 Results of the viability indicators (Life Length, Average Comfort, Physiological Wellbeing) for each agent in the experiment, broken down by Bond Condition, reported as the mean value across all experiments. Green results indicates agents who were bonded in that condition, grey indicates agents who were unbonded.

A6 nevertheless reported an increased Life Length (+8%) when bonded with higher-ranking agents. Despite the reduced *Life Length*, the Bonded group saw improved *Average Comfort* levels (0.66) and significantly improved *Physiological Wellbeing* (0.84) when compared to Control. Overall, viability benefits were still reported for the Bonded group on two of three measures.

5.4.2 Grooming Distribution

Table 5.2 shows the distribution of *Grooming* performed by each agent (the "groomer"), and the recipient of *Grooming* as a percentage of all grooming performed by each groomer.

Regardless of bond condition, the majority of grooming behaviours was directed towards the highest-ranked agents (A1 and A2) regardless of whether agents were bonded with them. Conversely, despite initiating *Grooming* with other individuals in the society, the lowest-ranking A6 was never the recipient of *Grooming* from non-bonded individuals. Instead, it was only the recipient of *Grooming* by other agents with whom it had a bond with (Bond Conditions B and D respectively).

recipient

BOND CONDITION A

		A1	**A2**	**A3**	**A4**	**A5**	**A6**
g							
r	**A1**	-	60%	13%	20%	7%	0%
o	**A2**	65%	-	19%	10%	6%	0%
o	**A3**	37%	54%	-	5%	5%	0%
m	**A4**	74%	13%	8%	-	5%	0%
e	**A5**	82%	12%	0%	6%	-	0%
r	**A6**	61%	13%	13%	8%	5%	-

BOND CONDITION B

	A1	**A2**	**A3**	**A4**	**A5**	**A6**
A1	-	45%	8%	4%	2%	41%
A2	60%	-	9%	7%	3%	21%
A3	48%	23%	-	19%	10%	0%
A4	71%	13%	9%	-	7%	0%
A5	79%	15%	0%	6%	-	0%
A6	68%	14%	8%	3%	6%	-

BOND CONDITION C

	A1	**A2**	**A3**	**A4**	**A5**	**A6**
A1	-	68%	13%	9%	10%	0%
A2	66%	-	21%	7%	7%	0%
A3	39%	17%	-	24%	20%	0%
A4	58%	16%	16%	-	10%	0%
A5	47%	12%	25%	16%	-	0%
A6	62%	19%	10%	6%	2%	-

BOND CONDITION D

	A1	**A2**	**A3**	**A4**	**A5**	**A6**
A1	-	78%	9%	9%	4%	0%
A2	80%	-	12%	5%	4%	0%
A3	60%	14%	-	14%	12%	0%
A4	54%	5%	5%	-	28%	8%
A5	64%	10%	0%	9%	-	17%
A6	46%	7%	9%	22%	16%	-

Table 5.2 Results of the amount of grooming performed by each agent (rows) on the other five agents (columns) as a percentage of the total grooming performed by the groomer. As an agent cannot perform *Groom* on itself, those corresponding cells are indicated with "−". Grey cells indicate agents who were bonded in each of the four bonded conditions.

In cases where agents were bonded, the distribution of *Grooming* behaviours shifted from being directed towards higher-ranked agents, towards affective bond partners. For instance, in Bond Condition B, Agent A3 performed 29% of its total *Grooming* lower-ranked agents (A4 and A5), and this increased to 44% when bonded with them in Bond Condition C. Similarly, A1 and A2 performed 41% and 21% of their overall *Grooming* interactions with the lowest-ranking agent (A6) when bonded with it, but zero social interactions otherwise.

Overall, we observe that, in line with expectations, *Grooming* was directed towards higher-ranking individuals, and that this distribution is offset when agents share an affective social bond. Accounting for the viability improvements for bonded agents, we suggest that increased *Grooming* between bond partners may have played a role in the viability of agents who shared a social bond.

5.5 Discussion

Overall, our results found support for the hypothesis that oxytocin's effects on specifically increasing bond partner preference for social interaction provides survival-related advantages in a small, rank-based society in physically-challenging conditions. Specifically in our agent model, we found viability of bonded agent groups to match or outperform non-bonded conditions, and that lower-ranked agents saw greater viability improvements through these social bonds.

We note one surprising result: that the third most-dominant agent in the group had the lowest rate of survival in control conditions, with all lower-ranked agents consistently outliving it. Watching experimental runs, we suggest that this was driven through an element of "competition" for resources between higher-ranking agents. We discuss this further below.

We also observed that, despite their relatively high ranks, the two most dominant individuals (A1 and A2) reported survival-related benefits in conditions where they shared an affective social bond. In terms of our model, this suggests that a positive affective relationship between socially-dominant individuals can provide further survival-related advantages in some environments.

As a result of these two observations, we hypothesised that dominant individuals, through their behaviours or affective relationships, may play a significant role in affecting the wider social dynamics and wellbeing of other individuals in a small society.

Watching the simulation runs, we observed behaviours related to altruistic and allo-parental behaviours between bonded individuals. In certain conditions, higher-ranked agents overlooked grooming with similarly-ranked agents to select a bonded (yet much lower-ranked) partner. We also observed how these altruistic tendencies do not always appear to be a viable approach when rank differences are much smaller among the least-dominant individuals, suggesting that such an approach towards lower-ranked bond partners is not always the most beneficial approach towards survival-related adaptation.

Finally, we observed some behaviours related to social amnesia—or "distrust" towards bond partners—from low-ranking individuals. Specifically, if an affective social relationship was not strengthened and maintained in the early phases of their life, the subsequent value associated with partners was affected in later stages.

5.5.1 Social Bonds Resulted in Increased Viability—Even for Dominant Individuals

In three of the four bonded conditions, we found how agent viability was improved on both an individual and group level for socially-bonded groups. This was also seen in conditions where the two most dominant agents (A1 and A2) were bonded with each other (Bond Conditions A and B): an unexpected result as dominant individuals do not experience the same benefits of social bonds as lower-ranked individuals. In Bond Condition C and D, we also saw viability benefits for Unbonded groups, suggesting some survival-related advantages for the wider social group in some conditions.

This observation led us to hypothesise that, in a small society as that used in our model, wider group viability and social dynamics may be impacted by the interactions and behaviours of smaller social groups—such as those sharing social bonds. We consider whether a tolerant approach between the two most dominant individuals in a society is actually more beneficial to their other bonded partners, and even (some members of) the wider society. For instance, the findings from Bond Condition B (when A1-A2-A6 were bonded) saw how a lack of competition (Subsection 5.5.2) between A1 and A2 significantly improved A3's viability (Figure 5.8), despite the latter agent not sharing a social bond with another. Conversely, the Control condition, when its higher-ranking agents were not bonded, saw A3 experience its poorest viability performance of all agents.

In this regard, a socio-positive relationship, akin to "tolerance", between dominant individuals may have trickle-down effects on the rest of their society. These wider effects of interactions between dominants may play a role in shaping wider-group dynamics, and may underpin the diverse social structures that emerge from the same social species (as seen in [117]). We retain this avenue of investigation for future work.

5.5.2 In Competitive Scenario, Dominance Rank is Not Indicative of Survival

As a result of higher-ranked agents having easier access to environmental resources, we had expected viability of agents to be correlated with their dominance rank in the Control conditions. However, we found how A3—the third highest-ranked agent—experienced the poorest viability in these conditions, with lower-ranking agents consistently outlived it. Watching the simulation runs, we suggest that this may be due a combination of several smaller interactions.

As the world environment contained only two resources, they would sometimes be occupied by the most dominant agents (A1 and A2) leaving A3 wandering between resources until one could be accessed. Furthermore, as agents are attracted to the highest-ranked available agents for *Grooming*, A3 would follow A1 or A2 in order to *Groom*. A1 or A2 would find food resources, which A3 would be unable to approach

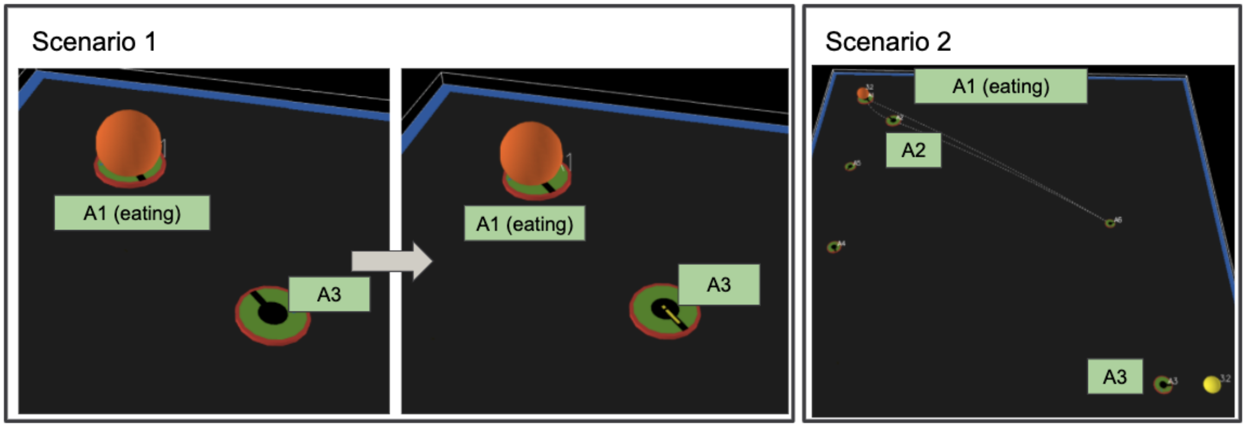

Figure 5.9 An example highlighting how conditions with affective bond partners affected the behaviours (and survival) of unbonded agents (in this case, A3). In Scenario 1 (left), the higher-ranking A1 (green disc) is eating at the food resource (orange sphere). As A3 (green disc) is unable to approach the food, it avoids the resource (Equation 5.9) and looks for another food resource to consume. In Scenario 2 (right), A1 is eating at a food resource (situated in the top left of the world), and a bonded A2 is able to approach it. This frees up the other food resource (bottom right) for A3 to approach and consume.

(due to the difference in rank) and would be forced to avoid it (Figure 5.9, Scenario 1).

As A3 dies, the social environment becomes less competitve for lower-ranked agents, increasing their likelihood to find and access food. This is further emphasised in results from Bond Condition C (A1-A2-A6 bonded), where A1 and A2 would share food resources, allowing A3 to eat at the other available food resource (Figure 5.9, Scenario 2).

The competitive scenario that emerges between the two dominant individuals when they are not bonded resulted in an additional social complexity for some agents (A3). Given the limited number of experimental runs, it is not clear as to why other lower-ranking individuals were not affected by these dynamics. Nevertheless these results may suggest that survival-related benefits associated with a higher rank are not always guaranteed, and emphasise how the dynamics between other agents can impact the behaviours and viability of the wider social group.

5.5.3 Dominant Agents Exhibit a "Parent-Child" Dynamic with Least-Dominant Individual

The highest-ranked agents, A1 and A2, would only groom with the lowest-ranked agent (A6) when bonded with it (Table 5.2), despite A6 providing no survival-related benefit for those agents. In some observations, it was noted that a dominant agent would overlook the opportunity to *Groom* with other similarly-ranked

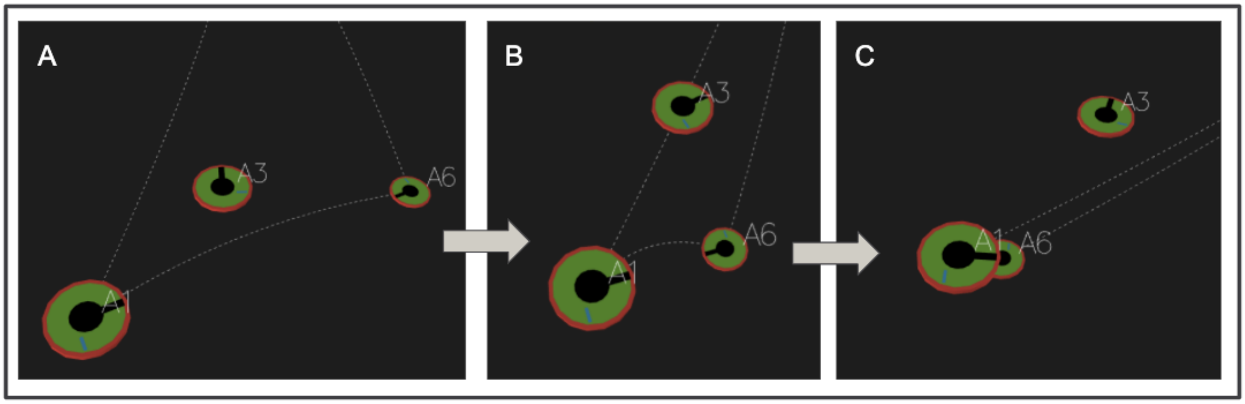

Figure 5.10 Screenshots of the "allo-parental" behaviour between the highest-ranking (A1) and lowest-ranked (A6) agents in our model. Dashed line between agents denotes the existence of an affective bond between them. A: Agent A1 is looking to *Groom*, and perceives both A3 and A6 in its vision. B: A1 selects A6 to *Groom* with, despite the presence of the higher-ranking A3. C: A1 performs *Groom* with A6, which releases oxytocin and strengthens their affective bond (Table 5.1). This then benefits the longer-term viability of agent A6 (Figure 5.8).

agents (Figure 5.10, A) in favour of its bonded, lower-ranked bond partner (Figure 5.10, B and C).

In terms of fitness benefits, such a decision appears to be counter-intuitive. A social bond with a subordinate agent appears to provide no benefit for the dominant agent—except for the potential maintenance of the bond itself—but provides benefits for the subordinate since the affective social bond is strengthened as a result. As bond strength is increased (and oxytocin is released), the lower-ranked agent improves its likelihood of finding and accessing food resources occupied by higher-ranked bond partners. The food sharing then further strengthens bond and facilitates oxytocin release.

This created a positive feedback mechanism of affective bond strength and prosocial behaviours—underpinned by oxytocin—and facilitated by the behaviours of the higher-ranked agents. Specifically, that a decision to *Groom* with a subordinate bond partner, which provides no benefit to dominant agents, facilitated greater viability of the lower-ranked agent in challenging conditions. We likened this to an altruistic, parental dynamic between these bonded agents. We also note that this type of bond also improved the viability of the higher-ranked agents (Figure 5.8), although the reasons for this are currently unclear.

5.5.4 Socio-Positive Behaviours Not Always the Optimal Decision

When the three lowest-ranked agents were bonded (Bond Condition D) the least dominant agent (A6) still showed viability benefits to being bonded with higher-ranked agents, albeit their (A4 and A5) relative rank compared to the other agents in the environment. However, A6's improved Life Length performance was

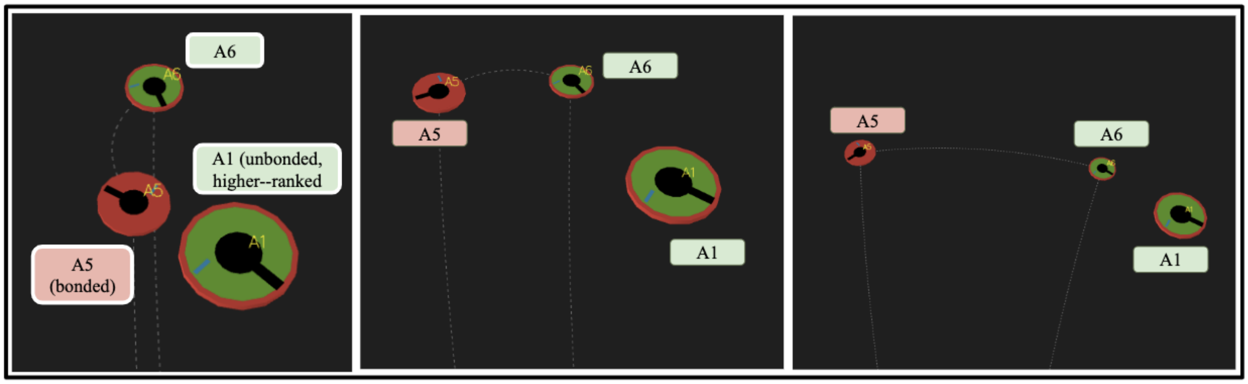

Figure 5.11 Screenshots of the "social amnesia" displayed by subordinate agents when affective bond strength is low. Agents are represented as doughnut-type shapes. A dashed line between agents indicates the existence of an affective bond. Left: Despite the fact that a bond exists between A5 and A6, internal oxytocin levels for A6 and the bond strength (DSI) between A5 and A6 is weak. Center: A6 is looking to *Groom* with another agent, and as a result of lower oxytocin levels, ignores its bond partner A5 in favour of A1. Right: A6 continues to follow A1, creating further distance between bond partners and reducing the opportunity for future *Grooming*. Colours of agents represent motivational states. Red = Hungry, Green = Lonely.

at the expense of A4 and A5's Life Length. Unlike Bond Condition B, when the dominant agents were bonded with the subordinate A6, the survival-related costs associated with "food sharing" and affective bond maintenance for lower-ranking agents is significantly higher.

Since these agents do not have the same access to food resources in our model—where higher-ranked agents have priority access—favouring a lower-ranked individual (who also has poor access to food resources) was a costly decision with little-to-no survival-related benefit (Table 5.8).

This suggests that this selfless behaviour (i.e. passive food sharing with lower-ranked individuals) may not be the optimal approach for lower-ranking individuals, particularly when the social and physical world poses a sufficient challenge. The socio-positive (or socio-mutual) behaviours exhibited was detrimental to their survival. Instead, it may be the case that more defensive, socio-negative behaviours—even towards affective bond partners—may be appropriate to ensure viability in challenging conditions. Understanding these dynamics further may facilitate understanding into the contextual differences in behavioural strategies between bond partners, and may play a role in the formation, maintanence, and loss of social relationships. We highlight this as an area for potential further work using our model.

5.5.5 Distrust and Social Amnesia: Subordinate Agents May Need Earlier Social Reinforcement

In some simulation runs where the most subordinate agent (A6) did not *Groom* with bond partners during the early stages, we observed behaviours akin to distrust or "social amnesia" of bonded agents. If grooming did not occur in the earlier phases—perhaps due to a lack of proximity between agents—oxytocin levels and bond strength (DSI) would fall to lower values. As a result, the valence of bond partners would be significantly reduced: with partner selection favouring higher-ranked, non-bonded agents, despite the presence of an affective bond partner (Figure 5.11).

In absence of early bonding to reinforce affective bond quality, these bonds appear to be "forgotten" by the subordinate agent. As a result, the survival-related benefits associated with social bonds is inhibited, with reduced food access, resulting in reduced viability for those agents. This mirrors the effects discussed in Subsection 5.5.3 .

This is in line with suggestions by [192] that oxytocin in mammals may potentially promote the affective processing of social stimuli in the early phases of social interaction. If such affective processing does not occur (as a result of lower levels of oxytocin, for instance), then higher levels of oxytocin in later phases may not be sufficient to introduce the affective quality in a social relationship. As a result, as we find in our model, the processing of social cues in later phases (of life) is hindered when such individuals do not experience social reinforcement in early phases.

Such flexible decision-making may highlight the importance of maintaining a strong bond with lower-ranked individuals, particularly if more dominant individuals are present in the environment. In these conditions, when bonds are not reinforced for subordinate agents, they may attempt to form new relationships with more dominant members of society. Though the current model did not include the formation of new bonds, implementing such a mechanism in future work may provide insight into the emergence of different social structures.

5.6 Summary

In this chapter, we tested one hypothesised effect of oxytocin's context-dependent role on the processing of social cues—-that it specifically promotes bond partner preference—and investigated its role on the viability and social interactions in a small society of agents. We tested these effects in a small, rank-based society using a number of pre-determined affective bond combinations in a physically-challenging environment.

We developed a new quantitative model of affective bond strength between dyads, called the *Dyadic Strength Index* (DSI). We also introduced an additional component to the model, called the *Social Assessment Component* which accounts for the perceived identity of others—through both social rank and affective relationship with others—in order to contextualise the execution of a social behaviour.

Our results found that oxytocin-mediated preference for affective bond partner interaction resulted in survival-related benefits for individuals with (and in some cases, without) affective bond partners. We also found that viability was not determined simply by a dominance rank or the presence of affective social relationships. Rather, the wider social context—in particular, the interactions of and between dominant individuals—may play a role in the emergent social dynamics and viability of the wider group.

Our observations raised two hypotheses for further investigation. Firstly, that wider social dynamics, structure and viability may be influenced by the interactions of dominant individuals or smaller groups within the society. Secondly, that socio-positive, or socio-mutual, interactions were not always the optimal strategy to determine viability and that, in some cases, a defensive or socio-negative approach may be necessary.

The complex dynamics and interactions resulting from a social hierarchy—such as that used in our model—can constitute a source of stress for some individuals. Due to the different ways in which social individuals experience and cope with stress (and as we have seen in this chapter), behaviours may not be universally neutral or positive. Negative interactions, such as aggression towards others, may be a possibility. These small-scale interactions—particularly from influential members of a social group—may also play a role in shaping social structure, and affect the survival and wellbeing of individuals in that society. For individuals in that society, the presence of, and interactions with, social bond partners, has been suggested to represent a source of comfort and safety in a stressful social and physical environments. This “safety cushion” associated with affective bond partners is a phenomenon known as “social buffering”, which describes how social support can provide an adaptive response to the potentially-deleterious effects of stress—and even future responses to stressful situations—in several ways. In turn, this adaptive response may provide advantages to an individual’s health and longevity over its lifetime.

In the following chapters, we begin our investigation into the role of social support on the regulation of “stress”. This regulation of the stress response through the social environment is a key concept that underpins social allostasis principles, and we investigate how this hypothesised stress-regulation of social support affects the long-term wellbeing of the social group.

Chapter Six

Investigating the Effects of Social Stress and Individual "Personality" Traits on Social Dynamics

6.1 Introduction

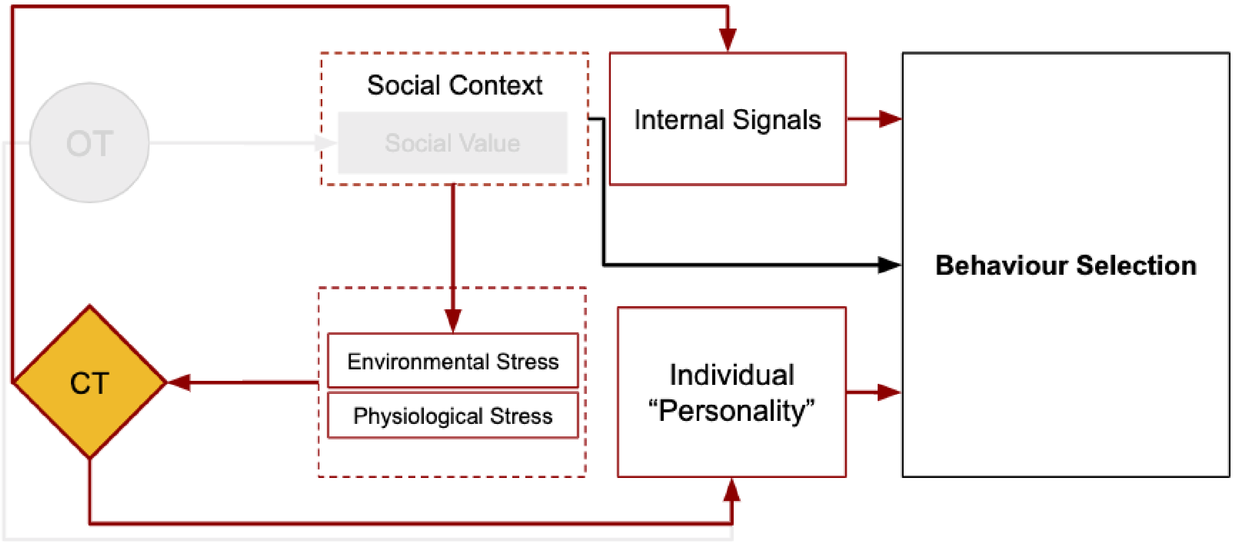

Figure 6.1 A high-level view of the effects investigated in this chapter. This investigation introduces a stress hormone of cortisol (CT) and a personality trait related to stress tolerance. We temporarily remove the effects of oxytocin from the model. These effects are reintroduced in Chapter 7. A more granular view of this model can be seen in Figure 6.2.

In Chapters 3 to 5, we investigated several hypotheses related to oyxtocin's effects on the attention and processing to social cues, in order to understand some potential hormonal mechanisms that may underpin social interactions and social relationships to promote long-term viability.

Despite early findings (Chapter 3) showing adaptive potential for the "prosocial" approach of oxytocin mechanisms—that oxytocin promotes prosocial behaviours by promoting attention given to social cues—subsequent chapters (Chapters 4 and 5) found that, with respect to viability, oxytocin's effects on social cues should vary across different contexts. This was further emphasised as we accounted for social context into our investigations in Chapter 5 through social hierarchy and affective bonds, investigating oxytocin's proposed effects on affective bond partner preference. Our findings from that chapter showed how, in a small social hierarchy, oxytocin-mediated bond partner preference for social interactions provides a mechanism of behavioural adaptation in a challenging environmental condition, promoting the viability of individuals with affective social bonds.

One significant limitation of those findings is that the investigation only considered the prosocial interactions associated with a social structure. For individuals in a society, a rank-based social structure presents an additional challenge that must be adapted, to in order to ensure survival for individuals in the society. Social bond partners can provide opportunity for prosocial interaction, who may also play a role in ameliorating the stress associated with difficult situations (called the "social buffering" hypothesis, which we investigate in the next chapter). In absence of these social bond partners, dominance societies can be a source of socio-negative interactions and stress, which may threaten the long-term survival for certain individuals in a society. It is not axiomatic to suggest that stress responses are experienced more in subordinates than dominant individuals. Whether higher-ranked or lower-ranked individuals in a (strictly rank-based) society experience the most stress is highly dependent on several factors, one of which is individual "personality" differences. One type of personality difference is an individual's tolerances to stress levels, which has been found to be associated with an individual's social rank.

However, whether an individual's responses to stress is a function of their social rank, and the social structure, or vice versa, is still clear. Nevertheless, these differences in personality may then determine how individuals respond to and cope with stress, and therefore affect their (pro and anti-)social interactions. In turn, these interactions may then play a role in determining the wider group dynamics and social structure.

Taking our observations from the previous chapter, we account for both the "stress" and socio-negative behaviours associated with living in a rank-based social structure, and investigate the role of different tolerances to stress (an abstraction of a "personality" trait) on the social interactions, group dynamics, and overall wellbeing of our small society.

Using our findings from Chapter 5 along with observations from ethology, we hypothesised that these different tolerances to stress will have a significant effect on the social interactions and dynamics in a small society. Specifically, that social dynamics will be determined by the behaviours and social interactions of dominant individuals on the rest of society. Furthermore, we hypothesised that, as a result of these social dynamics, overall group wellbeing would be affected by these different stress tolerances.

We test these hypotheses in a small, rank-based society using several different conditions of individual stress tolerance—both related and unrelated to dominance rank—across several environmental conditions related to different degrees of food availability. We extend on our previous model by including a "stress system"—an abstraction of several physiological and behavioural effects of cortisol—as well as an associated "tolerance" of this hormone in each of the six agents.

Our results find support for the hypothesis that these rank-related tolerances to stress play a role in determining social dynamics. Contrary to our expectations, we find that in some conditions, social dynamics were driven by the interactions of subordinate individuals, and not dominant individuals as we had expected. In addition, we report survival-related benefits to the group when accounting for the effects of a "stress system". However, we found no support for our hypothesis that different tolerances to stress will affect overall group wellbeing.

6.1.1 Social Dominance Rank and Stress

Social hierarchies and dominance rank are natural by-products of life in a society, but come with costs as well as advantages. Dominance rank in these societies is oftentimes (but not always [193]) associated with the most survival-related benefits. Taking a Darwinian perspective, a higher dominance rank in a society is associated with increased resource access (i.e. food) as well as greater or easier access to grooming partners, resulting in overall health and longevity [108], [171], [172].

As a result of these fitness-related benefits, it was thought that subordinates always experience the largest amount of "stress" (using elevated cortisol levels as one measure), since they would be subject to the largest amount of physical and psychological stressors [117], [194]. However, this relationship between subordinance and elevated stress levels is not always obvious. Correlations between dominance rank and cortisol levels have been found to be affected by social structure. For instance, cortisol levels have been found to be higher in subordinates in stable societies (where rank is inherited or not challenged), but higher in dominants in unstable societies (where rank is challenged or must be maintainted) [117]. Even in stable societies, it has been suggested [195] that lower cortisol levels are not direct markers of social dominance, but markers of individuals with the necessary personality traits to take advantage of the social environment.

In addition, gender [30], environmental and social challenges, [108], [117], [196], [197], the presence of close-kin or social coping mechanisms [108], [115], [117], and individual personality differences [116] all play a role in determining how individuals cope with, and adapt to, "stress".

With this in mind, one aspect of individual differences may be the physiological and behavioural responses to stressful physical and social environments [116]. This individual stress response has been seen associated with individual dominance rank [198], and can be seen as a "personality trait" [195], [199] that may play a role in individual stress coping mechanisms (such as displacement via aggression). These individual responses may then play a role in determining the wider social dynamics and structure.

In summary, there is no definitive relationship between dominance rank and levels of stress across species. Whether it is dominance rank that determines an individual's response to stress (and/or relative cortisol levels) or vice-versa is still unclear. What is clear are the individual differences related to rank [115], [195] and the variety of social dynamics for different populations of the same species [117]. In turn, these different personality types may play an important role in shaping overall social structure, which may explain some of the variance seen in real-world societies [108], [117].

6.1.2 Aims & Hypothesis

In this chapter, we aim investigate how one rank-related "personality" trait—an individual's stress tolerance—can affect the social dynamics (defined as the frequency and direction of socio-positive and socio-negative behaviours) and viability-related performance of individuals in a rank-based society.

In order to conduct this investigation, we include a stress-related hormone (cortisol) with several physiological and behavioural effects. We tested these effects using six different conditions of an individual's tolerance to stress, which we considered to be a loose abstraction of an individual's personality, across four different environmental conditions related to food availability.

Based on our findings from our previous chapter along with some observations in biology, we hypothesised that social dynamics and group viability will be significantly affected by rank-related stress tolerances. Furthermore, we hypothesise that these social dynamics will largely be determined by the interactions of dominant individuals in the society, via a "trickle-down" effect of stress and social behaviours. Specifically, aggressive dominant agent(s) will result in an aggressive society, and non-aggressive dominant agent(s) will result in a non-aggressive society.

6.2 Extensions to the Agent Model

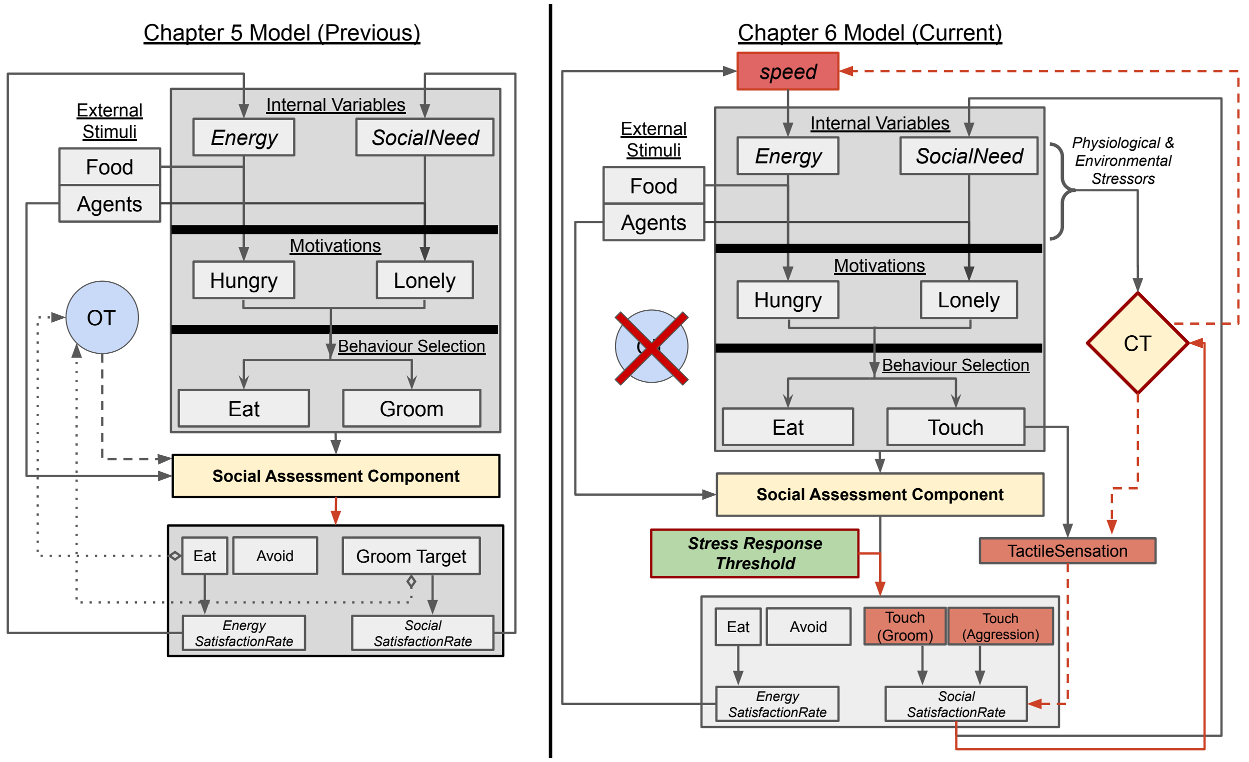

Figure 6.2 Illustration of the agent model used in this investigation (right) compared to the previous chapter (left). Extensions to the previous model are denoted in red. Oxytocin's effects are not included in the current model. A hormone of cortisol (CT) is introduced in this model, released as a function of physiological and environmental stressors (Equation 7.2). CT modulates the intensity of performing tactile behaviours (*TactileSensation*, Equation 6.7) and speed of movement (Equation 6.4). *Stress Response Threshold* is a value to determine how much CT agents can tolerate before they become stressed (Section 6.4). Tactile behaviour is performed through either *Groom* or *Aggression*, depending upon the stressed state. Dashed line indicates the modulation of values. Tactile behaviours (*Groom* and *Aggression*) also have effects on the recipient agent, as seen in Figure 6.3.

Figure 6.3 shows the agent model used in this investigation. We highlight a number of changes from the previous model.

Firstly, we introduce hormonal mechanisms related to cortisol, which adapts agent behaviour in several ways. We discuss this further in Subsection 6.2.1.

Secondly, we introduce a socio-negative behaviour of *Aggression* to the agent model which, along with *Groom* from previous iterations, constitutes a sub-behaviour of *Touch* (previously *Groom*, which was strictly socio-positive.)

Finally, we endow each agent with an internal threshold value, called the *StressResponseThreshold* (Sub-

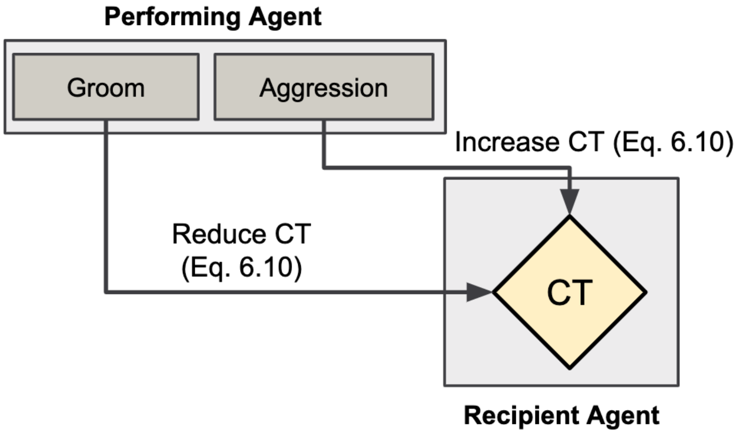

Figure 6.3 The effects of the two different tactile behaviours on the cortisol levels of recipient agents. *Grooming* reduces cortisol in the recipient agent, while *Aggression* increases cortisol. The rate at which cortisol is increased or decreased is a function of the intensity of the interaction, as seen in Equation 6.7.

section 6.2.2). This value describes how much cortisol an agent can withstand in its physiology before it undergoes a state of "stress". This stressed state then determines the type of social behaviour (socio-positive or socio-negative) being performed. We use this threshold as one feature of an individual's "personality", i.e. the tolerance to stress. This may also be considered an abstraction of an "activation level" of the sympathetic nervous system [200], [201]. We describe each of these changes in more detail below.

To focus our investigation on the effects of these mechanisms, we have removed the effects of oxytocin and social bonds used in previous chapters (Chapters 3 to 5) from this current investigation.

6.2.1 The Role and Release of Cortisol

We introduce a stress-related hormone of cortisol to the agent model. As we discussed in Section 1.2, the biological hormone of cortisol is a steroid hormone seen to be released in (anticipation of) stressful or challenging situations. Prolonged elevation of cortisol has maladaptive consequences [136], but is an important hormone of physiological and behavioural adaptation [29], [133]–[135]. In our model, cortisol is included as a hormone with these adaptive and stress-related effects.

Motivated by biological literature where cortisol is associated with a lack of (perceived) control [17], [108], [117], the cortisol release ($CTReleaseRate$) in an agent's physiology is related to their perceived control of their physiological (internal) and physical/social (external) environments. In other words, the less control an agent perceives that it has, the larger the rate of cortisol release.

Specifically in our model, this rate is calculated by two components:

- An agent's perceived "physiological" control, calculated using the mean error of its internal variables (i.e. "What is the intensity of my internal deficits?")
- An agent's perceived "environmental" control, defined by the sum of *available* resources (food or other agents, i.e. "What resources can I currently access to satisfy these deficits?")

.

More specifically, for each agent A:

$$CTReleaseRate_A = \left(\frac{\sum\limits_{i \in (P)} d_i}{\kappa} - \frac{\sum\limits_{i \in (S)} \hat{S}_i}{2}\right) \times c \tag{6.1}$$

where $\sum d_i$ is the physiological stressor, calculated by the sum of errors of both physiological variables (P) *Energy* and *SocialNeed*, $\sum \hat{S}_i$ is the sum of perceived "available" resources (S), κ is the sum of the upper limits of both internal variables, and c is a fixed constant to regulate the value of *CTReleaseRate*.

"Perceived" availability of resources is defined as resources ($Food$ or $Agents$) that an agent perceives it can access with no challenge. It is calculated using the Social Assessment Component (as discussed in Section 5.2.5, and Equation 5.7). As agents have no social bonds in this investigation, the value of *AgentVal* from Social Assessment Component is equivalent to the difference in rank between agents.

For food resources, if *AgentVal* of an agent at the food resource is negative (if $AgentVal < 0$), that food resource cannot be accessed (as this equates to a higher-ranking agent at the resource). If a food resource is unoccupied or a lower-ranked agent is perceived to be at that resource (if $AgentVal > 0$) then that food resource is seen as available.

The perceived availability of another agent as a grooming option is calculated using the inverse of the *AgentVal* value. The larger the $rank_{diff}$ between agent A_P and a potential grooming partner A_R, the larger the cortisol release in agent A_P. In this way, seeing higher-ranking agents will increase the *CTReleaseRate*, but seeing subordinates reduces it. Therefore:

$$\sum d_i = d_{energy} + d_{socialNeed} \tag{6.2a}$$

$$\sum \hat{S}_i = \hat{S}_{\text{agents}} + \hat{S}_{\text{food}} \tag{6.2b}$$

where $\hat{S}_{\text{agents}} = C_{\text{agents}} \times (1 - AgentVal)$, and

$$\hat{S}_{\text{food}} = \begin{cases} C_{\text{food}} & \text{if } AgentVal \leq 0 \\ \text{else } 0 \end{cases} \tag{6.3}$$

In addition to this release rate, cortisol levels can also be decreased by performing social behaviours (*Grooming* and *Aggression*) on other agents, and further increased by receiving *Aggression* (discussed in detail below).

Cortisol then adapts two aspects of agent behaviour. The first, an abstraction of cortisol's effects related to increased movement and energy metabolism [117], is that cortisol modulates the speed at which an agent moves through the environment. This also increases the energetic cost associated with movement. The second effect of cortisol is that it increases the intensity with which tactile behaviours are performed. This intensity has different physiological effects on both the actor and recipient of a particular social interaction.

First Effect of Cortisol: Adapting Movement Speed & Energetic Cost

Cortisol adapts the speed at which an agent moves through the environment—an abstraction of biological cortisol effects related to energy metabolism and increased movement [117]. The more cortisol present within an agent's physiology, the faster it moves. Moving at a faster rate uses up more energy. Therefore, the rate at which *Energy* depletes becomes a function of the speed of which it moves at (Equation 6.6):

$$speed_t = speed_0 \times (1 + (CT \times c)) \tag{6.4}$$

Where $speed_0 = 0.5$ (the default speed at the start of the experiment), t is the current time step, and c is a constant value used to regulate the value of *speed*. The rate of energy depletion, called the *EnergeticCost*, is then increased as a function of the speed of the agent.

$$EnergeticCost_t = EnergyLossRate \times (2 \times speed_t) \tag{6.5}$$

At the current time step t, the value of the internal variable *Energy* is updated by this updated *EnergeticCost*.

$$Energy_t = Energy_{t-1} - EnergeticCost \tag{6.6}$$

where *EnergyLossRate* is 0.003, the default rate of energy loss.

Second Effect of Cortisol: Adaptating the Intensity of Social Behaviours

Similar to the modulatory effects of OT in Chapter 4, cortisol adapts the intensity of the tactile behaviours (*Groom* and *Aggression*), affecting the physiology of both agents involved in tactile contact as follows:

$$\mathit{TactileIntensity} = b_{touch} \times CT \times c \tag{6.7}$$

where b_{touch} is the intensity of the winning behaviour *Touch*, and c is a constant use to regulate the *TactileIntensity* value. This *TactileIntensity* value is then used to correct the *SocialNeed* after either of the tactile behaviours (*Groom* or *Aggression*) have been performed:

$$\mathit{SocialNeed}_t = \mathit{SocialNeed}_{t-1} + (\mathit{TactileIntensity} \times c) \tag{6.8}$$

The *Groom* and *Aggression* behaviours have opposing effects on the cortisol levels of the recipient agent (A_R). When *Groom* is performed by agent A_P, cortisol of the recipient agent is reduced at a rate proportional to the *TactileIntensity* received. However, when an agent receives *Aggression*, its cortisol is increased proportionally to *TactileIntensity*, while cortisol of the performing agent A_P is reduced. Specifically:

$$CT_{A_P} = CT_{A_P} - TactileIntensity_{A_P} \tag{6.9}$$

$$CT_{A_R} = \begin{cases} CT - (TactileIntensity_{A_P} \times c) & \text{if } \mathit{Groom} \\ CT + (TactileIntensity_{A_P} \times c) & \text{if } \mathit{Aggression} \end{cases} \tag{6.10}$$

where A_R is the recipient agent and c is a constant value used to regulate values.

Taking inspiration from real-world dynamics, this *TactileIntensity* results in recipient agents being “pushed” a short distance relative to the intensity of the tactile behaviour. Specifically, for *Aggression*, large *TactileIntensity* values has the potential to significantly displace recipients of the socio-negative behaviour.

6.2.2 Stress Tolerance & Behaviour Adaptation

In our model, the *StressResponseThreshold* (θ_{stress}) represents a type of “personality trait” in an agent loosely associated with their tolerance to stress. Specifically, this value, which can be between 0–1, determines the amount of cortisol an agent can “withstand” in its system before it undergoes a state of “stress” (an abstraction of the sympathetic nervous system activation [200], [201]). When cortisol exceeds this value, an agent is deemed to be “stressed”, else it is not stressed.

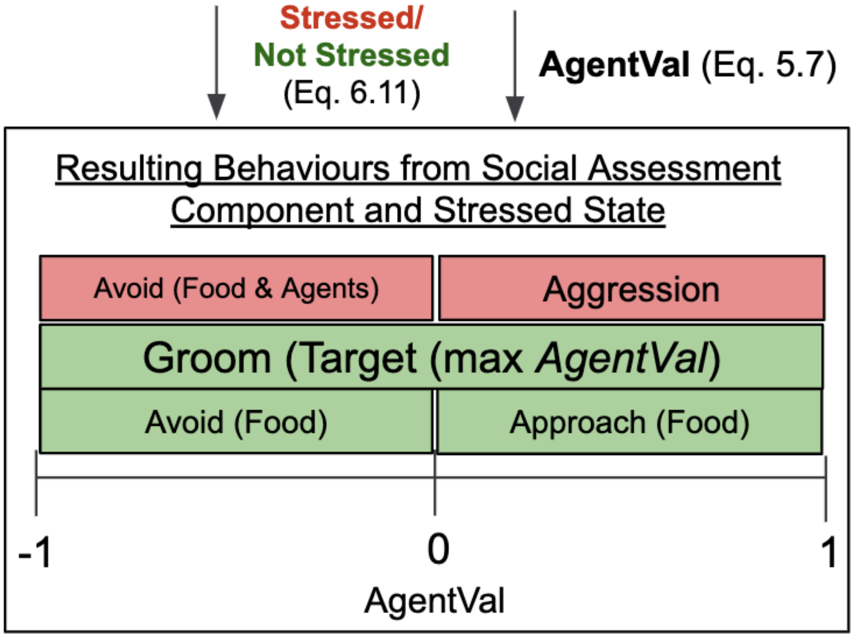

Figure 6.4 Diagram showing how approach/avoid behaviours and Groom/Aggression behaviours are adapted when agents are stressed (Equation 6.11). Green bars indicate behaviours when agents are not stressed, and red behaviours indicate the behaviours when agents are stressed. The Social Assessment Component returns the value of *AgentVal* (Equation 5.7). As no affective bonds are present in this investigation, *AgentVal* is equivalent to the normalised difference in agent rank.

$$stressed\text{-}state = \begin{cases} not\text{-}stressed & \text{if } CT < \theta_{\text{stress}} \\ stressed & \text{if } CT \geq \theta_{\text{stress}} \end{cases} \tag{6.11}$$

This *stressed* state adapts agent behaviour in numerous ways. Firstly, it adapts the type of tactile behaviour (*Touch*) that an agent looks to perform (*Groom* or *Aggression*) in order to satisfy its *SocialNeed*:

$$Touch = \begin{cases} \text{Groom} & \text{if } not\text{-}stressed \\ \text{Aggression} & \text{if } stressed \end{cases}$$

Secondly, when *Aggression* is not possible (due to subordinance), the angle at which a subordinate turns from a dominant is modulated by the level of cortisol. In other words, when cortisol levels are high and subordinate agents are stressed, these agents turn away and avoid interactions with dominants. We can consider this an abstraction of the fight-or-flight response [21]:

World Condition	**Food Availability Description**	**Number of Food Resources**	**Total Nutritional Availability**
A	Very Poor	One	4
B	Poor	Two	8
C	Good	Three	12
D	Abundant	Four	16

Table 6.1 Table describing the food availability in each of the world conditions. Each food resource has an initial nutritional value of 4 units. Screenshots of each of these conditions can be seen in Figure 6.5.

$$\angle_{\text{avoid}} = \begin{cases} 90 \times CT & \text{if } \textit{not-stressed} \\ 90 \times (1.5 \times CT) & \text{if } \textit{stressed} \end{cases}$$

When cortisol levels are below the threshold value θ_{stress}, the behaviour *Touch* is performed through a *Groom* behaviour. When an agent's cortisol value exceeds θ_{stress}, and agents are stressed, tactile behaviours are performed as *Aggression* towards subordinate agents (i.e. when $AgentVal > 0$) or, where perceived agents a higher-ranked ($AgentVal < 0$), agents will *Avoid* dominants at an angle modulated by their cortisol levels. This behaviour selection can be seen in Figure 6.4.

6.3 Experimental Set Up

Experiments were conducted using the same simulated platform used in previous chapters and the same society of six agents. In line with the setup of the rank-based society in Chapter 5, each agent held a different dominance rank, implemented as a value between 1 (highest) to 6 (lowest).

We tested six different conditions related to agents' *StressResponseThreshold* value: one *Control* condition with no *StressResponseThreshold*, three conditions where all agents had the same *StressResponseThreshold* value (called *Low*, *Neutral*, and *High*), and two conditions where agents had different *StressResponseThresholds*: both directly *Low→High condition* and inversely *High→Low condition* correlated with agent rank. These different *StressResponseThresholds* loosely relate to individual "personality" differences, where a higher-tolerance loosely corresponds to a less-anxious, more relaxed agent, and a lower-tolerance corresponds to a more anxious agent. The specific values for each of these conditions can be found in 6.2.

We conducted our experiments across a total of four different environmental conditions. Each of these world conditions contained a different amount of food resources and related to varying challenges of the

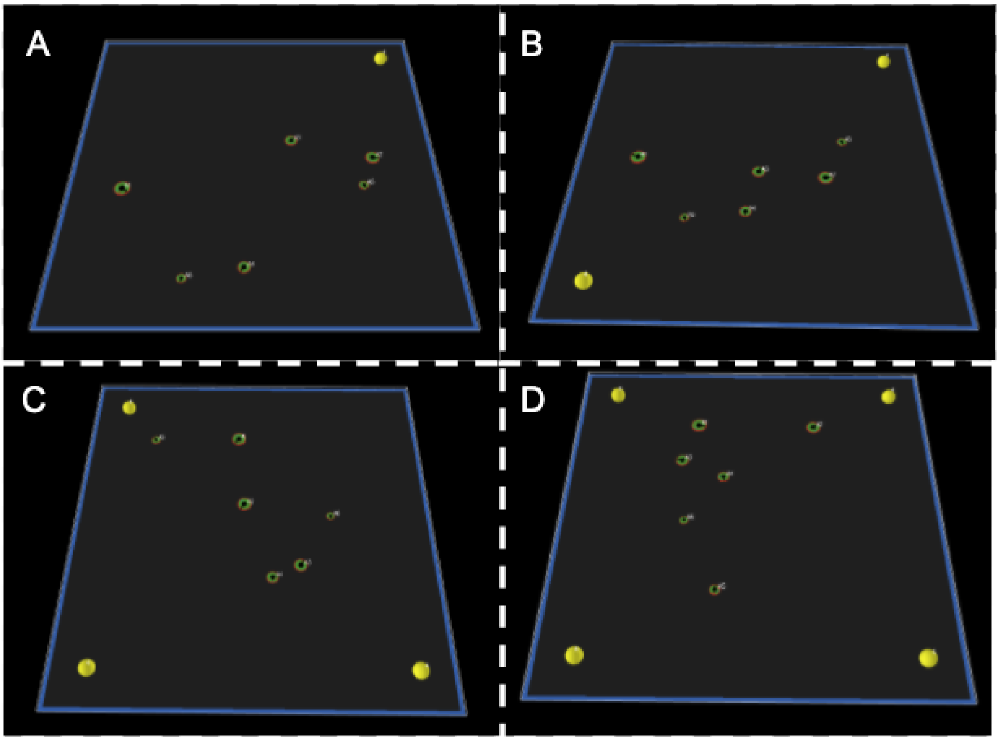

Figure 6.5 The four different world conditions used in this investigation of varying food availability. Food resources are represented as yellow discs. Agents are represented as doughnut-type shapes. Further description can be found in Table 6.1.

	Agent(Rank)					
Stress Threshold Condition	**A1**	**A2**	**A3**	**A4**	**A5**	**A6**
Control	-					
Low (L)	0.2					
Neutral (N)	0.5					
High (H)	0.7					
Low→High (LH)	0.2	0.3	0.4	0.5	0.6	0.7
High→Low (HL)	0.7	0.6	0.5	0.4	0.3	0.2

Table 6.2 Table showing the different *StressResponseThreshold* conditions, along with the θ_{stress} values of each agent per condition as described in Section 6.2.2. When an agent's cortisol exceeds this value, they become stressed and adapt their behaviours in a number of ways (see Figure 6.4).

physical environment (Very Poor, Poor, Good, and Abundant). Fig 6.5 shows screenshots of these four world conditions, and these conditions are further described in Table 6.1.

We measured our results using the three viability indicators (Life Length, Average Comfort, Physiological Wellbeing, Section 2.4.1). We also reported a number of behaviour-based metrics. *RateOfAggression* describes how "aggressive" an individual or society is by calculating the total number of *Aggression* behaviours performed as a percentage of the total social behaviours (*Grooming* and *Aggression*) by that individual agent or group:

$$\mathit{RateOfAggression} = \frac{\sum \mathit{Aggression}}{\mathit{Groom} + \mathit{Aggression}} \tag{6.12}$$

To further understand the relative stress of the society and taking inspiration from approaches in real-world societies [117], we also calculated the cortisol of each agent A_n as a percentage of the cortisol level of the most dominant agent A1:

$$\mathit{relative}CT_{A_n} = \frac{\mu(CT_{A_n})}{\mu(CT_{A_1})} - 1 \tag{6.13}$$

where $\mu(CT_{A_n})$ is the mean CT level of agent n, and $\mu(CT_{A_1})$ is the mean *Cortisol* level of the dominant agent A1.

To correct for different life lengths of agents, we report encounters of *Aggression* on *Aggression/1000* time steps, which we call the *AggressionFrequency*:

$$\mathit{AggressionFrequency}_{A_n} = \frac{\sum \mathit{Aggression}}{1000} \tag{6.14}$$

Each simulation run was run for a total of 15,000 time steps, capturing data at each time step. The 4 world conditions and 6 *StressResponseThreshold* conditions resulted in 24 experimental combinations. Using experience from prior investigations, each condition was ran for a total of 20 simulation runs each. This results in 480 simulation runs and approximates to 70 hours of total simulation run time.

Results are aggregated across all simulation runs both at an individual and group level. We use Pearson's Correlation Coefficient to describe correlations between experimental variables. For between-group comparisons, we use one-way analysis of variance. Statistical significance is reported at $p \leq 0.05$ and denoted with an $*$. Where appropriate, we also report significance at the 0.01 level, denoting it with $**$. We also supplement quantitative data with observations of the simulation runs in real-time and discuss our observations where appropriate.

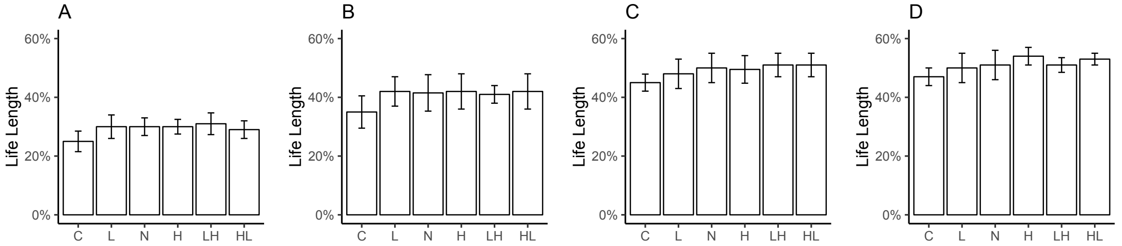

Figure 6.6 Mean Life Length, aggregated by group. Columns indicate *StressResponseThreshold* condition. C: Control group (no cortisol mechanisms). L: Low. N: Neutral. H: High. LH: Low→High. HL: High→Low. Error bars indicate standard error of the mean (SEM). Mean Life Lengths (alongside SEM) for each condition: A = 29% (± 3%), B = 41% (± 6%), C = 49% (±4%), D = 51% (±4%)

6.4 Results

6.4.1 Viability Indicator Results

Accounting for cortisol mechanisms within the agent's physiology resulted in increased *Life Length* of the overall group when compared to Control (Figure 6.6), regardless of world or *StressResponseThreshold* conditions. We saw improved *Life Length* of the overall group with increases in food resources (World Conditions A-D, r = 0.83), but found no correlation between mean group *Life Length* and the different *StressResponseThreshold* conditions (r = 0.17).

We observed little differences in the *Average Comfort* levels of the society at a group level (Figure 6.7, x-axis). We find no statistical difference between *Average Comfort* levels in groups with and without cortisol mechanisms ($r = 0.052$). At a group level, *Average Comfort* levels were not affected by the *StressResponseThreshold* or world conditions related to food availability.

At an individual agent level, we find a statistically-significant correlation between an agent's rank and *Average Comfort* levels (Figure 6.8, bottom row), $r = 0.70, p < 0.01$) across all conditions, including the control group. Therefore, *Average Comfort* levels were determined by an agent's rank, but not by the different *StressToleranceThreshold* or world conditions.

Compared to control, agents showed small improvements in their *Physiological Wellbeing* when endowed with cortisol mechanisms at a group level (Fig. 6.7, y-axis). At an aggregated group level, we note little differences between the average *Physiological Wellbeing* and *StressResponseThreshold* values in each world condition. At an individual level across all conditions, we find a statistically-significant correlation between individual rank and an agent's *PhysiologicalWellbeing* ($r = 0.79$, $p < 0.01$): the higher their rank, the higher their overall *PhysiologicalWellbeing*. This is true in both *Control* and non-control conditions.

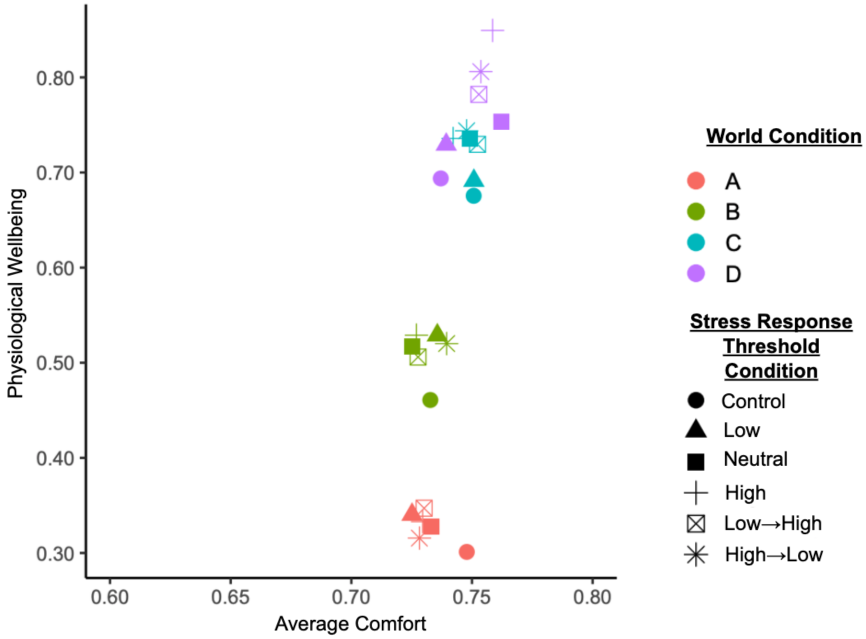

Figure 6.7 Mean Comfort and Physiological Wellbeing values for the overall group, aggregated by World Condition (colour) and Stress Response Threshold condition (shape). Note that axis values have been adjusted for visualisation purposes.

In summary, when cortisol mechanisms are present, we see improvements to viability at both a group and individual level on two of three measures—*LifeLength* and *PhysiologicalWellbeing*—when compared to control. Differences to these metrics were not correlated with different *StressResponseThreshold* conditions. Though we see no difference for group level *AverageComfort* between groups, we note correlations between rank and *AverageComfort* (Figure 6.8), across all world conditions.

6.4.2 Stress (Cortisol) Levels

We note an inverse correlation between increased food resources (World Conditions A-D) and mean group cortisol (CT) levels (Fig 6.9, left). We note a moderately-strong correlation between these two variables (Table G.1), observed across all *StressResponseThreshold* conditions. We therefore find an association between food availability and decreased stress (CT) levels for the group.

Focusing on conditions where agents had different *StressResponseThresholds (Low→High and High→Low)*, we observed negligible changes in mean CT levels for the group (Fig 6.9, right), highlighted by poor results from the statistical analysis (Table G.1, bottom). Therefore, we find that the tolerance of the stress hormone (*StressResponseThreshold*) did not affect the mean levels of cortisol.

Due to the "control" associated with higher social rank, we expected higher-ranked agents (A1 and A2) to

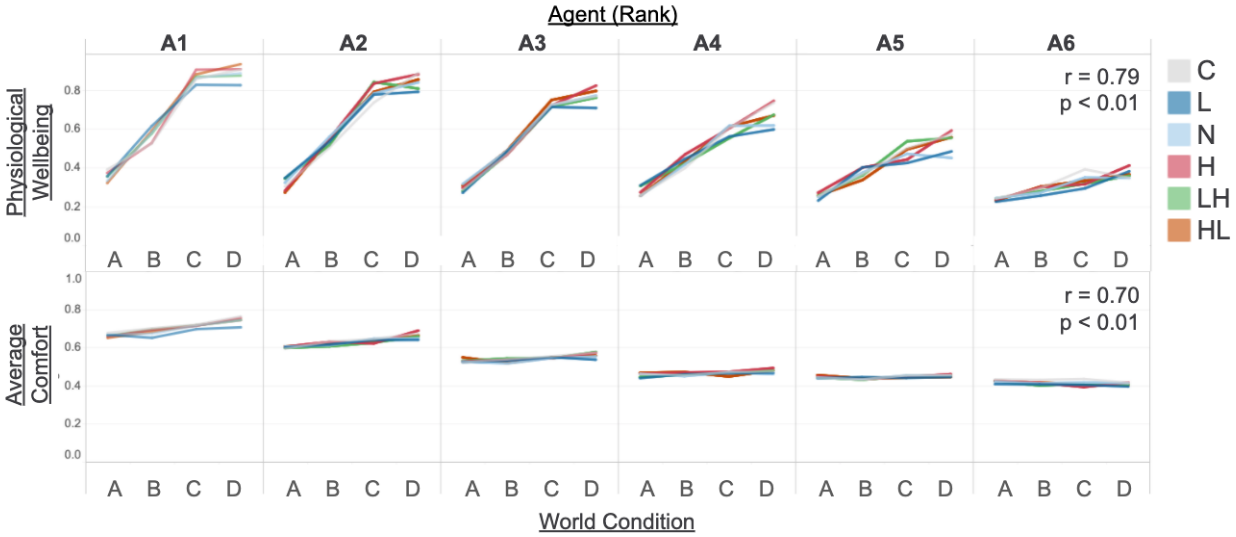

Figure 6.8 Mean Physiological Wellbeing levels (top) and mean Average Comfort levels for each agent (A1 to A6) across all world conditions (A, B, C and D), sorted by agent rank. r-values denote correlation between agent rank and the corresponding metric. p-value denotes results of significance testing. Line colours denote the *StressResponseThreshold* condition. C = Control. L = Low. N = Neutral. H = High. LH = Low→High. HL = High→Low.

report the lowest levels of cortisol (Equation 7.2). We find support for this across all *StressResponseThresholds* conditions, with an inverse trend observed between an agent's rank and their mean CT levels (Table 6.3). This trend becomes more statistically significant as food availability increases (from world conditions A to D, Table F.1).

Despite improved world conditions, mean CT levels remained elevated for subordinate individuals (A5 and A6), across all of the different *StressResponseThreshold* conditions (Table 6.3). We noted a large variance of mean CT levels for dominant agents (A1: ±0.29, A2: ±0.21) between world conditions, but a smaller variance for the subordinates (A5: ±0.13 and A6: ±0.10). Taking the view that CT levels can be seen as a function of perceived control, the smaller variance for subordinates may indicate that their perception of control in the environment was not improved despite the increased food availability.

Mirroring the approach used by [117], we also calculated the relative cortisol levels (*RelativeCT*) of individuals as a percentage of A1's mean cortisol levels (Figure 6.10). We observed that relative cortisol levels remain fairly similar between *StressResponseThreshold* conditions, with increased variance as food resources increased (World Conditions A to D).

These results showed that rank-related differences in cortisol levels were magnified as food availability increased, but not as a result of increasing stress tolerance (from *Low, Neutral, High*). These differences are also seen in conditions when agents have different *StressResponseThresholds (Low→High* and *High→Low*,

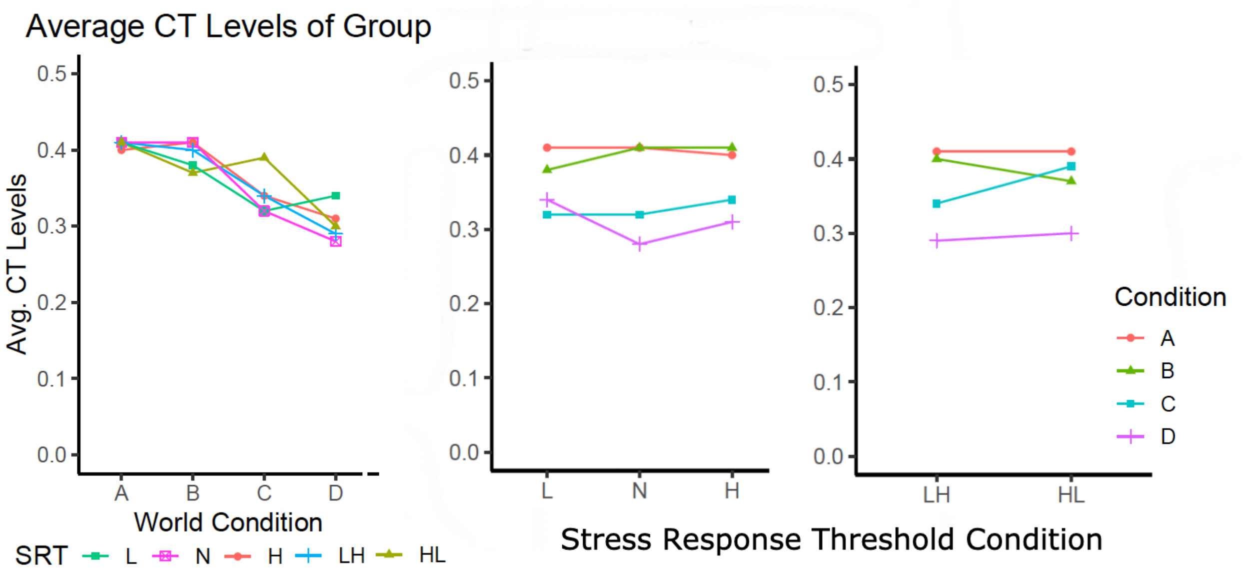

Figure 6.9 Mean cortisol levels of each group. Left: Cortisol levels across world conditions, broken down by *StressResponseThreshold* (SRT) conditions. Middle and Right: Cortisol levels across different *StressResponseThreshold* conditions, broken down by world conditions. Statistical significance results can be seen in Table G.1. L = *Low*; N = *Neutral*; H = *High*; LH = *Low* → *High*; HL = *High* → *Low*.

Figure 6.10, bottom).

In other words, a reduced physical stressor (increased food resources) resulted in larger variances in stress (cortisol) levels between agents in each condition. However, this variance was not affected by the different *StressResponseThreshold* conditions.

6.4.3 Social Interactions

Table 6.4 shows the *RateOfAggression%*—the "aggressiveness" of a society within each condition—at an aggregated group level, and the *AggressionFrequency* of each agent, across all simulation runs.

Results find that as food availability is increased, mean *RateOfAggression* of the overall society decreases, with a strong, statistically-significant correlation between these (r = 0.881, aggregated across *Low, Neutral, High StressResponseThreshold* conditions). We also note an inverse correlation between an increased tolerance to stress (*StressResponseThreshold*) and the *RateOfAggression* of the society for all agents (r = -0.956). This is despite similar CT levels (Table 6.3), Therefore, as group tolerance to stress increases (from *Low*, *Neutral*, and *High* conditions), aggression rates decrease at a group level.

World Condition	**Stress Response Threshold**	**A1**	**A2**	**A3**	**A4**	**A5**	**A6**
A	**Low**	0.47	0.52	0.51	0.64	0.60	0.66
	Neutral	0.49	0.53	0.58	0.53	0.57	0.63
	High	0.52	0.50	0.54	0.55	0.59	0.59
	Low→High	0.56	0.52	0.51	0.60	0.58	0.62
	High→Low	0.51	0.47	0.54	0.55	0.58	0.62
B	**Low**	0.46	0.51	0.54	0.61	0.66	0.66
	Neutral	0.47	0.54	0.57	0.61	0.64	0.66
	High	0.46	0.49	0.52	0.47	0.65	0.64
	Low→High	0.51	0.48	0.48	0.55	0.65	0.66
	High→Low	0.47	0.49	0.49	0.57	0.60	0.65
C	**Low**	0.38	0.44	0.54	0.56	0.62	0.65
	Neutral	0.37	0.45	0.48	0.55	0.59	0.62
	High	0.40	0.49	0.55	0.56	0.57	0.64
	Low→High	0.41	0.50	0.52	0.52	0.65	0.63
	High→Low	0.42	0.44	0.51	0.59	0.64	0.66
D	**Low**	0.27	0.37	0.50	0.55	0.62	0.69
	Neutral	0.30	0.40	0.42	0.50	0.54	0.62
	High	0.28	0.33	0.45	0.46	0.58	0.61
	Low→High	0.28	0.34	0.37	0.48	0.56	0.63
	High→Low	0.31	0.37	0.43	0.48	0.57	0.65

Table 6.3 Mean cortisol values (between 0 and 1) for each agent across all conditions, aggregated across all simulation runs.

With regards to *AggressionFrequency*, as food availability increases, *AggressionFrequency* of the most dominant agent A1 decreases (World Condition A: 8.4–25.7; D: 5.3–14.6). At the same time, *AggressionFrequency* increases for the middle-ranked agents such as A3 (World Condition A: 6.5–13.1; D: 9.4–18.9) and A4 (A: 7.0–15.1; D: 8.7–16.7). We suggest that the reduction of *Aggression* performed by the dominant agents A1 and A2 across world conditions is a result of their reduction in overall cortisol levels (Table 6.3), which contributed to the reduced *RateOfAggression* of these societies.

Focusing on the two conditions where agents have different *StressToleranceThresholds*—(*Low→High and High→Low*)—we saw no significant difference between *RateOfAggression* at an aggregated group level. However, we see differences in the agents accounting for the aggressive encounters in these two conditions.

In three of the four world conditions (A, B and C), we see a reduced *AggressionFrequency* from A1

World Condition	**StressThreshold**	**RateOfAggression%**	**A1**	**Δ**	**A2**	**Δ**	**A3**	**Δ**	**A4**	**Δ**	**A5**	**Δ**	**A6**	**Δ**
A	Low	30.3%	25.7	-	18.7	-	13.1	-	15.1	-	12.2	-	0.0	-
	Neutral	20.1%	12.7	*-13*	10.0	*-8.6*	10.0	*-3.0*	8.1	*-7.0*	5.2	*-7.0*	0.0	*0*
	High	15.9%	8.4	*-4.3*	6.8	*-3.2*	6.5	*-3.5*	7.0	*-1.1*	7.1	*1.9*	0.0	*0*
B	Low	31.5%	25.9	-	18.8	-	18.9	-	16.9	-	9.9	-	0.0	-
	Neutral	22%	14.9	*-11*	12.4	*-6.3*	10.0	*-8.9*	12.2	*-4.7*	7.1	*-2.8*	0.0	*0*
	High	16%	9.2	*-5.7*	8.8	*-3.6*	8.3	*-1.7*	8.9	*-3.3*	7.6	*0.5*	0.0	*0*
C	Low	28.5%	21.2	-	26.1	-	24.6	-	22.5	-	7.1	-	0.0	-
	Neutral	16.8%	12.5	*-8.7*	12.8	*-13.3*	10.4	*-14.2*	10.4	*-12.1*	6.9	*-0.2*	0.0	*0*
	High	14.2%	7.6	*-4.9*	8.8	*-4.0*	10.9	*0.5*	8.8	*-1.6*	5.2	*-1.7*	0.0	*0*
D	Low	23.0%	14.6	-	17.9	-	18.9	-	16.7	-	10.0	-	0.0	-
	Neutral	11.2%	7.8	*-6.8*	8.0	*-9.9*	7.8	*-11.1*	9.6	*-7.1*	4.8	*-5.2*	0.0	*0*
	High	9.4%	5.3	*-2.5*	5.9	*-2.1*	9.4	*1.5*	8.7	*-0.9*	5.8	*1.0*	0.0	*0*

Table 6.4 RateOfAggression% of overall society in each condition and Aggression Per 1000 Time Steps performed by each agent. The Δ value indicates the difference in Aggression Per 1000 time steps from the condition preceding it in the table. Green values indicate reduced *AggressionFrequency* vs. preceding *StressThreshold* condition, red values indicate increased *AggressionFrequency* vs. preceding condition.

World Condition	**StressThreshold**	**RateOfAggression%**	**A1**	**Δ**	**A2**	**Δ**	**A3**	**Δ**	**A4**	**Δ**	**A5**	**Δ**	**A6**	**Δ**
A	Low→High	19.9%	11.6	-	7.9	-	7.3	-	12.8	-	6.9	-	0.0	-
	High→Low	20.8%	9.8	***-1.8***	8.9	*1.1*	11.8	***4.6***	10.0	***-2.8***	8.6	***1.7***	0.0	*0*
B	Low→High	19.3%	9.7	-	9.3	-	11.7	-	10.4	-	8.8	-	0.0	-
	High→Low	18.1%	8.2	***-1.5***	12.0	***2.7***	10.1	***-1.5***	16.2	***5.8***	8.1	***-0.7***	0.0	*0*
C	Low→High	17.4%	8.4	-	13.3	-	13.5	-	13.0	-	9.8	-	0.0	-
	High→Low	18.1%	7.6	***-0.8***	10.4	***-2.8***	13.2	***-8.6***	16.9	***3.9***	11.1	***1.3***	0.0	*0*
D	Low→High	9.4%	4.5		6.8	-	6.3	-	10.7	-	6.4	-	0.0	-
	High→Low	11.8%	5.4	*0.9*	7.8	*0.9*	11.0	***4.7***	10.4	*-0.2*	8.6	***2.2***	0.0	*0*

Table 6.5 Aggression % of society within each condition and Aggression Per 1000 Time Steps performed by each agent for the two non-uniform *StressResponseThreshold* conditions. The Δ value indicates the difference in Aggression Per 1000 time steps from the condition preceding it in the table. Results in bold are visualised in the network graph (Fig. 6.11).

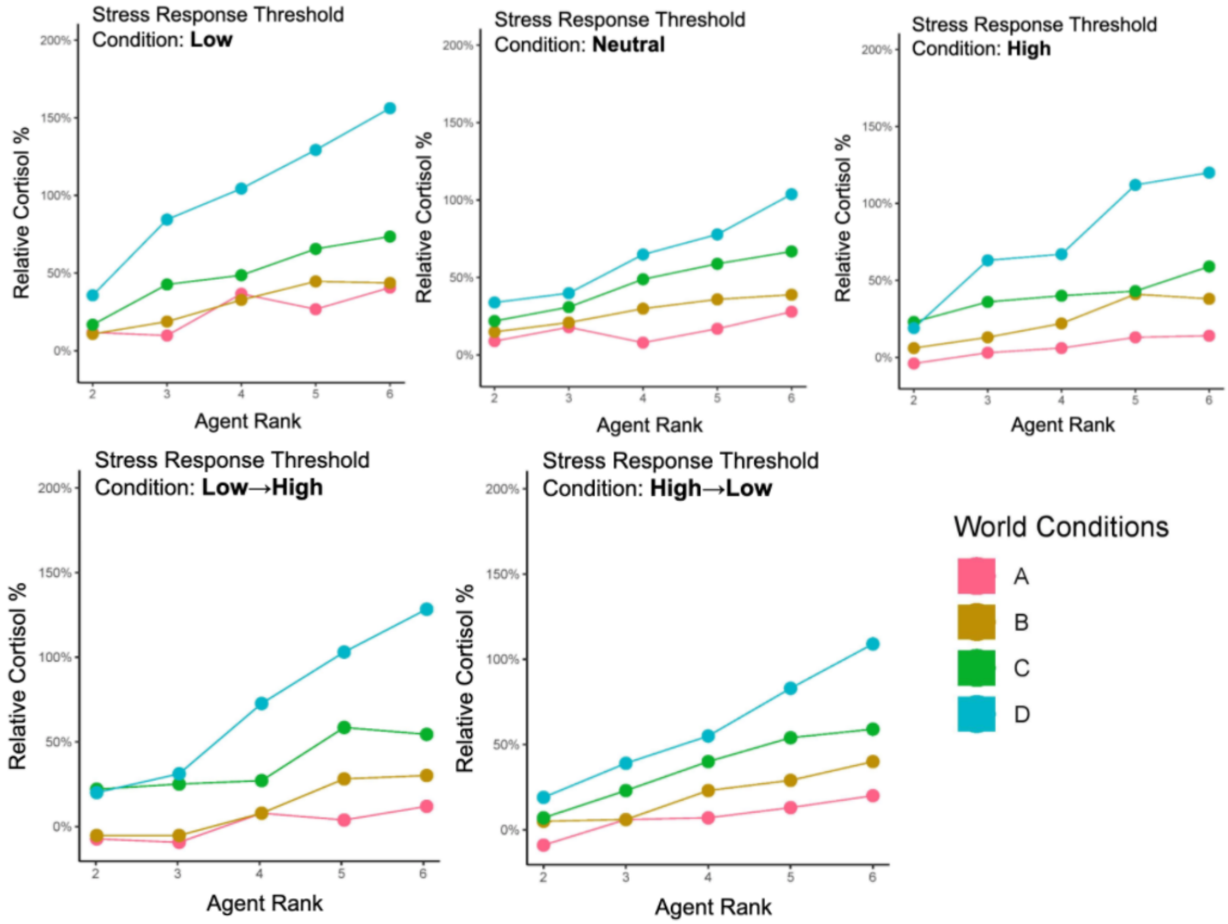

Figure 6.10 Relative differences in cortisol levels of Agents A2 to A6 (x-axis), as a percentage of the cortisol levels of A1 (not shown) in each condition. Relative cortisol levels are broken down for each *StressResponseThreshold* condition across all world conditions (line color).

(-1.8, -1.5 and -0.8 respectively). (Table 6.4). A reduction in aggression from A1 is offset by an increase in aggression from other agents (World Condition A: A3 (+4.6), A5 (+1.7); B: A2 (+2.7), A4 (+5.8), C: A4 (+3.9), A5, (+1.3), D: A3 (+4.7), A5 (+2.2)), resulting in societies reporting similar *RatesOfAggression*. Differences between aggression encounters in these two conditions are seen in visualised in the network graph in Figure 6.11.

In summary, the aggression of dominant agents (A1 and A2) contributes to the overall "aggressiveness" of societies when all agents have the same tolerance to stress (*Low, Neutral, High*, but not in conditions where *StressResponseThreshold* is determined by rank (*Low→High, or High→Low*). In these conditions, though societies show similar levels of aggregated group *RateOfAggression*, granular analysis highlights how the agents responsible for performing the aggression is different between conditions. These mixed findings

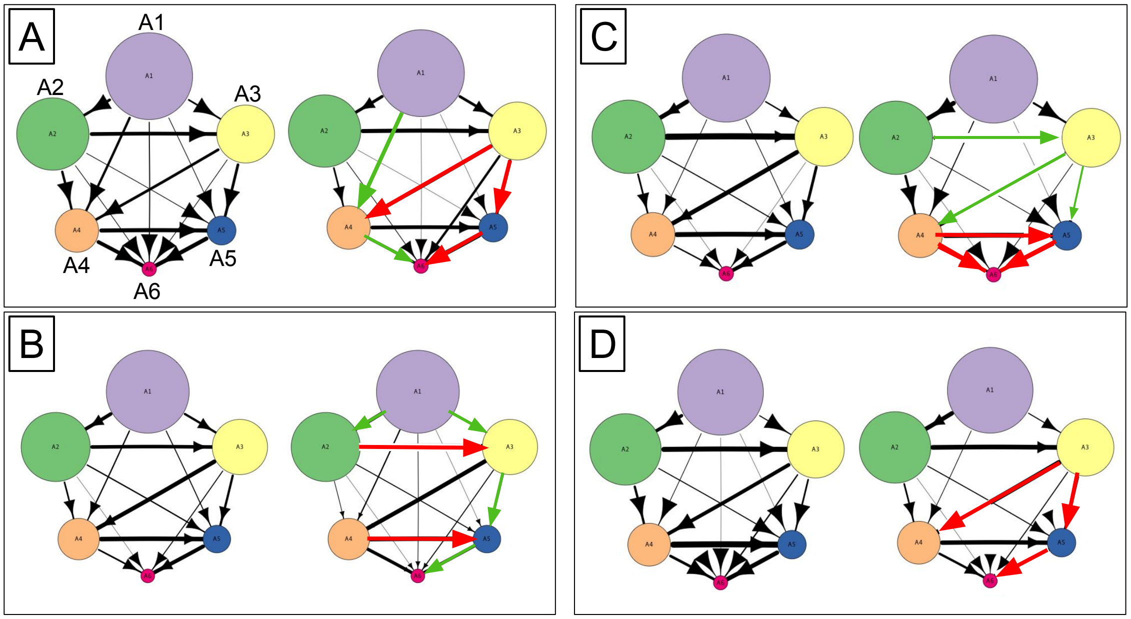

Figure 6.11 Distribution of aggression in the *Low* → *High* (Left) and *High* → *Low* (Right) *StressResponseThreshold* conditions, as a percentage of all aggression performed per condition. Thicker lines indicate higher percentage of aggression from dominant to subordinate agent. Notable increases and decreases of aggression frequencies compared to the *Low* → *High* condition are shown by red and blue lines (right), respectively (Table 6.4.)

show contextual support for our hypothesis that the interactions from dominant individuals determine social dynamics.

6.5 Discussion

Overall, our results showed contextual support for our hypothesis that rank-related responses to stress play a role in determining social dynamics and interactions. However, contrary to our hypothesis, we highlight that these dynamics are not simply dictated by dominant individuals. Instead, the individual stress tolerances throughout the entire society played a role in having an effect "up the hierarchical chain". We discuss some of our observations below.

We found no support for our hypothesis that group viability would be affected by individual, rank-related stress tolerances. Across all conditions, we found similar mean viability performance, regardless of the *StressResponseThreshold* condition.

Nevertheless, we found that both individual and group viability was improved across all conditions when cortisol was present in an individual's physiology, compared to control conditions when it was not. Including these hormonal mechanisms in our agent model improved viability on two of three measures (*Life Length* and *Physiological Wellbeing*), with no detrimental effect on *Average Comfort*. Therefore, although cortisol was added to this model as a stress-related hormone, its effects on physiological (in terms of its "metabolic" effects on movement and *Energy* usage) and behavioural adaptation (in terms of aggressive, socio-negative interactions) resulted in some viability benefits for our agents. Though we did not hypothesise this effect, this supports our suggestion from Chapter 5 where we considered how socio-negative interactions may be a viability-relevant behavioural adaptation). We consider the benefits of such mechanisms in the future development of our socially-adaptive model.

One important consideration of our results is that the Social Assessment Component, which determined whether agents would perform socio-positive or socio-negative interactions with others, used absolute values to determine which behaviours would be performed. Using a more probabilistic or stochastic approach to these behaviour rules (such as those used in similar work [72]) may have resulted in a different distribution of the types of interactions between agents. For instance, our results found how the "aggressiveness" of dominant agents contributed to the overall "aggressiveness" of the society in some conditions (i.e. when all agents had the same *StressResponseThreshold* value). We hypothesise that a stochastic rule in the Social Assessment Component would reduce the frequency of aggressive interactions from the dominant agents in these conditions, which would then further reduce aggressive interactions in the rest of the society. We reserve this line of research for future investigation.

Across the three uniform conditions related to stress tolerances (*Low, Neutral, High*), we found that both *RateOfAggression* and *AggressionFrequency* decreased across all world conditions as tolerance to stress increased for all individuals. This is despite the fact that mean cortisol levels remained similar throughout each of these conditions. These findings provide some support for the hypothesis that stress levels themselves may not be enough to determine social dynamics. Rather, social dynamics are determined by how tolerant different individuals in the society are to their own stress levels, and the availability of appropriate coping mechanisms.

In addition, cortisol for subordinate agents did not change across world conditions, despite easier physical conditions. We hypothesise that this was due to a lack of appropriate stress coping outlets for these agents, i.e. dominant agents could perform *Aggression* to reduce cortisol levels, but these outlets were not available for subordinates. This suggests that a less-challenging physical environment may not explicitely reduce stress for all individuals in a society the same way and that additional stress coping mechanisms may be required for some individuals to experience the same benefits.

Focusing specifically on the results from the the two conditions of non-uniform *StressResponseThresholds*, *Low→High* and *High→Low*, we found that, despite similar *Aggression* rates of the society, the agents responsible for performing this aggression varied between these two groups. When higher-ranking agents had higher tolerances to stress (in the *High→Low* condition), the majority of Aggression was performed by lower-ranking agents.

Similarly, higher-ranked individuals appeared to benefit from improved world conditions more than subordinate agents, resulting in significantly less *Aggression* from higher-ranking agents. Though we hypothesised that reduced *Aggression* from dominants would result in reduced aggression for the overall society, we found that this was not the case and, in some conditions, lower-ranking agents with lower tolerances to stress continued to perform large amounts of aggression towards other agents.

We also observe strong inverse correlations between an agent's rank and overall cortisol levels, which become more significant as environmental conditions improved. In our model, improved world conditions facilitated a reduction in stress for dominant agents, but not subordinates. This reduction in stress reduced their overall aggression and therefore overall aggression rates of the society.

In terms of real-world societies, these different social dynamics based on different tolerances to stress and physical world conditions may highlight some features which determine the dynamics of social species [117]. We discuss some of our findings along with qualitative observations below.

6.5.1 "Calmer" Dominants Results in Increased Aggression

To investigate our hypothesis that individual, rank-related responses to stress would affect social dynamics, we focus on the two rank-related *StressResponseThreshold* conditions: *Low→High* and *High→Low*, where stress tolerance was a direct and inverse function of rank, respectively (Table 6.2).

In these conditions, we found similar *RatesOfAggression* of both societies, regardless of environmental conditions (Table 6.5). This is despite similar cortisol levels of agents (Table 6.3). Further analysis found how the agents responsible for the aggression was different between conditions.

Despite fewer "available" social outlets for subordinate agents when compared to dominant agents (i.e. A4 has two subordinate agents to which it can perform *Aggression*, while A1 has five), subordinate agents still accounted for the highest *RatesOfAggression*. The limited number of agents for subordinates (e.g. A4) to perform *Aggression* on resulted in more frequent aggression focused on specific individuals (Fig. 6.11, Table 6.5).

Taking into account the similar cortisol levels between conditions (Fig. 6.10), and the shift in aggression further down the hierarchy, these results show some support for the hypothesis that it the individual tolerances to stress hormones (i.e. "personality types"), and not the level of stress hormone itself, that determines the dynamics of a society [108].

However, we had hypothesised that these dynamics would be dictated by the interactions and stress response of dominant individuals. These results indicate that this may not be the case. Instead subordinate responses to stress played a significant part in determining how aggressive a society may be perceived. Specifically, a "relaxed" society may only be so if all individuals in that society tolerate stress at a similarly-high level (Table 6.4). Individual "personality" differences (i.e. if some individuals are more or less "anxious" than others) may disrupt the balance of an otherwise-relaxed society.

Though we do not include changes in dominance rank in our model, this distribution and density of aggressive encounters may have implications on the stability of a society (i.e. where increases in aggression between agents may be seen as rank challenges). Extending this model to account for these changes may help us further understand how individual differences can have a "bottom-up" effect on emergent social structure and hierarchy. Nevertheless, given the biological grounding of our model, our results and model may be able to guide ethologists into how and where rank challenges are likely occur in biological systems.

6.5.2 Increased Stress Results in Food Theft

One reason as to why the physically-challenging environments (World Conditions A & B) were made increasingly difficult for subordinate agents through the emergence of food theft from dominants.

Previous chapters (Chapters 3, 4 and 5) used a simplistic model of "food sharing" where multiple agents could eat at a food resource with no direct detriment to either party. In our observations, when cortisol levels were low and agents were not stressed, food sharing was observed (Figure 6.12, top).

However, when cortisol levels of agents were elevated and dominant agents were stressed, the intensity of the tactile interaction (Equation 6.2.2), resulted in dominant agents performing *Aggression* and "pushing" subordinate agents away from the food resource that it was eating; stealing the food resource. As these subordinates are then unable to approach the resource, they would need to spend additional time and *Energy* locating another food resource.

For dominant individuals, this displacement of stress through aggression provided an advantageous effect on viability—through the monopolisation of a food resource—with a deleterious effect on subordinates, who suffered both an increase in cortisol, which also increased *Energy* depletion, and would have to spend time attempting to locate another food resource. From the perspective of the dominant agent in this situation,

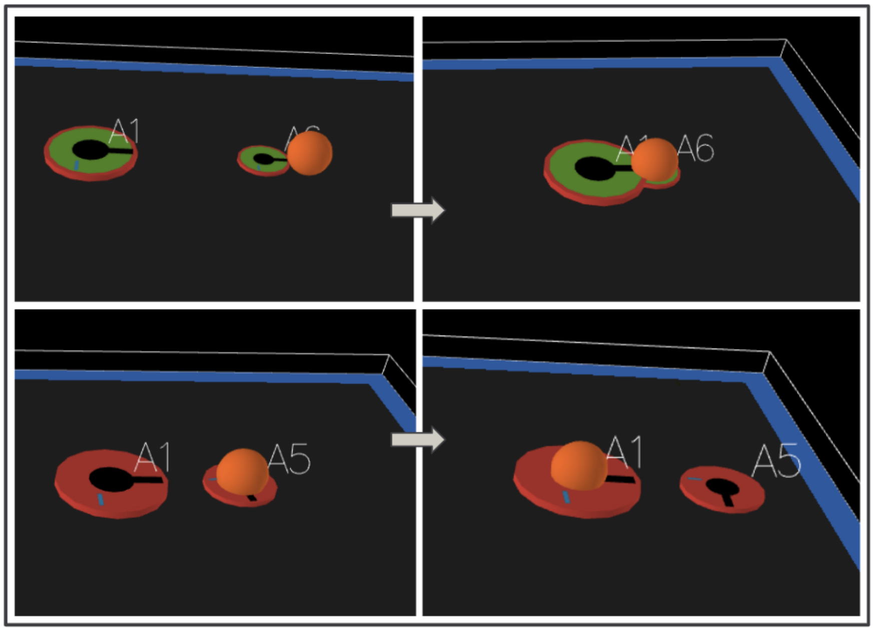

Figure 6.12 Screenshots showing food sharing behaviours when dominants are not stressed (top), and "theft" of food resources by dominants when they are stressed (bottom). Green agents (doughnut-type shapes) are not stressed; red agents are stressed. Top: Agent A5 is eating at the food resource (orange sphere), and a non-stressed A1 approaches to share the resource. Bottom: Agent A6 is eating at the food resource (left), but as A1 is "stressed", the intensity of the *Aggression* pushes A5 away from the resource, where A1 then eats. A5 then needs to locate a new resource.

such interactions are in line with our suggestion from Chapter 5 that socio-negative behaviours may be required to facilitate viability in some situations.

This socio-negative behaviour resulted in a more stressful environment for subordinates, and highlighted how how stress-related interactions from dominants could directly affect the survival of subordinates. To investigate the precise impacts of these interactions, future work should consider analysing the specific relationship between the amount of aggression received (including aggression received at food resources) and its effect on viability.

6.5.3 Easier Environments Reduce Stress for Dominants, not Subordinates

We found that the correlation between dominance rank and cortisol levels strengthens as environmental conditions provide more food resources Table 6.3. Conditions with poor food availability (A and B) may reflect uncertain or unstable environmental conditions for our agents, where even the dominant individuals do not possess large amounts of control. As a result of the lack of predictability of the environment, cortisol levels remain chronically elevated in all agents—even dominants—putting them in a state of stress. When all individuals in the society have *Low* tolerances to stress, this stressed state creates higher rates of *Aggression* towards other agents (Table 6.4).

We then expected cortisol levels for all agents to reduce as food availability increased, reducing the stress associated with limited resource availability (Equation 7.2). Although mean cortisol levels were lower for the overall group (Fig. 6.9), this reduction was driven by a reduction in stress of dominant agents. Despite the "easier" physical conditions, cortisol levels remained elevated for subordinate individuals across all conditions. We found that improved physical conditions were still not sufficient to reduce stress for agents who did not have an adequate amount of social coping mechanisms in this rank-based society (where rank plays an important role in resource acquisition).

To summarise, improved physical conditions appear to alleviate stress levels for the dominant individuals (A1, A2), but not subordinates (A5, A6) who lacked the level of social "control" as the dominants. These results support the suggestion by [115], [117], [194], who posit that psychological and social stress may play a greater role in determining stress levels than the physical conditions, especially for subordinate individuals who may require additional social support. We take these findings into account and highlight them for further investigation.

6.5.4 Bottom-Up Stress Contagion: Stressed Subordinates Cause Stressed Dominants

Despite our hypothesis that group stress levels and dynamics would be determined by the behaviours and interactions of dominant agents (A1 and A2) on the rest of society, we found that in some conditions—when food availability was poor and stress tolerance was *Low*—it was subordinate agents, not the dominants, that played a significant role in the emergent social dynamics.

In most experimental runs, cortisol levels of lower-ranked agents would often become elevated before that of dominants (Appendix A). Since elevated levels of cortisol increased movement speed (Equation 6.4), and stressed subordinate agents avoid dominants (Figure 6.4), it then became challenging for dominant agents to *Groom* subordinates in order to satisfy their *SocialNeed*, as the latter would continue to *Avoid* dominants while they remain stressed (Figure 6.13, B).

This resulted in dominants "chasing"—and failing to catch—subordinate agents moving at faster speeds (Figure 6.13, C). An inability to satisfy their *SocialNeed* resulted in increased stress levels of these dominant agents, causing them to also become stressed (Figure 6.13, D). As a result, dominant agents then perform

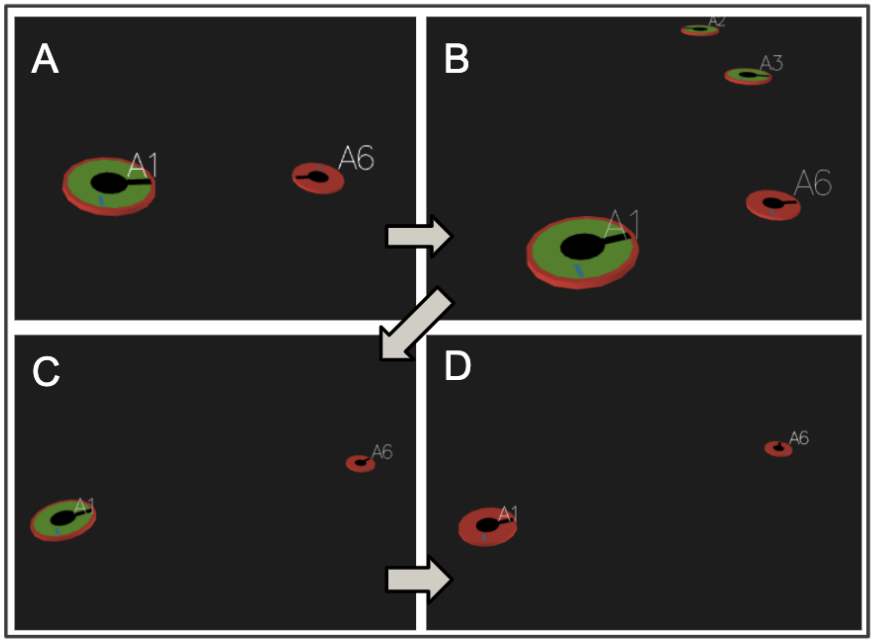

Figure 6.13 Screenshots showing how stressed subordinate agents caused dominants to become stressed. Green agents (doughnut-like shapes) are not stressed. Red agents are stressed. A: A1 (not stressed) begins to approach A6 (stressed) in order to *Groom*. B: The stressed A6 perceives the dominant A1 and *Avoids* it as a result of its stressed state. C: As higher levels of cortisol increase movement speed (Equation 6.4), A6 is able to move away from A1, resulting in larger distance between the two. D: Shortly after, as A1 is unable to satisfy its internal *SocialNeed*, it also becomes stressed.

Aggression on others to displace their stressed state (see Section 6.2.2). Therefore, in these conditions, the behaviour of stressed subordinates resulted in stressed dominants.

When subordinates have higher tolerances to stress than dominants (the *Low→High* conditions), they would be less likely to undergo a state of stress, resulting in less avoidance of dominants and therefore less "chasing". As a result, we saw lower *AggressionFrequency* from dominants (A1: World Condition D, *StressResponseThreshold* = *Low*: 14.6; World Condition D, *StressResponseThreshold* = *Low→High: 4.5*), despite dominant agents having the same tolerance to stress in both of these conditions.

This was contrary to our hypothesis. Rather than a stressed, dominant individual dictating the social dynamics via a "trickle–down" effect of stress and aggression, in these conditions, it was "stressed" subordinates that affected the subsquent social dynamics.

Though this was contrary to our hypothesised direction of effect, these findings further support the hypothesis that individual differences to stress affect social dynamics [108], and that—in a small society of agents such as the one investigated here—even the lowest-ranking agents can play a role in shaping overall group dynamics.

In terms of our investigation, we highlight the emergence of affective contagion resulting from the interactions of one individual within the society. Such observations may also highlight the importance of accounting for and analysing the independent behaviours of all agents in a society.

6.6 Summary

In this chapter, we investigated how different tolerances to stress—including rank-related tolerances to stress—would affect the social interactions, group dynamics and overall wellbeing of a small rank-based society. We considered this stress tolerance to be a type of "personality" trait. Using our observations from Chapter 5, we hypothesised that social dynamics and group viability would be determined by the stress tolerance and interactions of dominant individuals in a society. Specifically, we had predicted that an "aggressive" dominant individual would cause an "aggressive" society via a "trickle-down" effect of stress and aggression.

To investigate our hypotheses, we introduced a stress-related hormone, cortisol, to our agent model—released as a function of physical and "psychological" stress—which played an adaptive role on physiology and behaviour. We endowed each individual with one "personality trait" associated with their tolerance to this stress hormone: where individuals would become "stressed" if cortisol exceeded this value.

We tested our hypothesis using our society of six agents, each with a different dominance rank, associated with increased access to resources. We tested six different conditions of stress tolerance—three where all individuals had the same tolerance to stress, and two where this tolerance was both directly and inversely associated with rank—across four different world conditions related to food availability of decreasing challenge.

Our results found how these strictly rank-based social dynamics and interactions were affected by both environmental challenges, and these individual "personality" differences related to how stress tolerance. We found support for the hypothesis that tolerances to stress, and not stress levels themselves, plays a role in determining social dynamics. However, we found no support for our hypothesis that different rank-related tolerances to stress would affect group viability.

We found survival-related benefits for individuals endowed with a stress hormone (cortisol) released as a function of physiological, environmental, and social stress, but which provides several adaptive effects on behaviour and physiology. As cortisol is one key hormone that underpins adaptation through social allostasis [29], we consider retaining these effects for future development of our agent model.

Despite easier physical conditions reducing stress (cortisol) levels for dominant individuals, we found how these effects were not experienced by subordinates, who instead reported elevated cortisol levels across all conditions. We suggested that one reason for this may be due to the lack of social coping mechanisms or appropriate outlets to reduce stress. Social coping mechanisms (such as opportunities to perform aggression, the presence of close-kin or bond partners) plays a significant role in ameliorating an individual's stress response [108], [117]. Currently, our model only accounts for the stress-inducing effect of other individuals in the society, with stress-reducing social coping mechanisms only available to dominant individuals (via aggression).

In addition to providing direct outlets for socio-positive interactions (such as grooming), the presence of social support has also been hypothesised to provide additional, stress-reducing effects for supported individuals through the mediation of the stress response in several ways. This has been coined the "social buffering" hypothesis of social support, and has been suggested to play a significant role in facilitating the healthier, longer lives associated with living in a society. In addition, these stress-reducing "buffering" effects of social support appear to be facilitated by oxytocin, though its precise effects are still not fully understood.

By reintroducing the oxytocin-mediated, affect-based social bonds from Chapter 5 in the remaining chapters, we aim to investigate how some of the hypothesised, stress-reducing effects of social support can affect the wellbeing and social interactions of individuals in this rank-based society. In doing so, we continue to investigate such mechanisms, proposed by affect-based social relationships, in the development of our social allostasis-inspired agent model.

Chapter Seven

The Social Buffering Hypothesis and the Stress-Reducing Effects of Social Support

7.1 Introduction

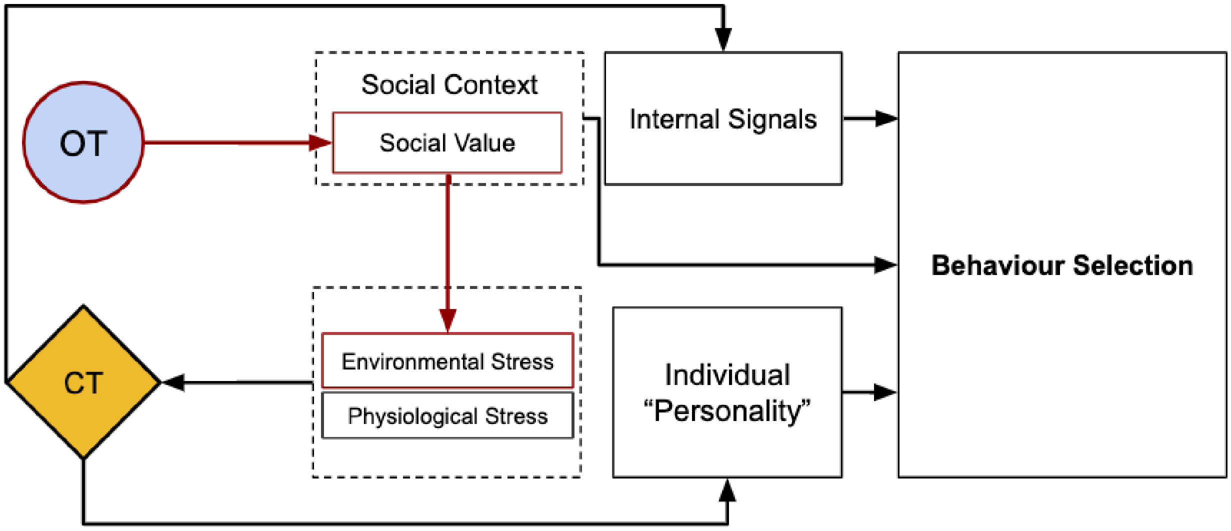

Figure 7.1 A high-level view of the effects investigated in this chapter. Red lines indicate extensions from the previous chapter. This investigation re-introduces oxytocin's (OT) effects on social partner preference (Chapter 5) which also has a stress-reducing effect. The rest of the model remains the same as Chapter 6. A more granular view of this model can be seen in Figure 8.1.

In the previous chapter, we had investigated the social interactions and dynamics in a society of agents when each individual had a different “personality” related to their tolerance to stress. We found that, for strictly-rank-based social interactions, social dynamics were affected by both the challenge of the environment, as well as the personality differences in the society. In these societies, we found how the stress levels of subordinates was significantly higher when compared to dominant individuals, even in easier environmental conditions.

We hypothesised that one reason for the increased stress was the difference in available social coping outlets available to individuals of different social rank. In the real world, social interactions are rarely ever based solely on rank. Instead, the formation and maintenance of close social support (such as bond partners) plays a role in determining social interactions and plays one part in the adaptive process.

The stress, fear, and "nervousness" associated with challenging environmental and social conditions may be offset through (the presence of) social support. The presence of social support, underpinned by oxytocin, has been found to have an anxiolytic, "calming" effect on individuals by reducing, or even eliminating, their stress response in both present and upcoming stressful situations. This phenomenon, known as "social buffering", plays a critical role in the longer-term wellbeing for social animals through adaptation to the social environment, and is a key aspect of social allostasis.

In this chapter, we build upon the previous investigation where social interactions were strictly rank-based. By introducing social bond partners, we investigate one of the stress-reducing effects of affect-based social support on the interactions and wellbeing of the society across several dynamic, physical environments. We investigate several combinations of social bond partners, related to bond partner rank, and two types of affective bonds with respect to bond strength (i.e. "fixed" and "variable"). "Fixed" bonds remain persistently strong throughout the lifetime of the bonded partners, whereas "variable" bonds are susceptible to deterioration when they are not renewed and can be strengthened or damaged as a function of social interactions. Therefore, we investigate the stress-reducing effects of both the presence and quality of affective bond partners.

Using the social buffering hypothesis along with our observations from Chapter 5 which saw viability benefits from the presence of social support, we hypothesise that the presence of social bond partners will result in viability-related benefits for socially-bonded individuals. We also hypothesise that these improvements in viability will be facilitated by a reduction in stress and socio-negative interactions.

Our results find that some support for this hypothesis. Specifically, that the stress-reducing effects of social support is dependent upon the wider social and environmental context, as well as the different type of affective bonds between individuals. In some conditions, we find improvements in viability and reduced stress for the wider social group, suggesting a wide-spreading, stress-reducing effect of social support in this small society.

7.1.1 Social Support and the “Social Buffering” Hypothesis

Positive social relationships play a significant role for the wellbeing of social creatures. For a variety of species, individuals who maintain close social bonds are associated with longer, healthier lives [18]–[20]. One way in which these relationships provide these beneficial effects is through a phenomenon known as “social buffering” [179], [202]. The social buffering hypothesis posits that social support provides a physiological and psychological “safety cushion” during stressful events for individuals with social support, reducing their response to stress, and adapting their perception, behavioural or emotional responses to future stressful situations.

Individuals with perceived social support have consistently been found to have reduced cortisol levels and improved cardiovascular health, compared to individuals who are not socially supported [203]–[206]. Additionally, stress responses in future stressful, unknown situations have also been found to be reduced in individuals who receive consistent social support [16], [207], potentially by perceiving the situation as less stressful or less threatening [208].

These findings suggest an anxiolytic effect of social support, likely as a result of one—or several—of oxytocin’s effects on physiological responses to stress-related hormones [209]–[211]. Beyond its physiological effects, oxytocin and social support also provide additional outlets for stress-reducing behaviours, by granting supported individuals opportunities to perform prosocial behaviours (such as friendly contact with a bond partner), further releasing oxytocin and strengthening the relationship between them, and thereby providing a feed-forward buffering effect on the stress response. Adapting to stressful situations in this way has survival-related physiological benefits: chronically-elevated levels of cortisol results in a higher susceptibility to disease and physiological dysregulation (Section 1.2). Therefore, the ability to reduce, and appropriately respond to, stress, can mitigate its maladaptive consequences, and be seen as a fundamental type of adaptation that facilitates longevity and wellbeing during, and in anticipation of, stressful or unknown situations.

7.1.2 Aims & Hypothesis

By accounting for one of the hypothesised effect of the “social buffering” hypothesis—that social support reduces the stress (cortisol) response—we aim to investigate the role of the presence of different types of affect-based social support on the wellbeing and social interactions of a rank-based society across environmentally-challenging, dynamic conditions.

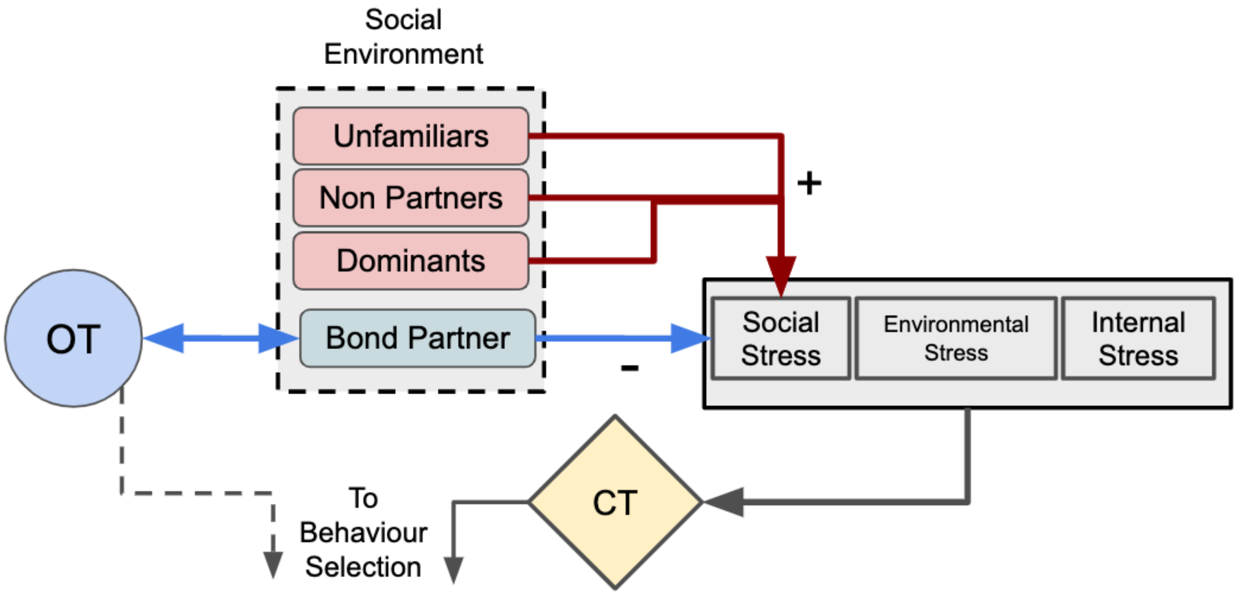

Figure 7.2 A high-level view of one potential effect of social buffering on reducing social stress, with an example of some causes of social stress from the social environment. While social stimuli such as unknown strangers, non-bonded, and dominant individuals can increase stress (red lines), the presence of a bond partner (blue line) can provide a stress-reducing effect. Our model does not currently account for unfamiliar individuals in the social environment, but we include them in this diagram as a potential stress-inducing presence. OT = Oxytocin, CT = Cortisol.

We test these effects across three different bond partner combinations related social rank, and two types of affective bonds with respect to bond strength (fixed and variable). We conduct these experiments across three different environmental conditions: one static, where food availability is consistent, and two dynamic conditions, where food availability increases and decreases over time.

Building on our findings of previous chapters (Chapters 5 and 6) and the findings from the natural world, we hypothesise that the presence of affect-based social bonds will provide survival-related benefits for individuals with these social bonds across these physically-challenging conditions. Furthermore, we hypothesise that this improved survival will be facilitated by a reduction in stress and socio-negative interactions between bonded partners.

We also aim to use this investigation to explore the effects of one type of adaptation through the social environment, in order to further the development of our social allostasis-inspired adaptive model.

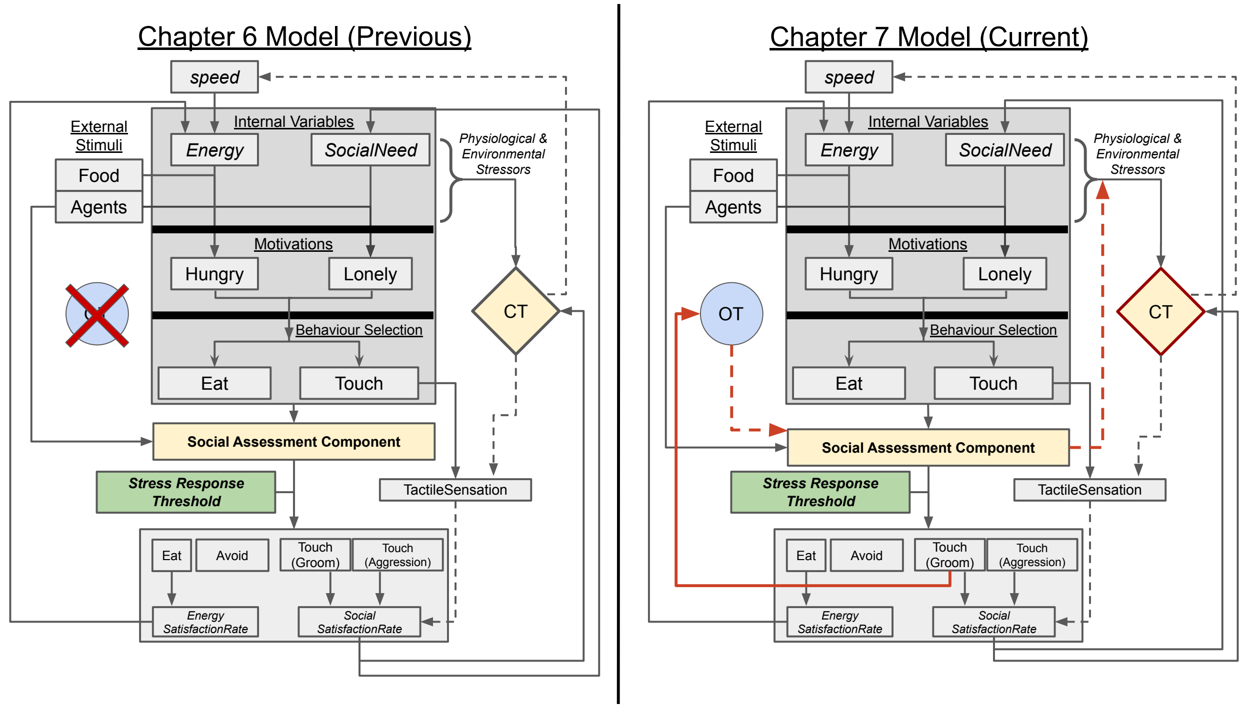

Figure 7.3 Illustration of the Action-Selection Architecture model used in this investigation (right) compared to the previous chapter (left). Extensions to the previous model are denoted in red. Oxytocin's (OT) effects on social partner preference from Chapter 5 are reintroduced in this model. "Social buffering" effects are calculated using the presence and strength of bond partners via the *Social Assessment Component* (Equation 7.5). OT is contextually-released through *Groom*.

7.2 Extensions to the Agent Model

7.2.1 Re-Introduction of OT and the Social Assessment Component

The model in this investigation reintroduces the Social Assessment Component (Figure 5.3) from Chapter 5, including the mechanisms of oxytocin (OT) and the Dyadic Strength Index *(DSI)* which describe the strength of the affective relationship between a pair of agents. We note a minor change to the release rate of oxytocin, though the effects of both OT and DSI mechanisms remain unchanged.

7.2.2 Oxytocin Release Rate

In previous investigations, oxytocin was released at a static value, depending upon the social behaviour being performed. For consistency in the way that we modelled hormonal release, the rate of oxytocin release is now a dynamic value, modulated by the intensity of the winning social behaviour *Touch*. This is in line with the modelling of cortisol, and is calculated as follows:

$$OT_t = OT_{t-1} + (OT_{\mathrm{d}} \times \mathrm{b}_{\mathrm{touch}}) \tag{7.1}$$

Where b_{touch} is the intensity value of the winning behaviour. We validated this new calculation across 50 experimental runs across all world conditions from the previous chapter with no significant differences on agent viability or mean oxytocin levels, and therefore retained this calculation for future investigations.

7.2.3 Stress Reduction via Affective Bonds

The stress-reducing effects of social support are calculated through the *AgentVal* calculation in the Social Assessment Component. To recall, CT is released with the following function:

$$CTReleaseRate_A = (\frac{\sum\limits_{i \in (P)} d_i}{2} - \frac{\sum\limits_{i \in (S)} \hat{S}_i}{2}) \times c \tag{7.2}$$

This rate accounts for both the internal stress, associated with the sum of physiological deficits $\sum d_i$:

$$\sum d_i = d_{energy} + d_{socialNeed} \tag{7.3}$$

and external stress $\sum \hat{S}_i$, associated with the lack of resource availability:

$$\sum \hat{S}_i = \hat{S}_{\mathrm{agents}} + \hat{S}_{\mathrm{food}} \tag{7.4}$$

As OT and affective bond partner presence (and quality) are included in determining “available” social resources when calculating *CTReleaseRate*, perceiving a bond partner reduces the stress response associated with stressful environments, by either reducing the amount of cortisol, or slowing its release rate. This is accounted for using the *AgentVal* (Equation 7.5:

$$\mathrm{AgentVal} = \mathrm{rank}_{\mathrm{diff}} + (\mathrm{bond}_{AB} \times (\mathrm{DSI}_{AB} \times OT)) \tag{7.5}$$

Therefore, in Equation 7.4, $\hat{S}_{\mathrm{agents}} = C_{\mathrm{agents}} \times (1 - AgentVal)$ and

$$\hat{S}_{\mathrm{food}} = \begin{cases} C_{\mathrm{food}} & \text{if } AgentVal \leq 0 \\ \text{else } 0 \end{cases} \tag{7.6}$$

In Chapter 6, the lack of social bonds resulted in *AgentVal* returning the normalised rank difference between two agents ($\mathrm{rank}_{\mathrm{diff}}$). Accounting for social bonds in this model, *AgentVal* is calculated using these additional parameters (OT and DSI), which adapts the stress response accordingly.

7.2.4 Fixed and Variable Bond Types

In our model, bond type is defined by the value of the bond strength (DSI; the *Dyadic Strength Index*) between two bonded agents.

“Fixed” bonds denotes a bond that remains persistently strong, with the DSI remaining at its maximum value of 2 throughout the simulation run.

“Variable” bonds denotes a bond with a changing DSI: experiencing a decay rate ($\mu = 0.9997$) over time in absence of social interaction, and subject to both strengthening and loss as a function of social interactions:

$$DSI_{AB} = \begin{cases} DSI + (TactileIntensity_A \times c) & \text{if } Groom \\ DSI - (TactileIntensity_A \times c) & \text{if } Aggression \end{cases} \tag{7.7}$$

7.3 Experimental Set Up

Experimental conditions were defined by three different parameters: three experimental conditions related to the presence (and type) of affective social bonds, four conditions related to the combination of socially-bonded agents, and three different environmental conditions related to food resource availability. Table 7.1 shows these different experimental conditions, and Figure 7.4 illustrates the bond combinations and types used.

We conducted three different experiments with respect to the presence and type of affective bond (Subsection 7.2.4). Experiment 1 served as our control group where all agents were unbonded. In Experiment 2, called the *Fixed Bond* condition, affective bond strength was fixed at its maximum value throughout the experiment (DSI = 2). In Experiment 3, the *Variable Bond* condition, the affective bond strength (DSI) is initialised at its maximum value of 2, but is subject to change as a function of social interaction, as well as decay in absence of social interaction (Subsection 7.2.4).

In Experiments 2 and 3, we tested three different bond combinations of agents: a parent-offspring relationship (A1-A2-A6), close-kin middle-ranked (A3-A4-A5) bonded group, and close-kin lowest-ranked (A4-A5-A6) bonded group. These bond combinations were biologically-inspired relationships, and taken from our investigation in Chapter 5 and can be seen in Figure 7.4.

To introduce uncertainty into the environment, experiments were conducted in two dynamic world conditions related to food availability, as well as one static environment to serve as our control condition. We therefore move away from the static environments used in previous investigations.

Experiment	Bond Type	Bond Combinations	World Conditions
1	None	-	Static / Seasonal / Extreme
2	Fixed	A / B / C	Static / Seasonal / Extreme
3	Variable	A / B / C	Static / Seasonal / Extreme

Table 7.1 Table of experimental combinations in this chapter. Bond combinations can be seen in Figure 7.4. World conditions can be seen in Figure 7.5. Experiment 1 has 3 total experimental combinations, and Experiments 2 and 3 have 9 experimental combinations each.

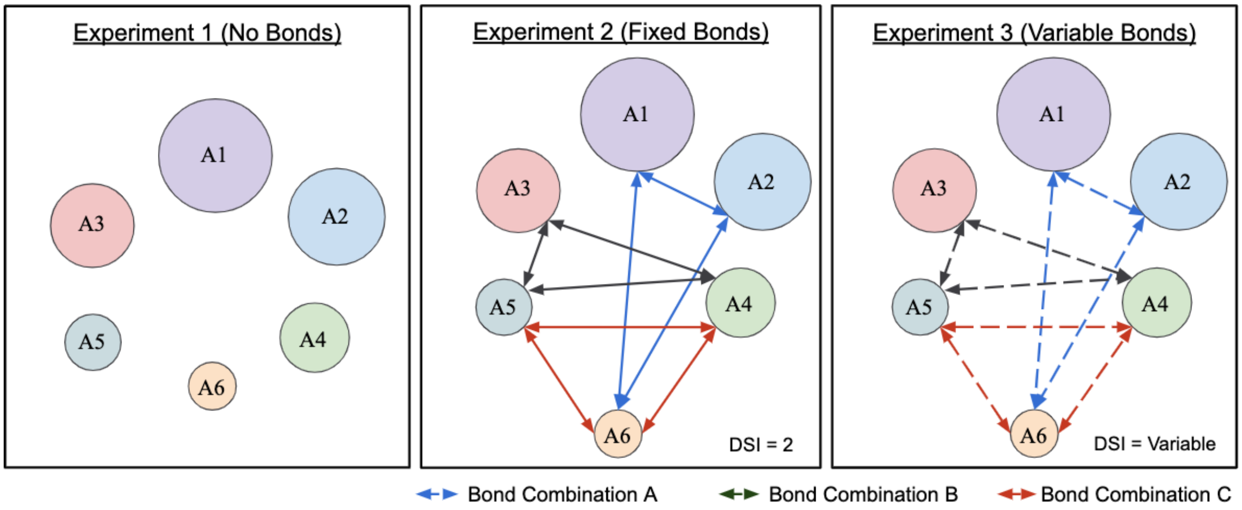

Figure 7.4 Illustrations of the three different types of affective bonds used in each experiment. Experiment 1 had agents with no bonds. Solid lines illustrates a fixed bond strength (DSI remains a fixed value of 2). Dashed lines illustrate variable bond strength (DSI changes over time).

The *Static* environment contained four food resources throughout the experiment, retaining the same set up as World Condition D as seen in Chapter 4 (four food resources).

In our first dynamic condition, called *Seasonal*, these four food resources steadily decreased and increased across fixed time periods (every 1000 time steps), as the world condition incrementally cycles between periods of very good (four) to very poor (one) food availability (Figure 7.5, orange arrows).

The second dynamic condition, called *Extreme*, food availability cycles between four food resources to a single food resource every 1000 time steps (Figure 7.5, blue arrows). Table 7.2 shows the number of food resources available in each cyclical change across all three environments.

The decision to model the dynamic world conditions in this way were inspired by the natural world. The *Seasonal* condition, as the name implies, was inspired by the natural seasons as it moved cyclically from world conditions of abundant food resources (for instance, Spring or Summer) to more extreme conditions

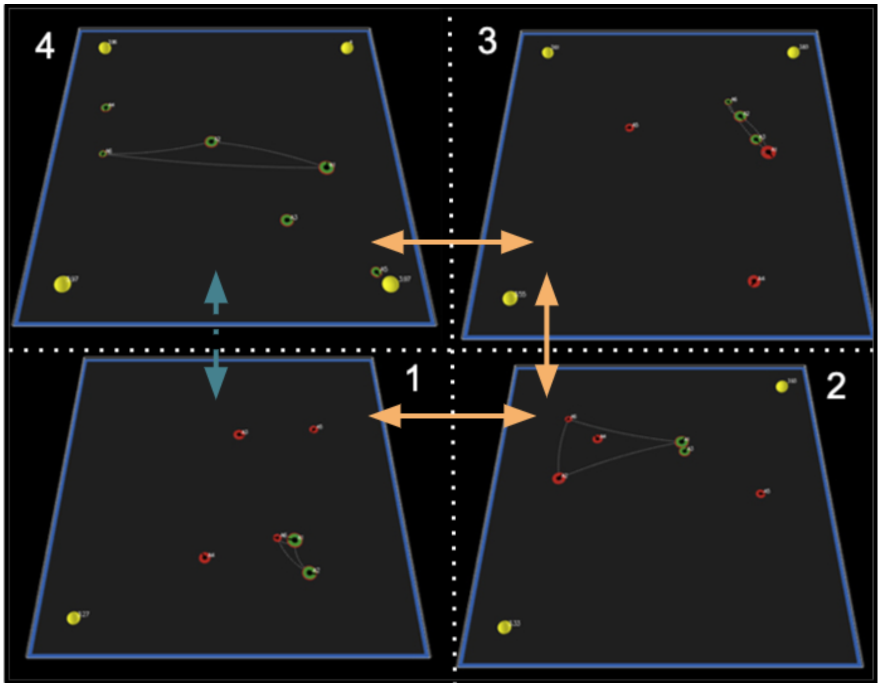

Figure 7.5 Screenshots of different food availability in each of the world conditions. *Static* environments remain fixed with four food resources (yellow discs, top left). The orange arrow denotes changes in the *Seasonal* condition. The blue arrow denotes changes in the *Extreme* condition. Details of each food phase can be seen in Table 7.2. Agents are rendered as doughnut-type shapes. Colour indicates their "stressed" state (green = not stressed, red = stressed). A dashed line between agents denotes an existing affective bond between them.

(Winter). The *Extreme* condition, on the other hand, was inspired by more dramatic changes (for example, as a result of a storm). We also took inspiration from similar experiments [40] who also performed investigations using similar dynamic conditions to assess the performance of their model.

In total, these combined parameters created 21 different experimental combinations: 3 in Experiment 1, 9 in Experiment 2, and 9 in Experiment 3 (Table 7.1). For each of these experimental conditions, we performed 20 simulation runs for 15,000 time steps each. This equated to approximately 10 hours of run time.

Results are broken down by the three bond type conditions (Experiment 1: No Bonds, Experiment 2: Fixed Bonds, Experiment 3, Variable Bonds). We measured results across the three viability-related metrics: *Life Length*, *Average Comfort*, and *Physiological Wellbeing* (Section 2.4.1). We also reported results on the distribution of social behaviours (*Grooming* and *Aggression*) throughout all simulation runs. We also measured internal hormonal levels (oxytocin and cortisol) for each agent at each time step. We complemented our quantitative results with personal observations of simulation runs.

Time Step	**Start**	0	2001	3001	4001	5001	6001	7001	8001	9001	10001	11001	12001	13001	14001
	End	2000	3000	4000	5000	6000	7000	8000	9000	10000	11000	12000	13000	14000	15000
World Condition	**Phase**	**1A/1B**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
Static		4													
Seasonal		4	3	2	1	2	3	4	3	2	1	2	3	4	3
Extreme		4	1	4	1	4	1	4	1	4	1	4	1	4	1

Table 7.2 Availability of food across different phases in each of the world conditions. This is also illustrated in Figure 7.5. Food availability is colour-coded for visualisation purposes: green indicates periods where maximum number of food resources (4) are available, and red indicates periods where minimum number of food resources (1) are available. Though the simulation ran for 14 phases, results are reported for phases where agents remained viable (1-12).

Statistical significance testing was performed between-groups using one-way ANOVA testing. We also used Pearson's Correlation Coefficient to report correlation strength between metrics where appropriate.

7.4 Results

7.4.1 Experiment 1 - No Bonds

Viability Indicators

In the control conditions (Experiment 1) when all agents were unbonded, viability performance of the overall society was greatest in the *Static* environments (Life Length (LL): 33%, Physiological Wellbeing (PW): 0.75). The results from the two dynamic world conditions were similar for both of these viability indicators (LL: 22%, PW: 0.49, 0.51 for *Seasonal* and *Extreme* conditions respectively). We observed no significant difference in Average Comfort across all world conditions.

Across all world conditions, we observed a moderately-strong correlation between an agent's rank and their mean Life Length (*Static*: r = 0.57, *Dynamic*: r = 0.64, *Extreme*: r = 0.68). Similar correlations were also observed for the individual agents' Physiological Wellbeing (r = 0.56 , r = 0.60, r = 0.65).

In sum, across all world conditions where social bonds were not present, agent viability was strongly correlated with their social rank.

Hormonal Trends

Table 7.3 shows the mean cortisol levels across all experimental conditions. In control conditions, we observed that mean group cortisol levels increased with world challenges (*Static*: 0.42, *Seasonal*: 0.50, and *Extreme*: 0.53). On an individual agent level, there was a strong correlation between mean cortisol levels and individual

		No Bonds (Experiment 1)						**Fixed Bonds (Experiment 2)**						**Variable Bonds (Experiment 3)**					
World Condition	**Bond Combination**	**A1**	**A2**	**A3**	**A4**	**A5**	**A6**	**A1**	**A2**	**A3**	**A4**	**A5**	**A6**	**A1**	**A2**	**A3**	**A4**	**A5**	**A6**
Static	**A**	0.27	0.31	0.40	0.50	0.58	0.66	0.16	0.16	0.31	0.39	0.49	0.36	0.21	0.24	0.36	0.43	0.52	0.55
Seasonal		0.49	0.53	0.61	0.60	0.65	0.74	0.27	0.29	0.45	0.5	0.61	0.57	0.42	0.43	0.48	0.53	0.56	0.48
Extreme		0.51	0.56	0.66	0.68	0.75	0.81	0.45	0.44	0.53	0.61	0.66	0.68	0.43	0.48	0.54	0.52	0.64	0.58
Static	**B**	0.27	0.31	0.40	0.50	0.58	0.66	0.29	0.3	0.22	0.24	0.25	0.54	0.19	0.22	0.29	0.34	0.44	0.6
Seasonal		0.49	0.53	0.61	0.60	0.65	0.74	0.46	0.45	0.48	0.47	0.53	0.56	0.44	0.46	0.48	0.49	0.5	0.55
Extreme		0.51	0.56	0.66	0.68	0.75	0.81	0.47	0.49	0.49	0.55	0.61	0.58	0.49	0.5	0.46	0.56	0.6	0.57
Static	**C**	0.27	0.31	0.40	0.50	0.58	0.66	0.26	0.26	0.35	0.25	0.27	0.35	0.26	0.28	0.36	0.38	0.47	0.49
Seasonal		0.51	0.56	0.66	0.68	0.75	0.81	0.47	0.44	0.45	0.54	0.53	0.58	0.41	0.47	0.48	0.47	0.48	0.54
Extreme		0.51	0.56	0.66	0.68	0.75	0.81	0.42	0.47	0.54	0.58	0.63	0.66	0.48	0.49	0.53	0.55	0.61	0.62

Table 7.3 Table showing mean cortisol levels of each agent across all three experiments, broken down by world condition (*Static, Seasonal, Extreme*) and Bond Combination (A, B, C). Values highlighted in grey indicate agents who were bonded in that condition.

rank (r = 0.89).

As the world challenges increased from *Static*, *Seasonal* and *Extreme*, the relative cortisol levels between the highest-ranking agent A1 and other agents decreased. Much like the findings from the investigation in Chapter 5, we observed that cortisol levels for subordinates (A5 and A6) remained relatively high across all experimental conditions.

7.4.2 Experiment 2 - Fixed Bonds

Viability Indicators

In this experiment, we introduced affective social bonds for three of our agents with fixed bond strength ($DSI = 2$), with the remaining three agents unbonded. For each of the three bond combinations, we saw increases in *Life Length* and *Physiological Wellbeing*, with non-significant changes to *Average Comfort* when compared to the control group, across each of the three world conditions. Tables 7.4 and 7.5 show the results of the viability indicators for the overall society as well as bonded and unbonded agents.

In *Static* world conditions, we observed *Life Length* increases between 59% and 72%, and a 27% increase in *Physiological Wellbeing* of the overall society. For the three bonded agents within this condition, *Life Length* performance increased by 58%–68% when compared to control, and *Physiological Wellbeing* increased between 27%–67%. For unbonded agents, we observed Life Length improvements of 40%–52%, and P*hysiological Wellbeing* improvements between 3%–29%. In sum, we observed viability benefits for both bonded and unbonded agents in these conditions compared to control.

Bond Condition	**World Condition**	**Life Length (Society)**	**vs. control**	**Life Length (Bonded)**	**vs. control**	**Life Length (Unbonded)**	**vs. control**
A	Static	53%	+72%	55%	+68%	50%	+52%
B	Static	49%	+59%	51%	+56%	46%	+40%
C	Static	50%	+61%	52%	+58%	47%	+42%
A	Seasonal	37%	+82%	39%	+81%	34%	+58%
B	Seasonal	34%	+68%	36%	+67%	31%	+44%
C	Seasonal	35%	+69%	37%	+69%	33%	+52%
A	Extreme	36%	+79%	38%	+77%	33%	+54%
B	Extreme	33%	+64%	35%	+63%	30%	+40%
C	Extreme	37%	+80%	39%	+78%	39%	+80%

Table 7.4 Table showing Life Length results from Experiment 2 (with fixed bonds) for the society overall, and broken down between bonded/unbonded agent groups. We highlight the largest changes versus control in grey.

In the first dynamic *Seasonal* condition, we found statistically-significant increases in *Life Length* (+68% to +82%, $p < 0.01$ for all conditions) and *Physiological Wellbeing* for the overall society when social bonds existed, compared to the control condition, with improvements between 14%–31% ($p < 0.01$). This increased performance was observed across all of the different bond combinations (Table 7.4).

For bonded agents, *Life Length* improvements were between 67%–81% compared to control, with unbonded agents benefitting from a mean *Life Length* increase of 51% vs. control. *Physiological Wellbeing* saw similar improvements: bonded groups saw improvements between 23%–48% ($p < 0.05$ for all conditions), with unbonded seeing a 5%–26% ($p < 0.05$) increase. Again, in these conditions, we note improvements to both Life Length and Physiological Wellbeing for both bonded and unbonded agents compared to control conditions.

In our second dynamic condition (*Extreme*), we again note increased performance for the entire society (Table 7.4). For bonded agents, *Life Length* saw improvements between 63%–78% ($p < 0.01$) vs. control, with unbonded benefitting with a 40%–80% ($p < 0.01$) increase. *Physiological Wellbeing* improvements ranged from 41%–84% ($p < 0.01$), and unbonded agents saw improvements of 12%–17% ($p < 0.01$).

In sum, regardless of bond combination, overall group viability increased across all world conditions when three of the agents were bonded through fixed bonds. In general, bonded agents outperformed unbonded

Bond Condition	**World Condition**	**Phys. Wellbeing (Society)**	**vs. control**	**Phys. Wellbeing (Bonded)**	**vs. control**	**Phys. Wellbeing (Unbonded)**	**vs. control**
A	Static	0.96	+27%	0.97	+27%	0.95	+29%
B	Static	0.95	+27%	0.96	+30%	0.94	+03%
C	Static	0.96	+27%	0.96	+67%	0.95	+03%
A	Seasonal	0.65	+31%	0.70	+36%	0.59	+26%
B	Seasonal	0.56	+14%	0.58	+23%	0.54	+05%
C	Seasonal	0.60	+22%	0.59	+48%	0.61	+05%
A	Extreme	0.64	+28%	0.77	+41%	0.55	+12%
B	Extreme	0.71	+38%	0.79	+63%	0.63	+17%
C	Extreme	0.73	+42%	0.72	+84%	0.74	+16%

Table 7.5 Table showing Physiological Wellbeing results from Experiment 2 (with fixed bonds) for the society overall, and broken down between bonded/unbonded agent groups. Largest differences versus control are highlighted in grey.

agents in terms of both viability indicators.

Social Behaviours

Table 7.6 shows the total amount of intra-bond *Grooming* and *Aggression* in each of the world conditions. For all bonded combinations, intra-bond *Grooming* increased as the relative world challenge increased. Agents in Bond Condition A (A1-A2-A6) exhibited the highest amount of intra-bond *Grooming* (63%–83%), whilst Bond Condition C (A4-A5-A6) exhibited the lowest amounts (17%–23%), suggesting the majority of *Grooming* for this group was performed with agents outside of their social bond group.

Intra-bond *Aggression* was highest in the *Seasonal* world condition for all bond combinations: account for 10% of aggression performed by agents in Bond Combination A, and 56% of all aggression performed by bonded agents in Bond Combination B.

Looking at the temporal distribution of *Aggression* performed by bonded agents in these conditions (Figure 7.6, top), we note that, despite increasingly difficult world conditions (in *Seasonal* and *Extreme* worlds), the absolute count of intra-bond *Aggression* remained fairly consistent throughout the experiments. We note that intra-bond *Aggression* increased for Bond Condition B in Phases 5 and 6 (with two and three food resources available, respectively), accounting for 56% of total aggression performed by those bonded agents.

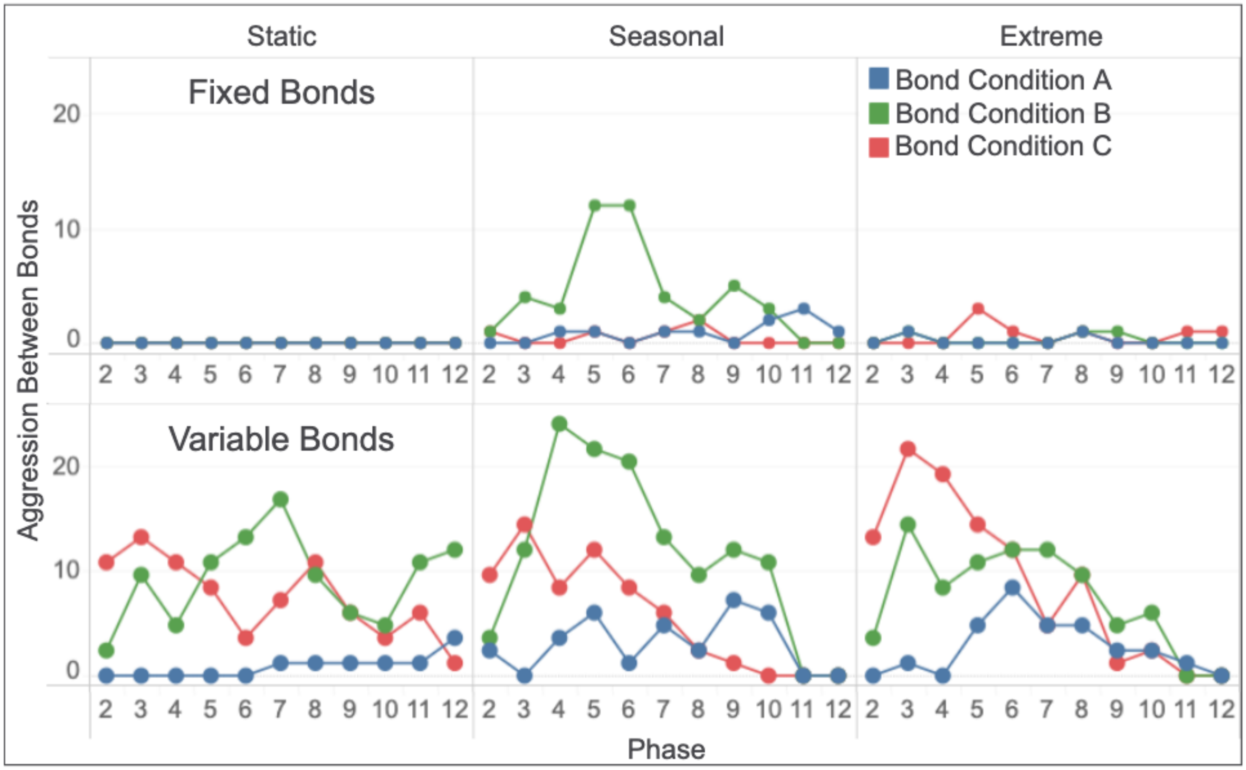

Figure 7.6 Total amount of intra-bond *Aggression* (y-axis) (*Aggression* performed between the three bonded agents) during each phase of food availability (x-axis), aggregated across all simulation runs.

Hormonal Levels

Across all bond and world conditions, we saw a reduction in mean cortisol levels for bonded and, in some conditions, unbonded agents (Table 7.3) vs. control.

In *Static* conditions, mean cortisol levels for bonded agents ranged between 0.28–0.37; 0.56–0.58 in the *Seasonal* environment, and 0.62–0.65 in the *Extreme* environments. We saw reduced cortisol levels in unbonded agents (*Static*: 0.36–0.42, *Seasonal*: 0.49–0.60, *Extreme*: 0.54–0.65). Therefore, the presence of affective bonds of fixed strength for three agents reduced the overall stress experienced by all agents in the society.

Subordinate agents experienced the greatest reduction in stress levels when bonded with other, higher-ranked agents. Specifically, in Bond Condition A (A1-A2-A6 bonded), cortisol levels of A6 were reduced by between 0.17–0.30 vs. control. Similar benefits were found for agents in other bond conditions (A5: -0.12 to -0.23 in Bond Condition B (A3-A4-A5); A6: -0.14 to -0.31 in Bond Condition C (A4-A5-A6)).

Figure 7.7 (top) shows the distribution of mean cortisol levels for all bonded agents. In the *Static* world

Bond Condition	**World Condition**	**Intra-Bond Grooming %**	**Bonded Grooming #**	**Intra-Bond Aggression %**	**Bonded Aggression #**
A	Static	63%	1,686	0%	121
A	Seasonal	73%	1,281	10%	125
A	Extreme	83%	1,438	2%	141
B	Static	42%	2,120	0%	80
B	Seasonal	45%	1,580	56%	84
B	Extreme	46%	1,072	8%	38
C	Static	17%	2,341	0%	0
C	Seasonal	24%	1,025	100%	5
C	Extreme	23%	1,249	100%	9

Table 7.6 Table showing the total amount of Intra-Bond Grooming and Aggression across each of the conditions in Experiment 2. Intra-Bond Grooming/Aggression values describes what percentage of total grooming/aggression by bonded agents was performed towards affective bond partners. Notable high and low values are highlighted in grey.

condition, mean cortisol levels of bonded agents remained fairly consistent for all bond combinations, with increases towards the final phases of experimental runs, likely due to limited food availability and death of bond partners.

In the *Seasonal* environment, we note an earlier increases in mean cortisol levels, and that this increase was related to the reduced food availability (Table 7.2). For Bond Condition A (A1-A2-A6, blue line), mean cortisol levels remained significantly lower than Bond Condition B and C, peaking during the most challenging world conditions (Phase 4, when only one food resource was available) and reducing again as the world conditions become easier (Phases 5 to 9). However, for Bond Conditions B and C, despite mean cortisol levels reducing as food conditions improve across phases (phases 5–7), they remain elevated for the remainder of the experimental runs.

Contrary to the trends from the *Seasonal* environment and our own expectations, mean cortisol levels in *Extreme* world conditions saw increases in periods of good food availability (Phase 3, 5, 7, 9), with reductions during periods of bad food availability (Phases 4, 6, 8, 10). This trend is most notable in Bond Conditions A and B. Bond Condition C also experienced similarly-high mean cortisol levels in the first period of bad food availability (Phase 3) and these levels remained chronically-elevated throughout the experimental run.

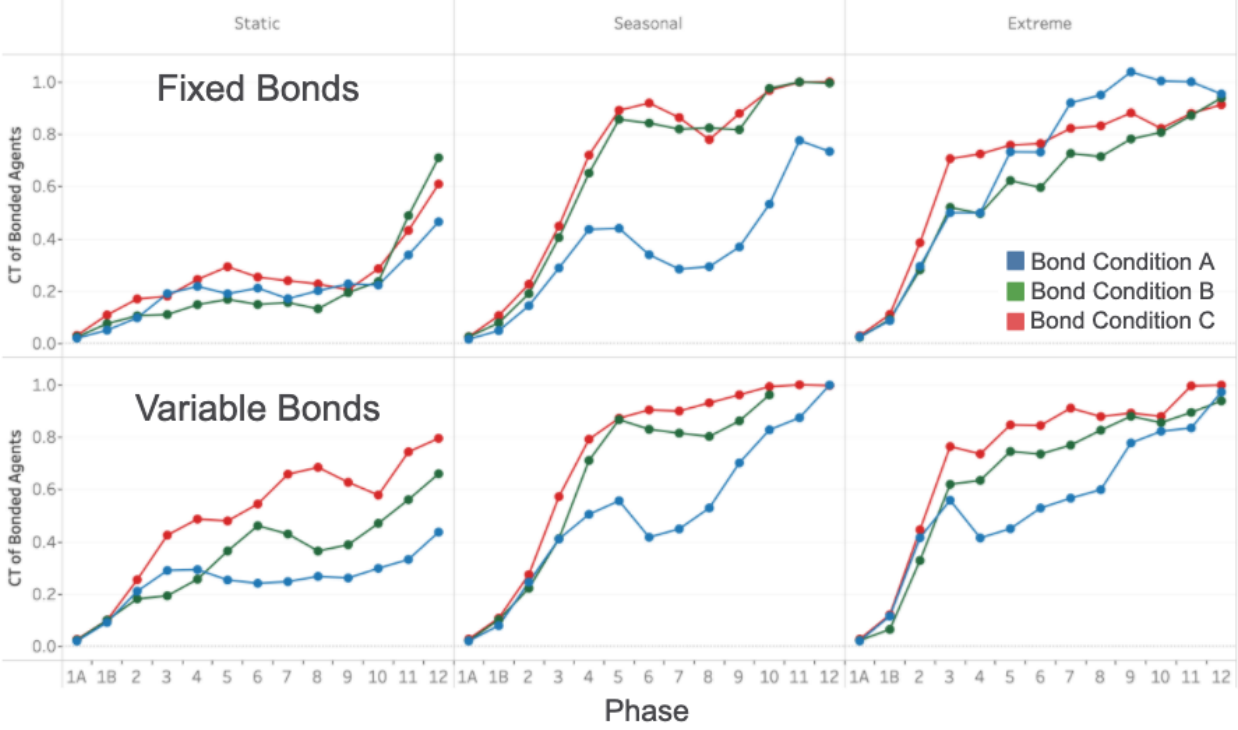

Figure 7.7 Mean cortisol (CT) levels for the three bonded agents across all world conditions. Top = Experiment 2, with fixed bonds. Bottom = Experiment 3, with variable bonds.

In terms of oxytocin levels, in conditions where food was abundant (*Static* conditions), oxytocin levels for bonded agents remained consistently high throughout the entire experimental run. Bonded agents in Bond Condition A (A1-A2-A6) experienced mean oxytocin level of between 0.85–0.93, Bond Condition B (A3-A4-A5) saw oxytocin levels between 0.82–0.93, and agents in Bond Condition C (A4-A5-A6) reported mean oxytocin levels between 0.83–0.92.

In sum, we saw lower mean levels of cortisol for both bonded and unbonded agents in these conditions when social support is present, compared to control. In the *Seasonal* conditions, cortisol levels loosely corresponding to the relative environmental challenges. In *Extreme* environments, the inverse was reported, with cortisol levels increasing during less-challenging periods. In this experiment, all bonded agents experienced similarly-high oxytocin levels throughout the experiments.

Bond Condition	**World Condition**	**Bond Type**	**Life Length (Society)**	**vs. control**	**Life Length (Bonded)**	**vs. control**	**Life Length (Unbonded)**	**vs. control**
A	Static	Variable	43%	+43%	46%	+42%	39%	+21%
B	Static	Variable	35%	+13%	37%	+13%	31%	+5%
C	Static	Variable	47%	+51%	46%	+40%	48%	+46%
A	Seasonal	Variable	31%	+50%	37%	+40%	27%	+2%
B	Seasonal	Variable	31%	+50%	33%	+25%	30%	+14%
C	Seasonal	Variable	30%	+48%	31%	+17%	30%	+14%
A	Extreme	Variable	37%	+79%	40%	+52%	32%	+22%
B	Extreme	Variable	31%	+54%	34%	+29%	28%	+5%
C	Extreme	Variable	31%	+52%	30%	+13%	32%	+22%

Table 7.7 Table showing Life Length results from Experiment 3 (with variable bonds) for the society overall and broken down between bonded/unbonded agent groups. Dark grey results highlight conditions where performance for unbonded agents was comparable, or better, than bonded agents.

7.4.3 Experiment 3 - Variable Bonds

Viability Indicators

In this experiment, the bond strength between individuals was initialised at its maximum value (2), experiencing a small decay rate at each time step, as well as being subjected to additional strengthening/weakening through social behaviours.

Despite not resulting in the same improvements of viability seen in Experiment 2, we nevertheless note increases in *Life Length* and smaller improvements in *Physiological Wellbeing* when compared to the control groups for all three bond conditions.

In *Static* conditions, we find viability benefits for the overall society when compared to control (Table 7.7). We observed *Life Length* improvements between 12%–51% compared to control. For all bonded conditions, this was a significantly lower result when compared to those of the fixed bonds (Table 7.4). We also observed small, yet statistically-significant increases in *Physiological Wellbeing* for the overall society (between 6%–15% improvement), with a small increase in *Average Comfort* (0.56%–3% increase, non-significant). Unlike conditions with fixed bonds, unbonded agents (LL: 48%) outlived bonded agents (46%) in Bond Condition C, with a significant difference in *Physiological Wellbeing* (Bonded: 0.76, Unbonded: 0.97, $p < 0.001$).

In both dynamic environments, we saw significant increases in both *Life Length* and *Physiological Wellbeing* for the overall society, with non-significant improvements in *Average Comfort* (Table 7.7). In the

Bond Condition	World Condition	Bond Type	Physiological Wellbeing (Society)	vs. control	Physiological Wellbeing (Bonded)	vs. control	Physiological Wellbeing (Unbonded)	vs. control
A	Static	Variable	0.78	6%	0.82	8%	0.77	4%
B	Static	Variable	0.80	9%	0.86	16%	0.78	2%
C	Static	Variable	0.79	15%	0.76	33%	0.97	4%
A	Seasonal	Variable	0.78	6%	0.57	9%	0.48	2%
B	Seasonal	Variable	0.78	9%	0.54	15%	0.53	2%
C	Seasonal	Variable	0.78	14%	0.50	24%	0.62	6%
A	Extreme	Variable	0.76	6%	0.61	11%	0.49	0%
B	Extreme	Variable	0.76	14%	0.60	22%	0.59	8%
C	Extreme	Variable	0.76	21%	0.53	37%	0.71	12%

Table 7.8 Table showing Physiological Wellbeing results from Experiment 3 (with variable bonds) for the society overall, and broken down between bonded/unbonded agent groups. Dark grey results highlight conditions where performance for unbonded agents was comparable, or better, than bonded agents.

Seasonal environment, *Life Length* improved by 50% in Bond Combinations A and B societies, and 48% in Bond Combination C societies, vs. control ($p < 0.01$). *Physiological Wellbeing* improvements ranged from 5.9%–13.6% ($p < 0.01$) for all groups vs. control.

Larger improvements in *Life Length* were noted in the *Extreme* world conditions, with *Life Length* improvements between 52% and 79% for the societies compared to control, and a 6.2% to 21% improvement in *Physiological Wellbeing*. Despite improvements in bonded agent viability, we see a number of conditions where unbonded agents showed greater viability in terms of *Physiological Wellbeing* (*Seasonal* and *Extreme* conditions for Bond Condition C).

In sum, we observe viability benefits in terms of *Life Length* and *Physiological Wellbeing* across all world and bond conditions, both for bonded and unbonded individuals. As with the results from Experiment 2, we did not observe any statistical differences in *Average Comfort*, and have not included them as part of the results.

Social Behaviours

Table 7.9 shows the distribution of intra-bond *Grooming* and *Aggression* interactions in experiments with variable bonds. In *Static* environments, despite the fact that overall amount of *Grooming* interactions was lower compared to Fixed Bond groups, the proportion of intra-bond *Grooming* was similar (Bond Condition A: 64%, B: 37%, C: 15%).

Bond Condition	World Condition	Intra-Bond Grooming %	Bonded Grooming #	Intra-Bond Aggression %	Bonded Aggression #
A	Static	64%	1,750	21%	66
A	Seasonal	60%	1,024	32%	91
A	Extreme	60%	473	35%	88
B	Static	37%	2,309	78%	120
B	Seasonal	42%	978	91%	116
B	Extreme	29%	730	78%	88
C	Static	15%	1,982	100%	75
C	Seasonal	17%	642	100%	83
C	Extreme	22%	708	100%	83

Table 7.9 Table showing the total amount of Intra-bond *Grooming* and *Aggression* across each of the conditions in Experiment 3. These values describe what percentage of total grooming/aggression performed by bonded agents was performed on affective bond partners. Notable differences to results of Experiment 2 (Table 7.6) are highlighted in grey.

However, we saw significant increases in intra-bond *Aggression* compared to Fixed bond groups. In Bond Condition A, *Aggression* towards bond partners accounted for 21%–35% of *Aggression* performed by these agents; 78%–92% for Bond Condition B, and 100% in Bond Condition C.

In contrast to the Fixed Bond groups, a significant amount of socio-negative behaviours performed by bonded agents was on their own bond partners. Looking at the temporal distribution of aggression behaviours (Figure 7.6, bottom), we identify significant increases in intra-bond *Aggression* across all world conditions, and for Bond Conditions B and C in particular. In the *Seasonal* world condition, we see that intra-bond *Aggression* for Bond Condition B increased in line with the environmental challenge (as food availability decreased until Phase 4), and then decreased as the environmental challenges were eased (Phases 5 to 8).

In the *Extreme* world condition, and for Bond Condition B and C in particular, we note that *Aggression* increases during periods of “good” food availability (Phase 3) and decreased during phases of bad food availability.

Hormonal Trends

We found similar levels of cortisol for both bonded and Unbonded agents in most conditions when compared to Fixed Bond groups (Table 7.3), with some notable differences.

In *Static* world conditions, bonded groups in Bond Conditions B and C reported increased mean cortisol levels when compared to unbonded agents (+0.09 and +0.19, respectively). Conversely, we observed lower mean cortisol levels for bonded agents in Bond Condition A in both *Seasonal* and *Extreme* environments, when compared to Fixed Bond groups (Figure 7.7, bottom). For *Extreme* world conditions, we saw earlier increases in mean cortisol levels for Bond Conditions B and C, which remained elevated throughout the experimental runs.

Despite similarities between cortisol levels, oxytocin levels for bonded agents were significantly lower (50% and 59%, $p < 0.001$) compared to Fixed Bond groups (Experiment 2). For individual agents with fixed bonds (Experiment 2), we found a moderate correlation between oxytocin levels and *Life Length*: ($r = 0.51$, $r = 0.70$, $r = 0.43$, $p < 0.001$ for all tests). In Variable Bond groups (Experiment 3), we found stronger correlations, with r-values of 0.55 ($p = 0.032$), 0.61 ($p = 0.041$), and 0.67 ($p < 0.001$) for Bond Conditions A, B, and C respectively. This suggests a potential relationship between bond quality and oxytocin levels, with a potential downstream effect on the reduction of cortisol and agent viability.

Overall, even though affective bond strength was of variable strength, we noted overall lower levels of cortisol for bonded agents when compared to control groups. This effect was seen regardless of the type of bond (Fixed or Variable), suggesting a stress-reducing effect when social support is available.

7.5 Discussion

Across both experiments with the presence of bond partners (Experiment 2 & 3), we observed significant improvements to the viability of bonded and, in some cases, unbonded agents. We observed lower levels of mean cortisol levels in conditions when agents had affective social bonds with others. In our model, we find that one of the “buffering” effects of social support on mediating the stress response resulted in significant advantages to agent viability. We propose that this type of behavioural and physiological adaptation, through interactions with the social environment, constitutes a type of “social allostasis”.

Compared to conditions where no bonds are present, the results of both experiments showed viability benefits on two of three measures (*Life Length* and *Physiological Wellbeing*) with no detrimental impact on the third (Average Comfort). In line with what we have also observed in previous chapters (Chapters 5 and 6) we found viability benefits for unbonded agents as well as bonded agents in a number of conditions. This suggests that affective relationships may have survival-related benefits for the wider society, and not just for the agents with affective bonds, suggesting a “contagion” of viability facilitated by affective social bonds.

We highlight several key findings. Firstly, despite the fact that only three of the agents in this model having social support available, their intra-bond behaviours and interactions still had benefits on the life

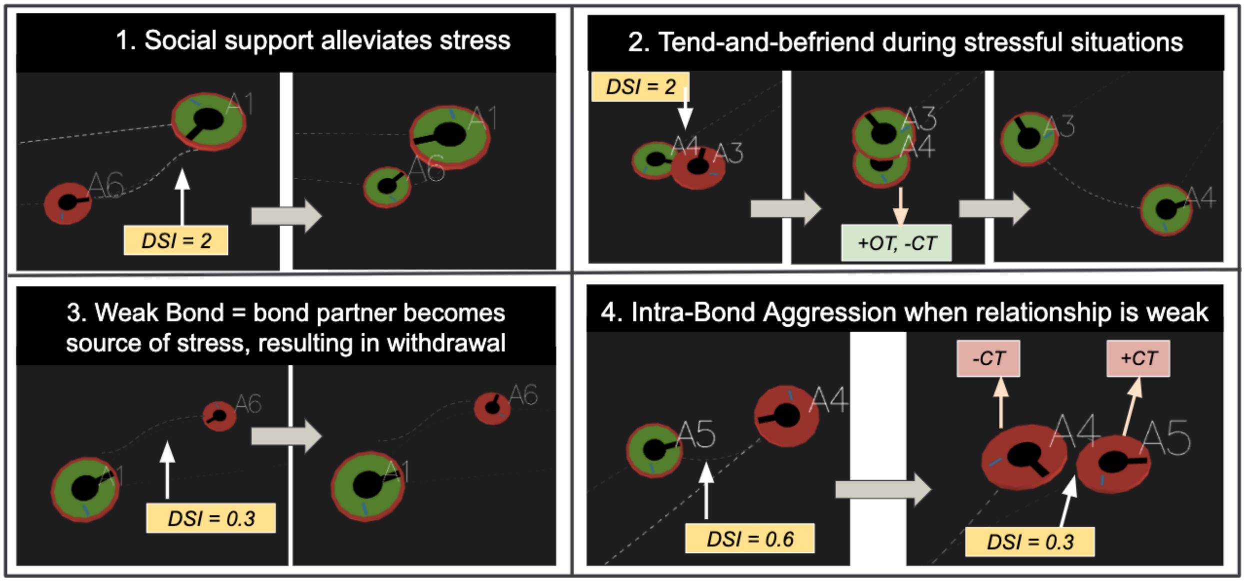

Figure 7.8 Screenshots of four different emergent social scenarios, which we discuss in detail in this section. Green agents indicate agents who are not stressed, red agents indicates agents who are stressed. Agents are facing in the direction of the black line. Scenario 1: Agent A6 is stressed, and seeing a bond partner (A1) with a strong affective bond reduces cortisol levels and alleviates A6's stressed state. Scenario 2: When bonds are strong or fixed, A4 approaches its stressed bond partner A3 to *Groom*, increasing A4's oxytocin (OT) and reducing cortisol (CT), alleviating stress. Scenario 3: A weakened affective bond between A1 and A6 results in A1 no longer providing the "social buffering" effect, instead becoming a source of stress resulting in A6 avoiding it. Scenario 4: When affective bonds are weak, bond partners are targeted to perform *Aggression* on, further reducing the bond strength (DSI) between them, and increasing stress levels (CT) on recipient.

quality and quantity of the wider society. Secondly, we identified behaviours related to the "tend-and-befriend" [30] hypothesis in some conditions of Fixed Bond types, whereby individuals seek out positive social contact during stressful conditions, rather than a fight-or-flight [21] approach. Thirdly, contrary to our prediction that variable relationships may be strengthened during periods of food shortages, we found instead that these bonds were often irreversibly damaged during challenging environmental conditions, as a function of both social and environmental stressors. Furthermore, in some conditions, intra-bond *Aggression* increased during periods of food abundance, not food scarcity. We also highlight how these affective relationships, if not extremely strong, can be impacted by the wider social context.

Below, we discuss some of our observations in detail. We have also included screenshots of behaviours that we have discussed in this section.

7.5.1 "Social Buffering" Effect for Some Individuals Provides Wider Social Benefits

From our results from Experiment 2 (Fixed Bonds) and in some conditions in Experiment 3 (Variable Bonds), we saw significant viability benefits in terms of *Life Length* and *Physiological Wellbeing* for the three unbonded agents as well as bonded agents who had social support available. These viability benefits were seen across all world conditions.

We considered that these stress "buffering" effects of social support for some agents—via its effects on mediating cortisol release and providing additional social outlets—may present society-wide survival-related benefits. In addition to improved viability, overall cortisol levels were lower for societies in Experiment 2 (Fixed Bonds, Table 7.3). Reduced stress for some agents reduced the stress of unbonded individuals: reducing their likelihood to perform *Aggression* behaviours on others, resulting in survival-related advantages for all. This effect was seen both down (Bond Condition A) and up (Bond Condition C) the social hierarchy.

As with our findings from Chapter 6, this may highlight the contagious nature of affective states (stress) in a small society: the calming effect provided by social support may therefore be felt by the wider society. These emergent social dynamics may give insight into the nature of real-life social structures (including humans).

For autonomous agents, these results suggest that it may not be the case that all agents require the same types of adaptive mechanisms in order to experience adaptation-related benefits. Taking these results with the results of previous chapters (Chapter 5 and 6), we hypothesise that (small) societies of socially-affective agents can benefit from a type of proxy-adaptation when other individuals in their society are endowed with mechanisms of adaptation. Such findings may have implications on future development of adaptive models.

7.5.2 "Tend-and-Befriend", Not "Fight-or-Flight", when Affective Bonds are Fixed

In the experiment where affective bond strength remained fixed throughout an agent's life (Experiment 2), we noted that for all bond combinations, intra-bond *Grooming* increased with the world challenge, from *Static*, *Seasonal* and *Extreme* worlds respectively (Table 7.6).

In the dynamic *Seasonal* and *Extreme* environments, rather than become stressed and perform *Aggression* towards others during stressful situations, we saw that bonded agents would seek out positive social contact with their own affective bonds: performing *Grooming*, releasing oxytocin and strengthening the bond between them. Figure 7.8 (top right) shows screenshots of this behaviour, and Figure 7.9 illustrates the internal

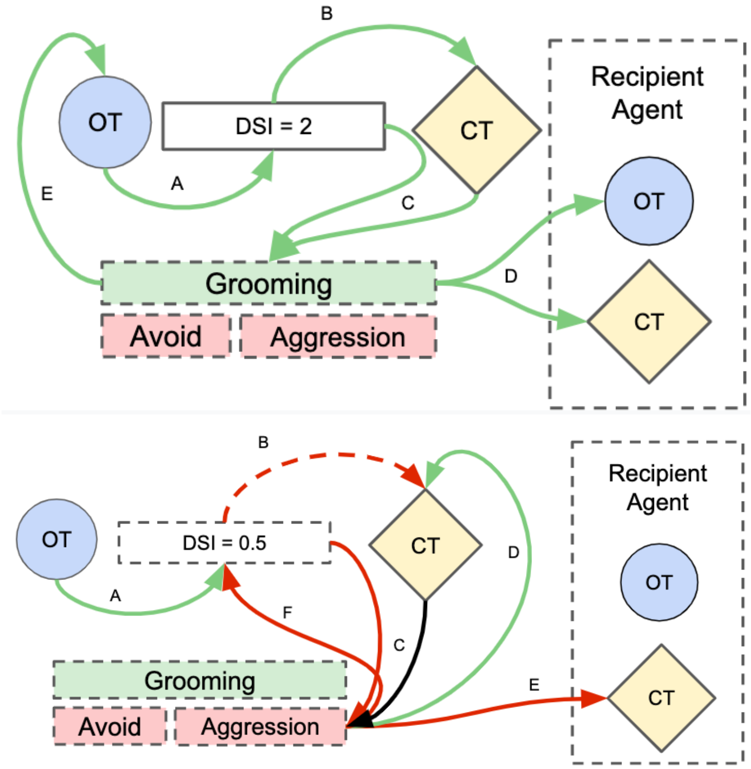

Figure 7.9 Illustration of two contrasting mechanisms resulting in the “tend-and-befriend” and “fight-or-flight” behaviours observed in Experiment 3. Top: Hormonal mechanisms during the “tend-and-befriend” behaviours (Figure 7.8, (2)). A: Oxytocin (OT) modulates preference toward partners with strong bonds (Equation 5.7). B: These bond partners reduce the stress (CT) response, reducing the likelihood of agents becoming stressed. C: Non-stressed agents and strong bond partners results in *Grooming* being executed to satisfy *SocialNeed*. D/E: *Grooming* increases OT in both actor and recipient agents, and also reduces CT in recipient. Bottom Hormonal mechanisms during the “fight-or-flight” behaviours (Figure 7.8, (4)).: A: Oxytocin modulates preference towards bond partner (Equation 5.7). B: A weaker bond reduces in a reduced buffering of the stress response, increasing the likelihood of agents becoming stressed. C: Stressed agents and weaker bonds results in *Aggression* (or *Avoid*) towards bond partners. D: *Aggression* reduces stress (CT) in actor agent. E: *Aggression* increases CT levels in recipient agent. F: *Aggression* reduces bond strength between agents. Note how in the bottom illustration, further secretion of OT does not occur.

mechanisms of these interactions.

This reinforcement of prosocial interaction between agents reduced the stress of both agents and contributed to their survival in future, more challenging situations. We considered this an important behavioural adaptation that facilitated their viability in difficult conditions.

Contrary to the “fight-or-flight” hypothesis of stressful situations [21]—which we had also observed in the experimental condition with variable bonds—we suggest that this type of coping behaviour during stressful situations may have similarities with the “tend-and-befriend” model [30]. This hypothesis proposes that, rather than the classic fight/flight response where individuals use socio-negative behaviours as coping strategies, social individuals—including humans—actively seek positive affiliative contact during times of stress. Furthermore, this coping strategy, underpinned by oxytocin, is an essential survival-related adaptation of behaviour that has been considered a maternal or biologically female-oriented response during stressful conditions.

This increase in intra-bond *Grooming* rates was significantly higher in Bond Condition A compared to the other bond conditions. As a reminder, we tested this bond combination as an abstraction of a parent-offspring relationship. Considering oxytocin’s implications on maternal and tending behaviours [212], [213] and that the “tend-and-befriend” model has been proposed as a behavioural adaptation typically found in females (with fight-or-flight being a more male-oriented response [30]), our results may give insight into some of the potential gender-related differences of oxytocin’s effects on behavioural adaptation in stressful situations in natural societies. However, further work is required to assess this further and we consider this as an area for future investigation.

In contrast, we had seen the fight-or-flight response in conditions with Variable Bonds (Figure 7.8, bottom right). This response, which we also saw was associated with lower oxytocin levels and increased cortisol levels 6.3, resulted in bonded agents performing higher rates of intra-bond *Aggression* in physically-challenging conditions.

Taking both observations and quantitative results, we consider the importance of strong bonds—such as those typically seen in parent-offspring relationships—and its implications on the different coping behaviours used during stressful situations. These considerations have also been raised by [179], who highlight the importance of bond stability as well as the degree of affiliation on the efficacy of social buffering effects.

7.5.3 Bond Reconciliation is Hindered in Challenging Environments

Given the social buffering hypothesis, we had predicted that bond partners who had strong affective relationships prior to periods of physical challenges would tolerate the physically-challenging conditions with a reduced stress response—as lower cortisol levels would allow these agents to stay “calmer” during stressful situations—and then seek out social support again during periods of good food availability.

While we found support for this hypothesis when bond strength remained fixed (Experiment 2), this was not the case if bonds were weakened or if oxytocin levels were lower for bonded individuals (Experiment 3). In these situations, the stress presented by challenging environmental conditions detrimentally affected the

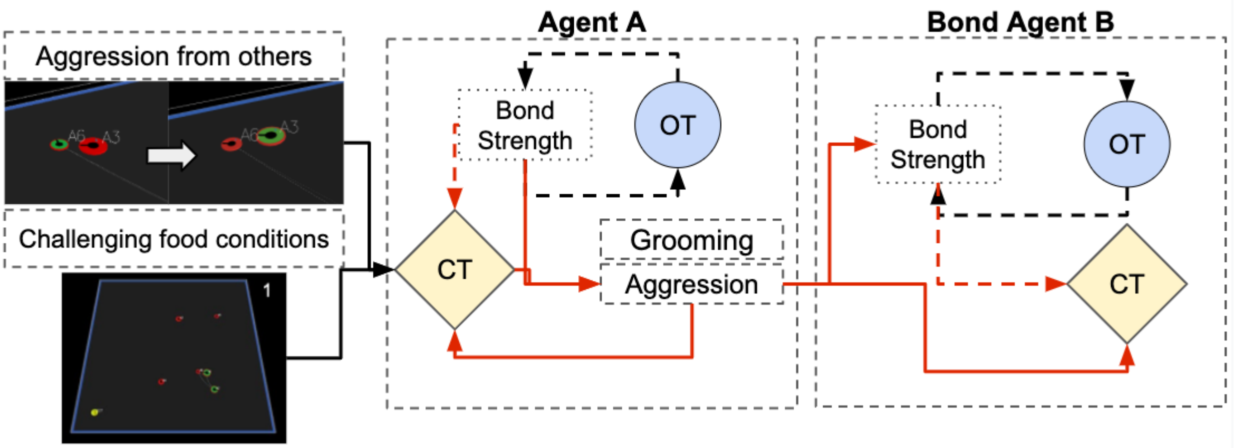

Figure 7.10 Illustration of how external social and physical contexts affect bond strength between bond partners. Receiving *Aggression* from higher-ranked individuals or being in a challenging world condition with limited food availability raises cortisol levels in agent A. A weaker social bond (as is sometimes the case in the Variable Bonds group) along with elevated CT levels results in *Aggression* towards bonded agents (B). Performing *Aggression* on bond partners reduces CT levels in agent A, increases CT levels in bonded agent B, and also "damages" bond strength between bond partners A and B. Increased CT levels and weakened bonds results in future behaviours skewing towards socio-negative interactions.

quality of social bonds to varying degrees–affective relationships were not strong enough to provide an appropriate stress-reducing effect. As a result, intra-bond aggression increased (Table 7.9) which reduced bond quality. The ongoing challenges from the *Seasonal* and *Extreme* environment then hindered the opportunity for any *Grooming* between bond partners; eventually resulting in a loss of social support. These mechanisms are further illustrated in Figure 7.10.

We found that, despite social bond partners offering some "buffering" effect if bonds were extremely strong (Experiment 2); weaker, unstable bonds were susceptible to being lost beyond reconciliation due to significantly challenging conditions. Again, this highlights how there may not be a universal effect of social buffering, but that such effects may be impacted by both environmental and affective contexts.

7.5.4 Variable Affective Relationships can be Broken by "Outsiders"

We observed that the quality of variable affective relationships was affected, not only by the interactions between the bonded agents, but the interaction on these bonds from other individuals in the society. Particularly if an affective bond was already "weak" (DSI ≤ 1), it had the potential to be irreversibly damaged by outside social influences.

When unbonded, higher-ranked individuals (such as A2 or A3) would become stressed in physically-challenging conditions, *Aggression* was directed towards lower-ranking agents (i.e., A4): increasing the stress level of the latter. For these lower-ranked agents, if their current bonds were not strong enough, they would often become aggressive towards their own (lower-ranking) bond partners. Figure 7.10 shows how these mechanisms emerge in both agents.

Rather than "tend-and-befriend" behaviour that we described earlier (see also: Figure 7.8, top right), these individuals would instead perform *Aggression* towards their own bond partners (Figure 7.8, bottom right), reducing their affective bond strength. This loss of affective relationship quality, then, had an impact on the availability of future coping mechanisms and bond partner viability.

In these situations, when bonds were weak or if agents were experiencing higher levels of stress, socio-negative behaviours from higher-ranked members of society had a "trickle-down" effect on their interactions with bond partners. This loss of affective relationship results in bond partners withdrawing from interactions in the future (Figure 7.8, bottom left), and highlights how an affective bond partner can go from being a source of social support to one of social stress. This behaviour highlights the importance of the wider social context on the quality and maintenance of affective relationships. While affective relationships may reduce stress in some conditions, their stress-reducing efficacy may be challenged by significant environmental and social difficulties. However, this observation may also be a result of the small number of agents in the society (6), or the limited number of agents bonded together (3 in each condition). Future work should consider larger societies or different group sizes of bonded agents to further understand whether this effect emerges in those conditions. For example, increasing the number of bonded agents (from 3 to 4, for instance) in this same society of six agents would both increase the potential availability of social support and reduce social stress for bonded agents, and we speculate that this would result in lower rates of intra-bond *Aggression* between bonded agents.

7.5.5 Aggression Rates Increase During Periods of Food Abundance

In the *Extreme* world conditions, and in particular for Bond Conditions B and C, we observed an increase in aggression rates when food availability was "good" (Phase 3, 5, 7, 9, etc).

Despite our prediction that grooming would increase during periods of reduced environmental stress, we instead saw that the magnitude of changes associated with the *Extreme* world condition did not allow the agents to adapt to the constantly-changing conditions fast enough. As a result, the sudden shortage of food would increase group stress, which then carried over into the next "season" of food abundance, increasing *Aggression* during these periods—including intra-bond *Aggression*.

This is not seen in conditions where the physical environment changes more gradually (*Seasonal*). In those conditions, we see *Aggression* increase and decrease roughly in line with the food availability.

In terms of the stress-reducing effects associated with social support in our model, we conclude that these effects were not one-size-fits-all, but rather dependent on the magnitude of environmental changes and the quality of affective bonds.

We propose that this may be a limitation of the effects investigated in the current model, and that such behaviours may be offset in real-life individuals who benefit from the more anticipatory, stress tolerance-related benefits of social support [16], [207], [208].

Conversely, this observation may shed light on emergent social dynamics when individuals do not, or cannot, experience the complete (physiological) benefits of stress-reduction. By lacking the ability to adapt and respond appropriately to future stressful situations, we consider that societies may quickly become dysfunctional, increasing their socio-negative behaviours during periods of stress, where it would be more beneficial for them to cooperate. We suggest that our future work should investigate the specific temporal dynamics between food availability, rates of aggression and resultant viability to understand this further.

7.6 Summary

Building on the strictly-rank-based social interactions from Chapter 6, we have investigated the role of "social buffering"—which proposes that social support mediates the stress response—and studied how different types of affect-based social bonds affect social interactions and wellbeing of individuals with social support. Building on findings from biology and our previous observations from Chapter 5, we had hypothesised that the presence of oxytocin-mediated social bond partners would benefit the viability of bonded agents—and particularly the viability of subordinate agents with affective bonds—across a variety of different physical conditions, through a reduction in overall stress and socio-negative interactions.

To investigate our hypothesis, we accounted for one hypothesised "buffering" effect of social support—that it mediates stress hormone release—and included this abstracted effect in our agent model.

We performed our investigation using our society of six agents across three different bond combinations, related to an individual's social rank, and two different types of bonds ("fixed" and "variable") related to the strength of the affective relationship between individuals. "Fixed" bonds remained consistently strong throughout the entirety of an individuals' life, whereas "variable" bonds were subject to strengthening and loss over time. We studied these bonds across three different environmental conditions, including two dynamic environments in order to introduce environmental uncertainty to our investigations.

Overall, we found support for the hypothesis that the stress-reducing effects of social support facilitates wellbeing during stressful situations through physiological and behavioural adaptation. Furthermore, we had also seen how the viability of non-bonded agents improved when other agents in the society had social support available to them. These findings may have implications on the development of future socially-adaptive agent models. Specifically, it may be that not all social agents in a society need to be endowed with the same (level of) adaptive mechanisms in order to benefit from their adaptive process. For physical agents in particular, this may have some considerable benefits in terms of time and resource (hardware) requirements.

Our findings suggest that, for real-world societies, the efficacy of social support on stress-reduction may not be universal, but that the effects may depend upon the relative environmental and social challenges as well as the quality of affective relationships. We found that "variable bonds" were susceptible to being damaged beyond reconciliation if the environmental challenge was significant enough and relationships were not strong. We also observed how these variable bonds may also be affected by external social influences.

For stronger relationships, we found how behaviours related to the "tend-and-befriend" approach promoted wellbeing during stressful situations over the more traditional fight-or-flight response, and found how these behaviours were also related to oxytocin levels. We suggest that this may be indicative of differences in stress coping behaviours between different biological sexes in the real world. We highlight that our model, with its abstractions of oxytocin grounded in biological literature, may be used as a tool to investigate this in more detail.

Contrary to our prediction that socio-negative behaviours would be reduced during periods of good food availability, we noted that this was not the case when the environment drastically changed. The emergence of increased aggression during relatively safe environments made it significantly more challenging to respond to future stressful situations. From a survival perspective, this appears to be counter-productive. However, we note that we have only observed one aspect of social buffering in this chapter. Only considering this aspect may limit the potential for adaptation when the environment undergoes such an extreme change.

Mediating cortisol release is one way in which social support has been suggested to regulate an individual's stress response. As [16], [207], [208] have stated, social support also plays an important role in the perception and response to future stressful situations. For individuals with such support, perceiving future stressful situations as less threatening, minimising the fear-based response, and being able (or willing) to tolerate more stress when social support is available plays a significant role in their ability to adapt to change and maintain long-term survival.

This adaptation of stress resilience—potentially driven by oxytocin's effects on the stress response system—also has implications on stress-related coping mechanisms [214], [215], which may influence overall social dynamics.

In the next chapter, we continue our investigation into the stress-reducing effects of social support and focus our attentions to this second stress-reducing effect of social buffering. Specifically, we focus on the role of oxytocin-mediated, affect-based relationships on the adaptation of an individual's stress tolerance, and investigate its effects on group wellbeing and social interactions.

Chapter Eight

Investigating the Adaptive Effects of Social Support and Oxytocin on Stress Tolerance

8.1 Introduction

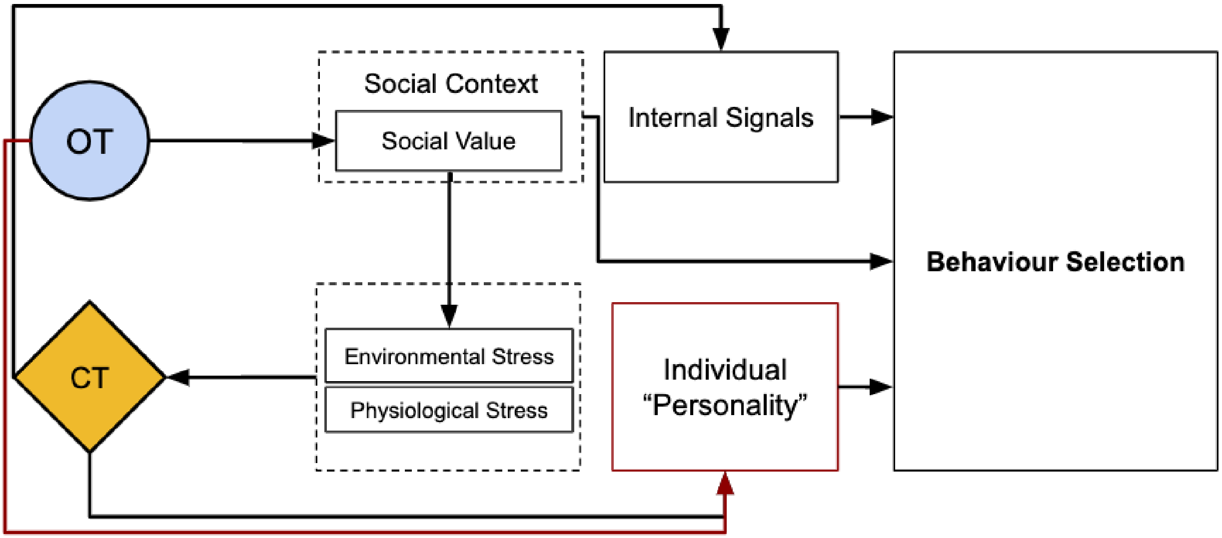

Figure 8.1 A high-level view of the effects investigated in this chapter. Red lines indicate extensions from the previous chapter. In this investigation, oxytocin is used to provide a buffer to individual stress tolerances (considered a "personality trait"). The rest of the model remains the same as Chapter 7. A more granular view of this model can be seen in Figure 9.3.

In Chapter 7, we investigated the role of one of the hypothesised effects of the "social buffering" hypothesis—that social support mediates the release of stress hormones—for individuals with social support—on the survival and social interactions of a small society of social agents. Investigating this effect using different types of affective bonds, we found that the presence of social support resulted in a reduction of overall stress levels and therefore provided advantages to the survival and wellbeing for bonded and, often, unbonded agents.

However, we saw that for some bond types, challenging environmental conditions posed significant detriments to the efficacy of social support, as well as the overall strength of the affective bond. We hypothesised that this was due to our model only accounting for one hypothesised effect of stress-reduction through social support. As we saw in Chapter 3—which contradicted some earlier findings (Chapter 2)—accounting for additional hypothesised effects in our model may be an important aspect of our investigation, in order to understand and contextualise some of the diverse observations of these mechanisms in biological societies.

In addition to attenuating the release of cortisol, social support's hypothesised role on stress reduction is also hypothesised to occur via an adaptive effect on stress tolerance. This adaptive regulation of stress tolerance (or stress resilience) has been found to be mediated through oxytocin mechanisms, which is secreted as a function of social interaction. More directly, oxytocin has been considered to regulate the activation of the stress response of the autonomic nervous system, which drives a number of physiological and behavioural responses.

While the presence of social support has been found to ameliorate the release of stress hormones during stressful or challenging (physical and social) situations, oxytocin may play an additional complementary role on modulating the internal tolerance of this stress system—which can also be considered a type of affective adaptation of the stress response.

However, the precise relationship between social support, the oxytocinergic system, and the internal stress tolerance ("resilience") of an individual is still unclear. Other systems, such as the reward-based and opioid systems, have also been considered to play a role in oxytocin's effects on stress-buffering and affect-based social interactions. While these precise, intermediary mechanisms are still unknown, we hypothesise that oxytocin's direct effects on the regulation of the internal "stress tolerance" system can provide stress-reducing effects, which can also promote the wellbeing of individuals with affect-based social support.

In this chapter, we account for this hypothesised effect of oxytocin-mediated social support on the regulation of the stress response. We investigate how adaptation of an internal stress tolerance through oxytocin—released as a function of social interaction—affects the overall social interactions and group viability of agents in a small society. We also test whether also accounting for this second effect of social support on stress-reduction is more beneficial to the long-term viability of these social agents, compared to the single stress-reducing effect we had accounted for in Chapter 7.

Our results found that accounting for this additional effect of social support resulted in significant survival-related benefits, compared to when only one effect was considered in Chapter 7. These survival-related benefits were seen for individuals both with and, in some conditions, without affective social bonds, mirroring our previous findings. We also found reduced rates of socio-negative behaviours in our societies compared to conditions where this adaptation was not present, and considered the implications of stress regulation on affect-based behavioural contagion.

In addition, we find that social support provided an anticipatory, adaptive regulation of the stress response, through a combination of hormonal and behavioural responses, which we propose played a key role in the long-term survival of individuals. This positive feedback loop of stress-reduction and behavioural adaptation suggested that there may be numerous effects of social support on long-term stress reduction, which may then have consequences on long-term physiological health for social individuals.

In terms of the development of our social allostasis-inspired model of adaptation, we find that accounting for two stress-reducing effects proposes significant benefits to the viability of our agent model, and we will retain these affect-based mechanisms in future iterations.

8.1.1 Social Support, Oxytocin, and Adaptation of Stress Tolerance

As investigated in Chapter 7, social support can provide stress-reducing effects during stressful events by providing immediate outlets for prosocial interaction [179] and by reducing the perceived threat of a stressful situation [208] resulting in reduced levels of stress (cortisol) for socially-supported individuals.

In addition to these effects during stressful conditions, it has also been hypothesised that social support can provide an adaptive "buffering" of the stress system (the autonomic nervous system) [208], [216]–[219]: responsible for (often involuntary) physical, physiological and behavioural responses to stress (such as fight-or-flight responses occurring through activation of the sympathetic nervous system [21]).

Adaptation of this internal stress system has been found to occur through the presence of social support [204], [219], as well as after both receiving and giving social support [220], [221], and has been found even when social support is not immediately available [16], [206]. This adaptation is potentially linked to stress "resilience" [222], which may play a role in adapting individual coping strategies during stressful situations.

Furthermore, regular interactions with social partners may provide individuals with a "sense of predictability" and control of their social environment [223], including predicted or expected support [224].

This perception of control may help regulate their stress response during future stressful events [16], [207], even in absence of immediate social support, through adaptation of stress tolerance [219].

In sum, social support has been seen to provide an internal "buffer"—either psychological, physiological, or both—on stress tolerance (which can be viewed in the context of a "personality trait", investigated in Chapter 5): reducing activity of the sympathetic nervous system, which may play a role on adapting stress-related coping strategies [225] and physiological health [203], [204], [206].

Underpinning this effect is oxytocin, which has been found to play a role in this attenuation of the stress response by suppressing the activity of the sympathetic nervous system [226]. It also been seen to play a role in the modification of stress-related behaviours as well as coping mechanisms [214], [215], and mediate the effects associated with separation distress [227].

Taken together, oxytocin's anxiolytic effects, released as a function of social support, appear to play a key role in the adaptation of stress tolerance.

Though there appears to be a well-established relationship between oxytocin, social support, and long-term health and wellbeing, the underlying causal mechanisms still remain unclear. Oxytocin systems promoting partner preference have also been to interact with the reward and opioid systems [228], as well as reinforcement learning of social interactions [229].

Whether oxytocin has a direct physiological effect on the regulation of the stress system [209], [210], [218] and health, whether these effects are strictly-psychological [222], or whether these stress-reducing effects are a consequence of its promotion of social interactions has been difficult to distinguish in biological systems. More likely, the stress-reducing effects of oxytocin-mediated social support may be multi-faceted: by both suppressing stress system activation, and facilitating a positive feedback loop that promotes stress-reducing prosociality [191], [230].

In terms of a model of adaptation through social allostasis [14], this anticipatory adaptation of the stress response may play a significant role in determining how individuals adapt their physiological and behavioural responses to stressful events, and underpin the long-term wellbeing of these social individuals. We take a high-level approach in our model, and include an abstraction of oxytocin's hypothesised effects on regulating the autonomic nervous system, by adapting the individual tolerances to stress through oxytocin mechanisms in our model.

8.1.2 Aims & Hypotheses

Continuing the investigation into the "social buffering" hypothesis, we aim to investigate how the adaptation of the stress response via oxytocin—released through social interactions with bond partners—affects the viability and social interactions of a small society of agents, across environmentally-challenging, dynamic conditions.

We test these effects across three different bond partner combinations related to social rank. The strength of these bonds are dynamic and subject to change through social interactions (i.e. the "Variable Bond" type from Chapter 7). To test the adaptive effects of these stress-regulating mechanisms, we conducted our experiments across three different physical environments: one static and two dynamic, with respect to food availability.

We hypothesise that the adaptation of an individual's stress tolerance through oxytocin mechanisms mediated by social support will provide viability-related performance benefits to individuals with affect-based social support. We predict that this will be driven through a reduction of socio-negative interactions compared to conditions where this stress-reducing adaptation is not accounted for.

Based on our findings from Chapters 6 and 7, we further hypothesise that this improved viability performance of socially-bonded agents will provide a wider effect on group viability, resulting in improved performance for agents without affect-based social support.

In terms of developing our model of social allostasis, we also aimed to investigate whether accounting for two of the stress-reducing effects of social support provides advantages to our model performance, compared to conditions when a single effect was included in Chapter 7.

8.2 Extension to the Agent Model

To investigate the effects of social support on the adaptation of stress tolerance discussed in Section 8.1.1, we make one addition to our agent model from previous investigations. Here, internal levels of oxytocin, secreted as a function of social interaction with bond partners, is used to modulate and adapt the default tolerance to stress (the *StressThreshold* value), which we consider an abstraction of the autonomic nervous system and a "personality trait". The term "tolerance" describes the level of cortisol required for an agent to become "stressed". We retain all previous effects of the model used in Chapter 7 (Figure 8.1).

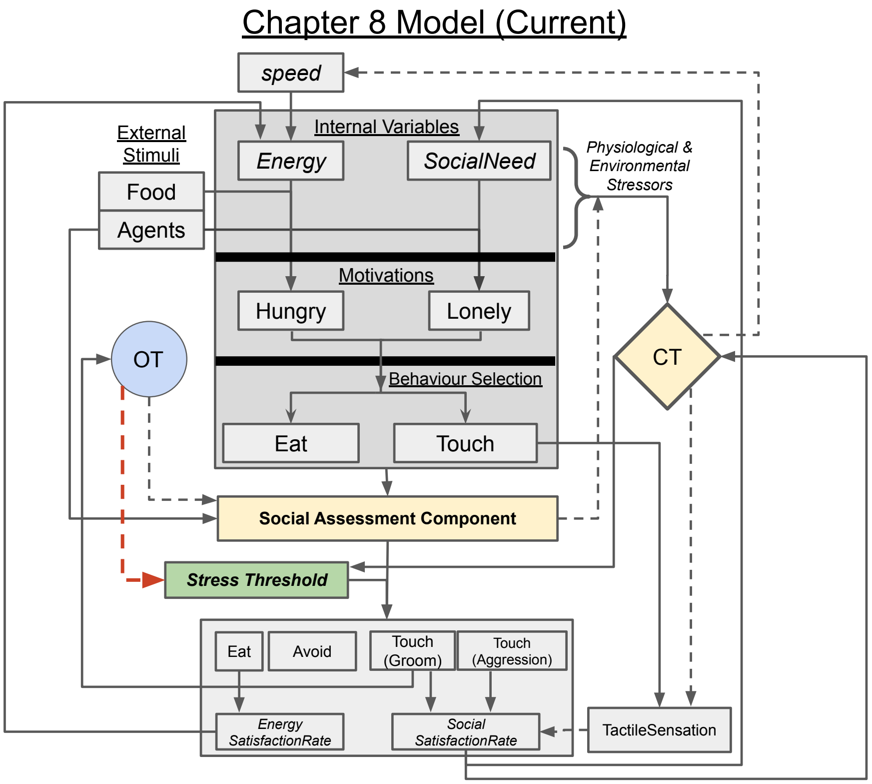

Figure 8.2 Illustration of the Action-Selection Architecture model used in this investigation. Extensions to the previous model are denoted in red. In this model, oxytocin (OT) adapts the *Stress Threshold* value of agents with affective bonds (Section 8.2.1), which is initialised at a default value (0.5) for all agents. This is an abstraction of the stress-regulating effects of social support discussed in Section 8.1.1).

8.2.1 Adaptation of Stress Tolerance via Oxytocin

In this model, an agent's *StressThreshold* value is adapted through its internal oxytocin levels. To reiterate, this *StressThreshold* value is an abstraction of "sympathetic nervous system"-type activation [231] and a "personality trait" that determines how much of the stress hormone (cortisol) an agent can withstand before it undergoes a "stressed" state.

We had investigated several different values of this *StressThreshold*: using both uniform and rank-related values (Chapter 6), with a uniform "neutral" value of 0.5 used in Chapter 7. In those investigations, *StressThreshold* values were fixed throughout all experiments.

In the present model, oxytocin—released as a function of positive social support and interaction—dynamically adapts this *StressThreshold* value over time. We consider this an abstraction of oxytocin's effect on the internal stress response system discussed in Section 8.1.1.

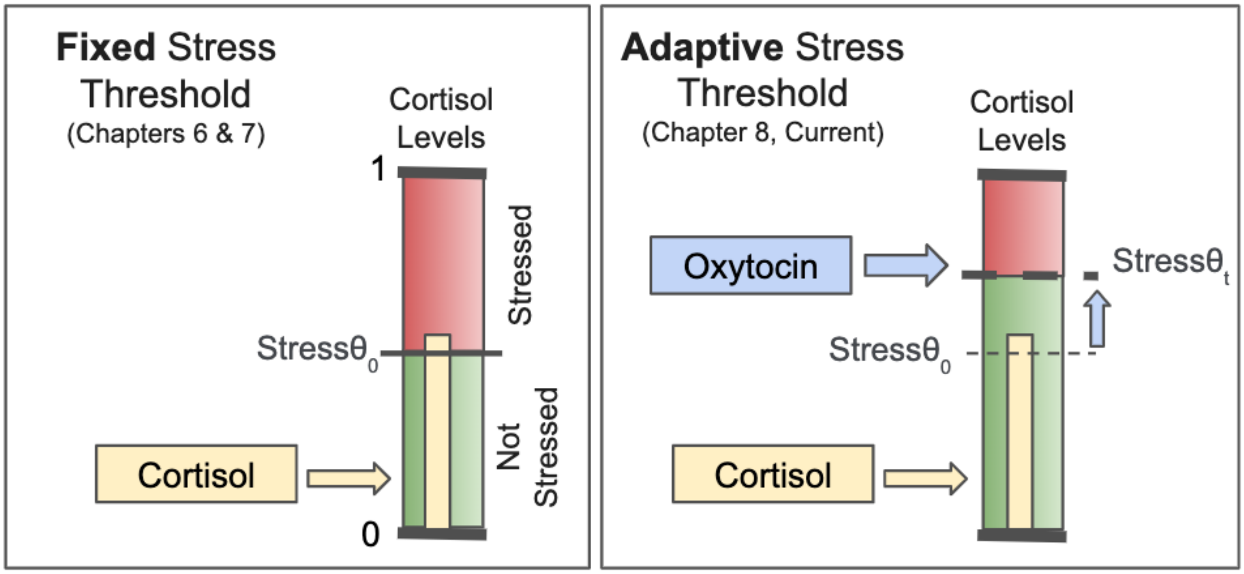

Figure 8.3 Illustration of how oxytocin adapts the internal tolerances to stress—the *StressThreshold*—of each agent. This stress tolerance (Stressθ) determines how much cortisol (yellow bar) an agent can withstand before it becomes stressed (red zone). Left: Our previous model, where the stress tolerance (Stressθ_0) was a fixed value (in this example, 0.5). Right: The adaptation of the Stress Threshold through oxytocin investigated in this chapter. The default Stress Threshold value (Stressθ_0) is modulated by oxytocin levels, with a value of (Stressθ_t) at each given time step (Equation 8.1). (Stressθ_t) can take the range 0.25–0.75.

Each agent was initialised with the same default *StressThreshold* value of 0.5. This was equivalent to the "Neutral" Stress Threshold conditions we had investigated previously (Chapters 6 and the value used in 7). For agents with an affective bond, this value is modulated by its level of oxytocin at any given time step:

$$\text{Stress}\theta_t = \text{Stress}\theta_0 \times (0.5 + OT) \tag{8.1}$$

Using this equation, Stressθ_t remains in the range 0.25 to 0.75, which is approximately equivalent to the lower and upper range of Stress Thresholds we had investigated in Chapter 6 (the *Low* and *High* conditions, respectively).

As with previous investigations, oxytocin was initialised for all bonded agents at its maximum value of 1. The *StressThreshold* for non-bonded agents remains fixed at the default value (0.5).

Condition Label	Bond Type	Description	Effects of Social Support	Chapter of Investigation
C	None	All 6 agents unbonded.	-	Chapter 7 (Experiment 1)
FB	Fixed	3 agents bonded, 3 unbonded. Affective bond strength remain fixed values (DSI=2).	Reduce CT Release	Chapter 7 (Experiment 2)
VB	Variable	3 agents bonded, 3 unbonded. Affective bond strength changes over time.	Reduce CT Release	Chapter 7 (Experiment 3)
VB+	Variable	3 agents bonded, 3 unbonded. Affective bond strength changes over time.	Reduce CT Release OT Adapts Stress Tolerance	Chapter 8 (Current)

Table 8.1 Table describing the different experimental groups discussed in our results. Experimental results from C, FB and VB groups are reused from Chapter 7. Results for the VB+ group are presented through the experiment in this chapter. OT = Oxytocin. CT = Cortisol. “Reduce CT Release” relates to the social buffering effects discussed in Section 7.2.3. “Adapt Stress Tolerance” relates to the modulation of the *StressThreshold* presented in this chapter (Section 8.2.1).

8.3 Experimental Set Up

As a direct extension to the previous chapter (Chapter 7), an almost-identical same experimental set up was retained for this investigation. Experiments are defined using two parameters: three conditions related to affective bond partner combination, and three world conditions related to food availability.

All experiments in this investigation retain the “Variable Bond” type used in Chapter 7. As a reminder, this bond type was initialised at its maximum value of 2, with bond strength (DSI) subject to change as a function of pro, and anti-social interaction from other bond partners.

We tested three combinations of social bonds as investigated in Chapter 7: A1-A2-A6 (Bond Combination A), A3-A4-A5 (Bond Combination B), A4-A5-A6 (Bond Combination C). These bond combinations can be seen in Figure 8.4.

Experiments were conducted across the same three world conditions as used in Chapter 7: a static world where food remains fixed in the environment (*Static*), and two dynamic environments where food either steadily decreases and increases periodically (*Seasonal*), or changes more dramatically (*Extreme*). Screenshots of these worlds can be seen in Figure 7.5, with a description of food availability in each world phase in Table 7.2. This created a total of nine experimental conditions.

We compare the results from this experiment to those of Chapter 7. We focus particularly on comparisons with the “Variable Bond” group from that investigation (Experiment 3 of Chapter 7) as the current model builds directly upon that bond type. For simplicity of comparison, we abbreviate the current and previous

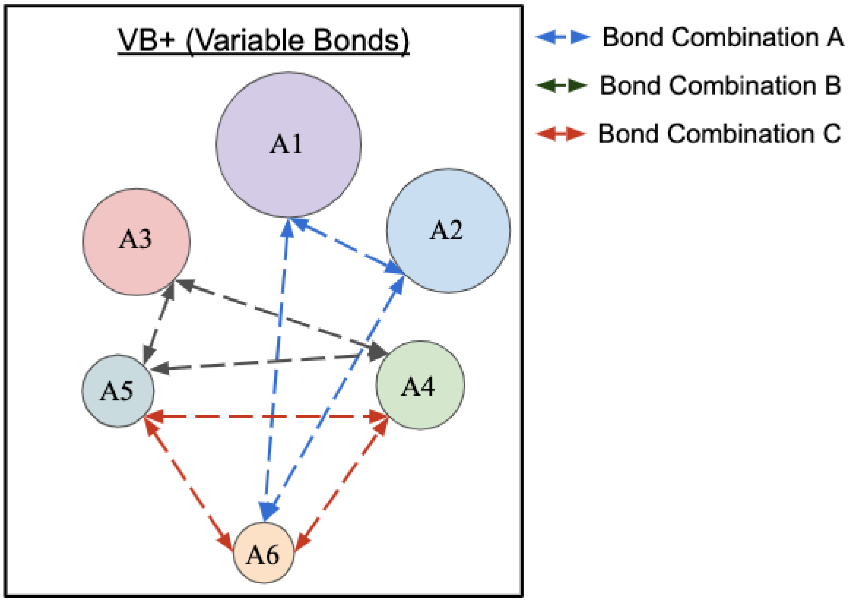

Figure 8.4 Illustrations of the three different affective bond combinations used in this experiment. These bonded agent combinations remain the same as those investigated in Chapter 7. Dashed lines illustrate variable bond strength (DSI changes over time). Agent number denotes rank (higher is better).

experimental conditions. Control conditions (with no affective bonds across the society, Experiment 1 in Chapter 7) are referred to as “C”, results from the Fixed Bond condition (Experiment 2, Chapter 7) are referred to with “FB”, and the Variable Bond results (Experiment 3, Chapter 7) are referred to as “VB”. We refer to the results from this experiment with “VB+”, denoting the extension made from the VB group. These condition names along with brief descriptions of each can be seen in Table 8.1.

In total, 20 simulation runs were performed for each one of the 9 different experimental conditions (three conditions of affective bond combinations, and three world conditions) resulting in a total of 180 simulation runs. As with previous experiments, simulations were run for 15,000 time steps, resulting in approximately 18 hours of run time. Despite the comparisons to previous conditions, we did not capture additional data from those experiments, instead using the results we had already obtained in that chapter.

We use the same three viability-related metrics as previous chapters: *Life Length*, *Mean Comfort*, and *Physiological Wellbeing* (Chapter 2.4.1). In line with our approach in Chapters 6 and 7, we also reported results on the distribution of social behaviours (*Grooming* and *Aggression*) at an aggregated and group level. We captured and reported internal hormonal levels (oxytocin and cortisol) for all agents at each time step, as well as *StressThreshold* values for each agent. We then analysis these results further with qualitative

Bond Condition	World Condition	Comfort (Bonded)	vs. VB	Comfort (Unbonded)	vs. VB	Physiological Wellbeing (Bonded)	vs. VB	Physiological Wellbeing (Unbonded)	vs. VB
A	Static	0.85	+18%	0.82	+17%	0.89	+9%	0.83	+8%
B	Static	0.86	+19%	0.83	+19%	0.97	+13%	0.93	+19%
C	Static	0.79	+10%	0.81	+16%	0.87	+15%	0.98	+1%
A	Seasonal	0.83	+26%	0.74	+17%	0.72	+26%	0.66	+38%
B	Seasonal	0.83	+26%	0.72	+14%	0.77	+43%	0.69	+30%
C	Seasonal	0.81	+23%	0.79	+25%	0.73	+46%	0.69	+12%
A	Extreme	0.74	+19%	0.69	+11%	0.72	+17%	0.66	+34%
B	Extreme	0.77	+24%	0.73	+18%	0.83	+38%	0.75	+27%
C	Extreme	0.70	+13%	0.68	+10%	0.73	+37%	0.73	+3%

Table 8.2 Table describing the results of the *Average Comfort* and *Physiological Wellbeing* metrics, broken down by bonded and unbonded agents, across all world conditions. Results are compared to these same conditions in the VB group (Chapter 7, Table 7.9).

analysis of simulations.

Statistical significance testing was performed between-groups using one-way ANOVA testing, with significance reported at p = 0.05. We use Pearson's Correlation Coefficient to report metric correlation.

8.4 Results

8.4.1 Viability Indicators

We observe significant improvements of the overall society across all three viability indicator metrics (*Life Length*, *Average Comfort* and *Physiological Wellbeing*), when compared to groups with no social bonds (Control), or with no adaptation of stress tolerance (FB and VB). These results can be seen in Figure 9.7.

We noted statistically-significant improvements in *Life Length* performance for the overall society when compared to the VB group, with improvements of between 20%–70% in *Static* environments, 56%–75% in the *Seasonal* environments and 37%–72% in the *Extreme* conditions ($p < 0.05$ for all conditions).

We noted that the improved *Life Length* performance of the society was driven by improved performance of both bonded and unbonded agents (Figure 8.6). Bonded groups saw *Life Length* improvements of 27%–71% in the *Static* environment, 38%–74% in the *Seasonal* environment, and 41%–63% in the *Extreme* environment, vs. the VB groups. Significant *Life Length* improvements were also noted for unbonded groups (Figure 8.6, bottom row): 13%–81% in the *Static* environments, 50%–67% in the *Seasonal* environments, and 32%–86%

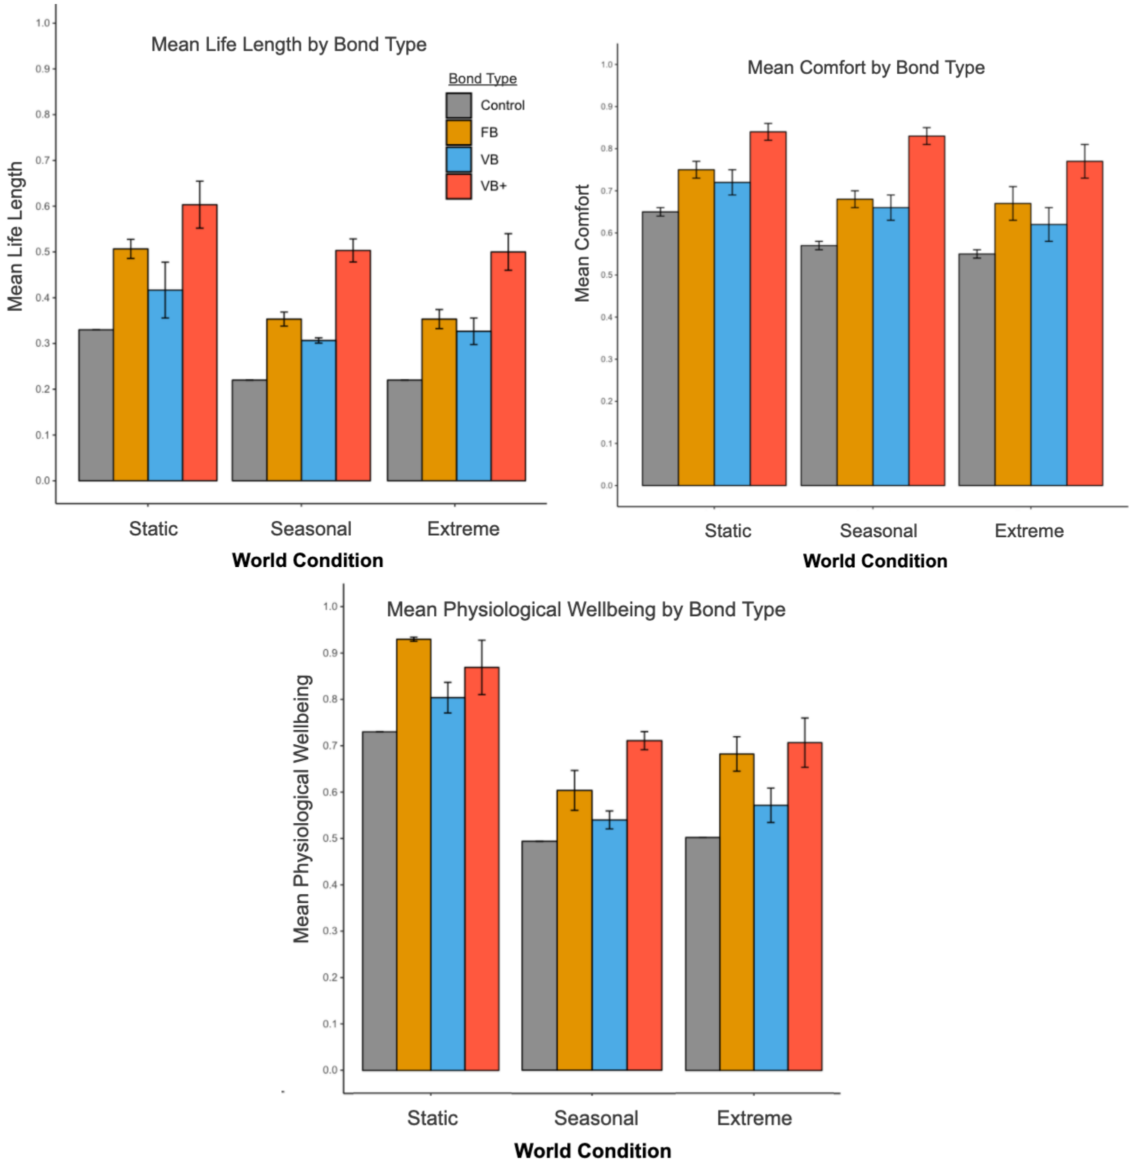

Figure 8.5 Mean values of all three Viability Indicator metrics, aggregated across all simulation runs. Results from Control, FB and VB groups are taken from Chapter 7. VB+ results are new results in this chapter. Control: Conditions with all agents unbonded (Experiment 1, Chapter 7). FB: Conditions of Fixed Bonds (Experiment 2, Chapter 7). VB: Conditions of Variable Bonds (Experiment 3, Chapter 7). VB+: Conditions of Variable Bonds with Stress Tolerance adaptation (current). Error bars show standard error of the mean (SEM) for each group.

in the *Extreme* conditions vs. VB conditions.

Despite no statistically-significant differences in Chapter 7, we observed statistically-significant improvements to *Mean Comfort* levels for the VB+ groups when compared to both the FB and VB groups. *Mean Comfort* levels ranged from between 0.82–0.85 in the *Static* environments, 0.81–0.84 in the *Seasonal* envi-

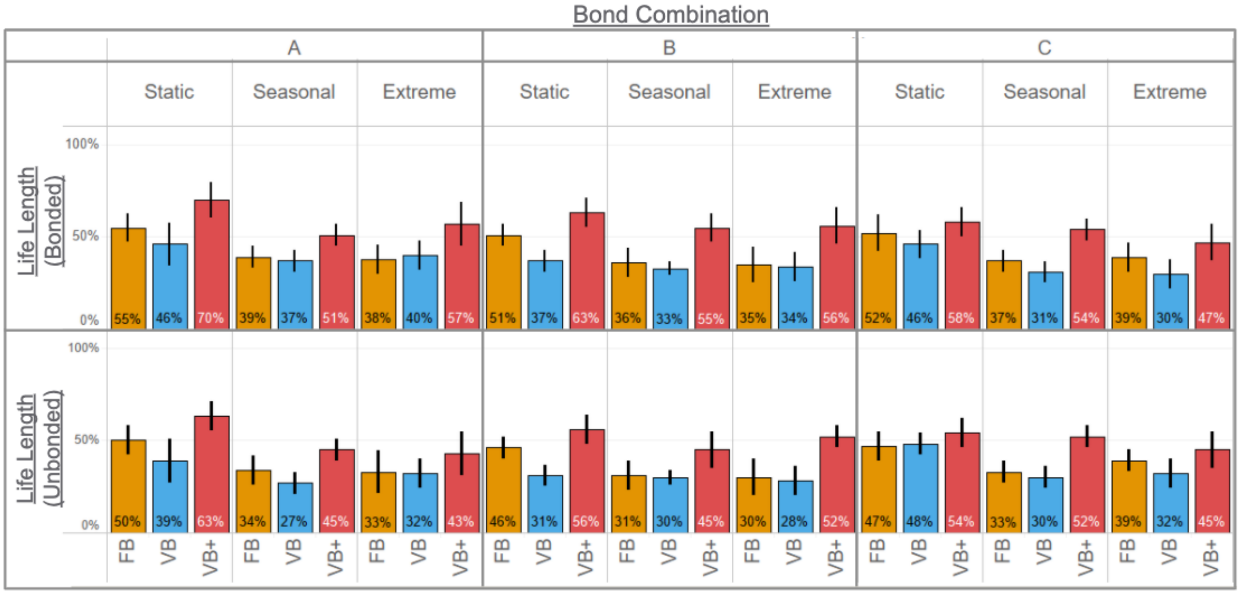

Figure 8.6 Mean Life Length values for bonded agents (top) and unbonded agents (bottom) across all world conditions. Error bars denote standard error of the mean (SEM) for each group. FB: Fixed Bonds (Experiment 2, Chapter 7). VB: Variable Bonds (Experiment 3, Chapter 7). VB+ Variable Bonds with Stress Tolerance adaptation (current chapter).

ronments and 0.73–0.80 in the *Extreme* environments. All results were statistically-significant differences compared to both the FB ($p = 0.03$) and VB ($p = 0.02$) groups. These improvements were seen in both bonded and unbonded groups vs. VB groups (Table 8.2).

We found statistically-significant improvements in *Physiological Wellbeing* at an overall group level vs. VB groups in the *Seasonal* and *Extreme* environments ($p < 0.01$), and non-significant improvements in the *Static* environment ($p = 0.065$). Improveed performance was driven by both bonded and unbonded agents, with bonded agents experiencing an improvement of between 9%–46%, and unbonded agents reporting an improvement of 1%–38%. Table 8.2 shows these results by each bond condition.

To summarise, we saw significant improvements in viability-related performance for both bonded and unbonded agents when bonded agents adapted their stress tolerances, compared to groups where this adaptation was not present (FB and VB groups, Chapter 7).

8.4.2 Social Interactions

Much like previous results, we found that the overall *Aggression* rates of the society to be correlated with the relative challenge of the world environment (Figure 8.7). *Static* world conditions had the lowest aggression rates (6%–9%), with higher ranges seen in the *Seasonal* (10%–14%) and *Extreme* (14%–18%) world

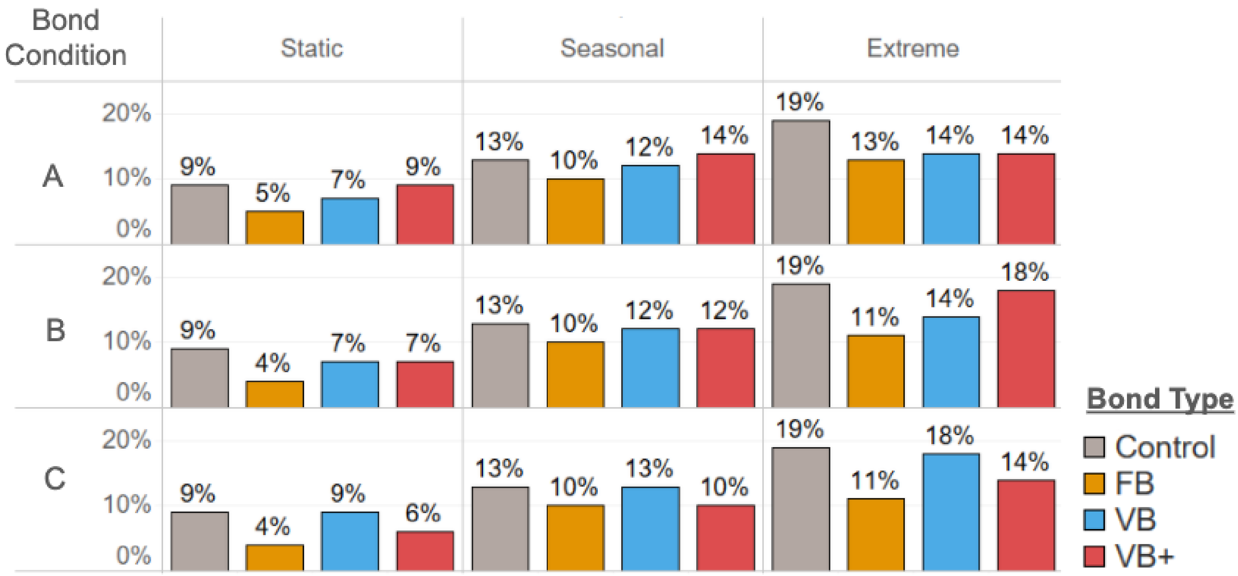

Figure 8.7 Mean Rates of Aggression (defined as the number of *Aggression* interactions as a percentage of all social (Grooming and Aggression) interactions of the overall society. Results are broken down by World Condition (columns) and Bond Condition (rows). Control, FB (Fixed Bond) and VB (Variable Bond) groups are results from Chapter 7. VB+ denote the results of the experiment in this chapter.

conditions.

Compared to the groups with no adaptation of the *Stress Threshold* (Chapter 7), these aggression rates are higher compared to the FB (Fixed Bond) groups, but comparably-similar to the VB (Variable Bond) groups (Figure 8.7). Specifically comparing these results to the VB groups, aggression rates were 2% higher in *Static* and *Seasonal* environments for VB+ groups in Bond Condition A (Figure 8.7, top row), but 3%–4% lower across all environmental conditions in Bond Condition C (Figure 8.8, bottom row).

Across all world conditions and bond combinations, we observed both higher rates of intra-bond *Grooming* and lower rates of intra-bond *Aggression* (Table 8.3) compared to the experiments of the previous chapter (Tables 7.6 (FB) and 7.9 (VB), respectively), highlighting a shift towards prosocial interactions for bonded agents in this experiment. In Table 8.3, we highlight notable differences to the VB group in grey.

Similar to the Variable Bond (VB) group (Table 7.9) intra-bond *Grooming* remained high in Bond Condition A in VB+ groups (Table 8.3), accounting for between 63%–68% of total *Grooming*. We also saw significant increases in intra-bond *Grooming* for agents in the other two bond conditions (B & C) vs. VB groups. For Bond Condition B, *Grooming* between agents increased from 37% to 54% in *Static* environments, from 42% to 45% in *Seasonal* conditions and from 29% to 47% in *Extreme* conditions. For

Bond Condition	World Condition	Intra-Bond Grooming %	Bonded Grooming #	Intra-Bond Aggression %	Intra-Bond Aggression #
A	Static	68%	1,755	25%	16
A	Seasonal	63%	859	27%	23
A	Extreme	65%	523	23%	18
B	Static	54%	1,634	76%	36
B	Seasonal	45%	824	52%	28
B	Extreme	47%	805	73%	52
C	Static	23%	1,021	100%	43
C	Seasonal	28%	603	100%	27
C	Extreme	35%	605	100%	53

Table 8.3 Table showing the total amount of Intra-Bond *Grooming* and *Aggression* across each of the Bond Conditions and World Conditions in this investigation. Intra-Bond Grooming/Aggression values describes what percentage of total grooming/aggression performed by bonded agents was performed towards their affective bond partners. Notable differences to results of VB groups (Experiment 3 from Chapter 7, Table 7.9) are highlighted in grey.

Bond Combination C, similar trends were observed, increasing from 15% to 23% in *Static* environments, 17% to 28% in *Seasonal* environments and 22% to 35% in *Extreme* environments. For Bond Combinations B & C specifically, intra-bond *Grooming* also saw significant increases compared to the Fixed Bond experiment (Table 7.6).

In sum, prosocial interactions between bonded agents in this experiment saw a significant increase compared to both groups from the previous investigation (FB and VB).

Intra-bond *Aggression* rates were lower in the VB+ groups compared to VB groups (Table 7.9), particularly for Bond Combinations A % B in the *Seasonal* and *Extreme* conditions. We also note a significant reduction in the total count of intra-bond *Aggression* encounters in these groups (Table 8.3, last column) compared to the VB group (Table 7.9, last column).

Therefore, both the rates and absolute numbers of intra-bond *Aggression* were significantly reduced when agents with affective social bonds were endowed with two of the social buffering effects (VB+), compared to groups with only one effect accountedfor 7 (VB). Combined with the improved intra-bond *Grooming*, social behaviours between bonded agents were much more prosocial in this experiment compared to previous experiments.

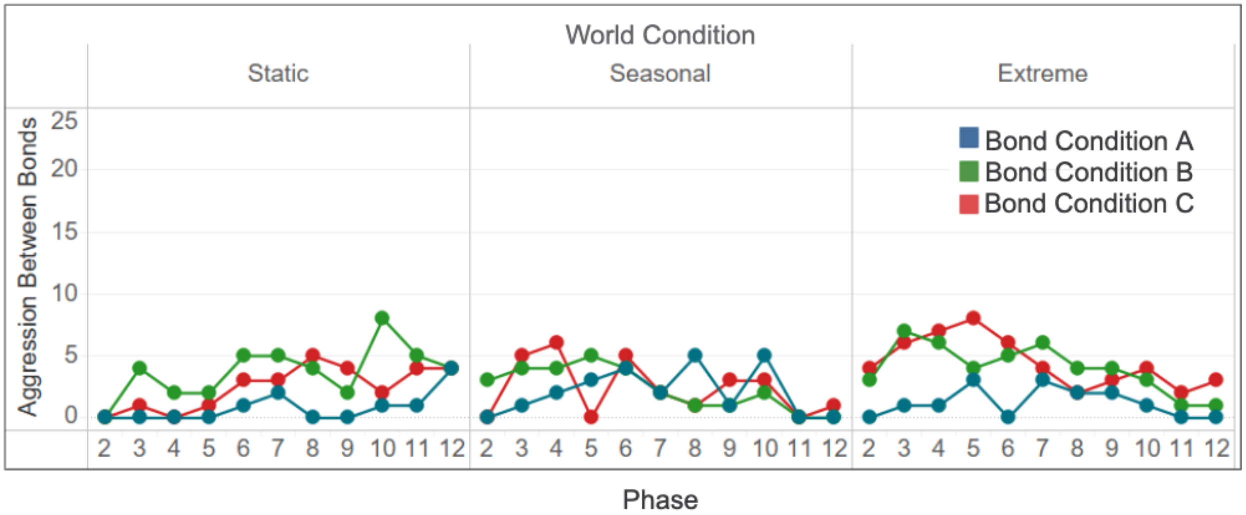

Figure 8.8 Total amount of intra-bond *Aggression* (y-axis) (*Aggression* performed between the three bonded agents) during each phase of food availability (x-axis), aggregated across all simulation runs. Bond Condition A: A1-A2-A6 bonded. Bond Condition B: A3-A4-A5 bonded. Bond Condition C: A4-A5-A6 bonded. Further data can be found in Table 8.3.

Unlike the intra-bond *Aggression* seen in Chapter 7 (Figure 7.9), socio-negative behaviours towards bond partners was less affected by the dynamic world environments, as seen by the trended intra-bond *Aggression* (Figure 8.8). Regardless of the world condition, there were non-significant changes in the *Aggression* between bonded agents across different phases. This was despite the *Aggression* rate of the overall society increasing across the challenging environments (Figure 8.7).

At an individual level, we find reduced rates of aggression for agents when they are part of a bonded group, compared to conditions when they are not (Table 8.4). This is most notable in Bond Conditions B and C, where *Aggression* reduced from between 9%–25% when agents were not bonded, to 0%–12% when these agents were bonded.

8.4.3 Hormone Levels

Figure 8.9 shows the comparison results of mean cortisol levels between bonded and unbonded agents from our previous investigation (FB and VB, Chapter 7) and the groups with stress tolerance adaptation in this chapter (VB+).

Mean cortisol levels for the VB+ group was lower when compared to both the fixed bond (FB) and variable bond (VB) groups. Lower cortisol levels were seen in both bonded and unbonded groups across all world conditions, although bonded agents experienced the largest improvements compared to the VB group.

		Aggression Rates					
Bond Condition	**World Condition**	**A1**	**A2**	**A3**	**A4**	**A5**	**A6**
A	Static	5%	4%	17%	17%	9%	0%
	Seasonal	18%	8%	18%	24%	10%	0%
	Extreme	13%	7%	21%	25%	14%	0%
B	Static	9%	13%	4%	2%	0%	0%
	Seasonal	21%	17%	6%	4%	2%	0%
	Extreme	27%	23%	9%	10%	1%	0%
C	Static	6%	9%	16%	4%	2%	0%
	Seasonal	22%	16%	22%	6%	2%	0%
	Extreme	28%	26%	29%	12%	3%	0%

Table 8.4 Mean rates of aggression for each individual agent across all bond and world conditions. Results in grey highlight agents who were bonded in those conditions.

Bond Condition A saw reductions in cortisol levels between 15%–21%, and Bond Condition C saw reductions of between 15%–49%. Bond Condition B saw reductions of between 9%–17%, with one exception (in the *Extreme* world condition), where cortisol levels were similar to VB groups (0.60 to 0.59).

Overall, and in line with previous results, we find overall lower stress levels for the society when half of the society has affective bonds. We suggest that cortisol levels in agents that shared an affective social bond resulted in reduced mean cortisol levels to for agents without social support.

For all bond combinations across all environments, we noted significant increases in mean oxytocin levels when compared to the VB group (Table 8.5, $p < 0.01$)). Bonded agents in Bond Condition A saw internal oxytocin levels increase between 7%–75%; Bond Condition B saw increases of between 46%–57%, and Bond Condition C saw increases of between 42%–57%. Given that the direct effects on oxytocin remained unchanged between these groups, this is a noteworthy result.

Figure 8.10 shows the trended mean levels of internal oxytocin for the three bonded agents in each world condition. Despite their relative high dominance ranks, agents in Bond Condition A reported the lowest mean oxytocin levels (0.55–0.64), compared to the lower-ranked agent combinations in Bond Condition B (0.67–0.74) and C (0.65–0.74). For the latter groups, this increased oxytocin was seen despite slightly higher cortisol levels (Figure 8.9) compared to Bond Condition A.

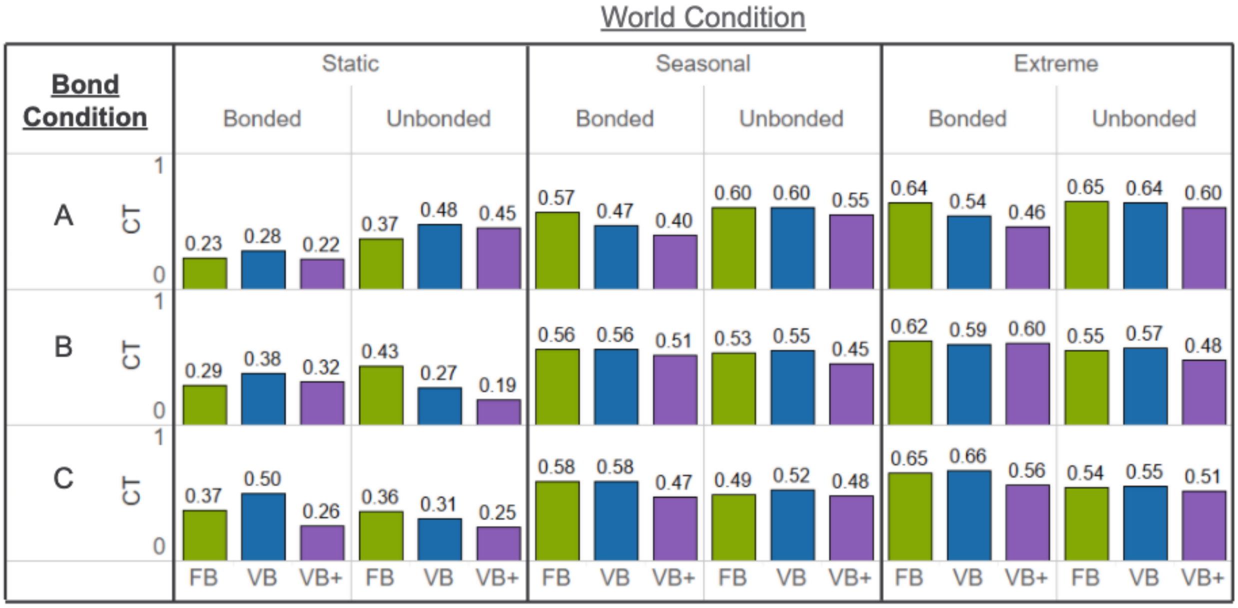

Figure 8.9 Comparison of the mean cortisol levels for bonded and unbonded agents across all bond types, aggregated across all simulation runs. FB: Fixed Bonds (Experiment 2, Chapter 7. VB: Variable Bonds group (Experiment 3, Chapter 7). VB+ Variable Bonds with tolerance adaptation, from the modeul used in this chapter. Cortisol levels exist in the range 0 to 1 and are only recorded for agents that are still alive in each condition.

8.4.4 Stress Threshold Trends

We find that the *StressThreshold* of bonded agents to decrease across the three challenging environmental conditions. In the *Static* environment (Figure 8.11, top), mean *StressThreshold* values were 0.41, 0.55 and 0.53 for bonded agents in Bond Condition A, B and C respectively; 0.32, 0.41, and 0.41 in *Seasonal* conditions and 0.33, 0.48 and 0.41 in *Extreme* conditions. Table 8.6 breaks down the mean *StressThreshold* values for each agent across these conditions. All unbonded agents in each condition maintained a fixed value of 0.50.

Despite the fact that the two most dominant agents (A1 and A2) were bonded in Bond Condition A, we found that the mean *StressThreshold* value of the bonded group was significantly lower in this group compared to both Bond Condition B (*Static*: $p = 0.034$, *Seasonal*: $p = 0.023$, *Extreme*: $p < 0.01$) and Bond Condition C (*Static*: $p = 0.04$, *Seasonal*: $p = 0.023$, *Extreme*: $p = 0.02$).

We note that this was a result of the two most dominant agents (A1 and A2) adapting their *Stress Threshold* values to low values, whilst the subordinate A6 adapted to a higher *StressTolerance* across all world conditions (Table 8.6). For lower-ranked bonded groups (Bond Condition B and C), mean *Stress Threshold* values remained similar between bonded agents.

Bond Condition	World Condition	VB	VB+	Δ
A	Static	0.49	0.64	+30%
	Seasonal	0.37	0.40	+7%
	Extreme	0.31	0.55	+75%
B	Static	0.51	0.74	+46%
	Seasonal	0.44	0.69	+47%
	Extreme	0.45	0.67	+57%
C	Static	0.47	0.74	+57%
	Seasonal	0.46	0.65	+42%
	Extreme	0.42	0.65	+55%

Table 8.5 Mean oxytocin levels for bonded agents across all conditions, aggregated across all simulation runs. VB = Variable Bond groups with no adaptation of the *Stress Threshold* (Experiment 3, Chapter 7). VB+ = Variable Bond groups with adaptation of *Stress Threshold* (current chapter). Δ = percentage difference between VB+ and VB oxytocin levels.

In other words, Bond Condition A experienced a larger variance (0.30, *Static* world condition) of their *Stress Threshold* values than the closer-ranked Bond Condition B (0.07, *Static* world condition) and Bond Condition C (0.04, *Seasonal* world condition).

The trended *Stress Threshold* values seen in Figure 8.11 shows a cyclical adaptation of stress tolerance through the dynamic (*Seasonal*, middle, and *Extreme*, bottom) environments. However, we also note that all groups did not adapt in the same way.

Specifically, during a period of poorer food availability in the *Seasonal* world condition (between 3K-4K time steps when only 2 food resources are available, Table 7.2, Figure 8.11), Bond Condition A adapted to a much lower mean *Stress Threshold* value, whereas the other two bond conditions maintained a steady threshold throughout this period.

Similarly, in the *Extreme* world condition (Figure 8.11, bottom), between 7K-8K (a period of “good” food availability with 4 food resources available, Table 7.2) and then 8K-9K (a phase of “bad” food availability with 1 food resource available), Bond Condition B (green line) adapted a higher *StressThreshold* value across these phases—from 0.40 to 0.47—while no such increase was observed for Bond Conditions A or C.

Taken together, we see how bonded groups adapted their *StressThreshold* to different values, and still reported viability benefits across all conditions (Figure 8.6, Table 8.2). These results indicate a diversity of

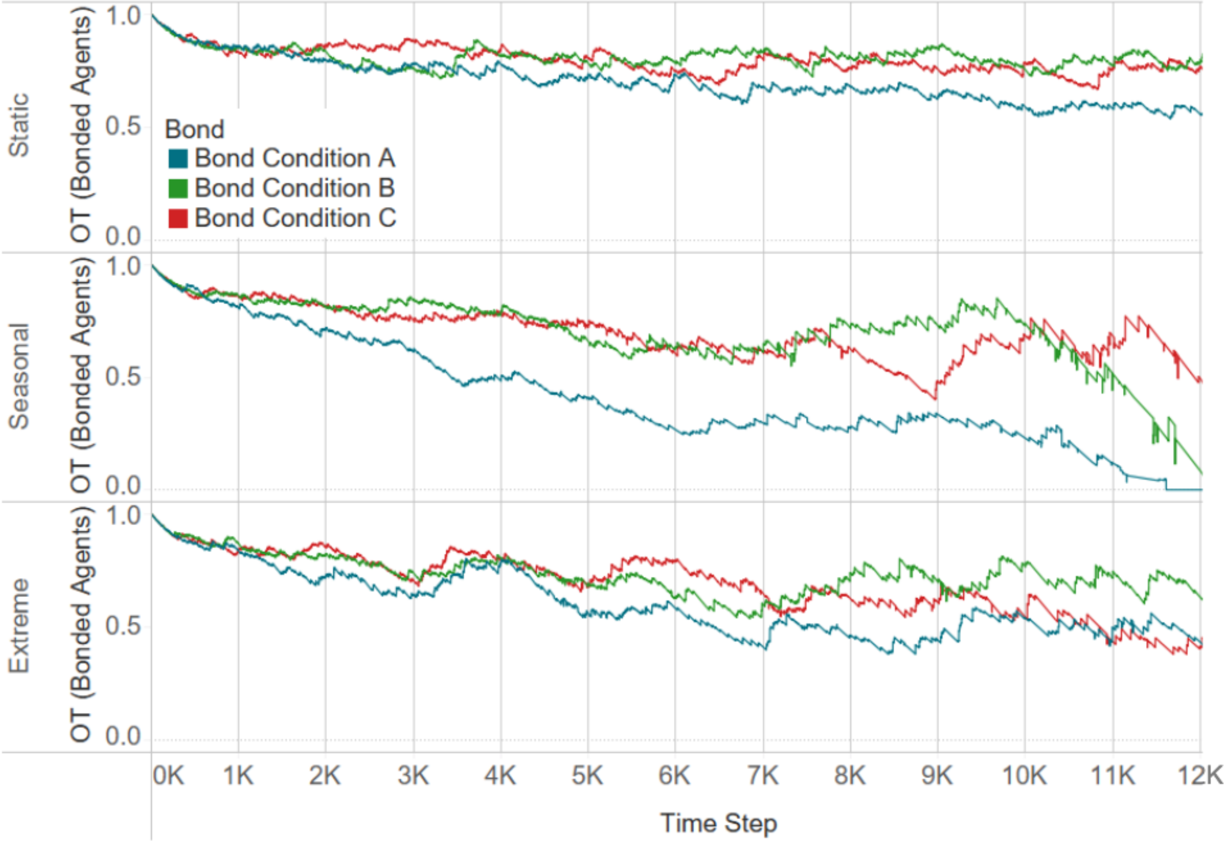

Figure 8.10 Trended mean oxytocin levels of the three bonded agents per world condition. Results are aggregated across all experimental runs and for the three bonded agents in each condition (Bond Condition A: A1-A2-A6. Bond Condition B: A3-A4-A5. Bond Condition C: A4-A5-A6. Though simulations were run for 15,000 time steps, results after 12,000 were not significant and not included.

physiological adaptation between different groups of agents, and show how different strategies to adapting a stress response provided advantages to the long-term viability of these agents.

8.5 Discussion

From our results, we found support for the hypothesis that oxytocin-mediated adaptation of stress tolerance would provide survival-related benefits to agents with social support. Specifically in our model, we found statistically-significant improvements across all three viability-related metrics (*Life Length*, *Average Comfort*, and *Physiological Wellbeing*) for overall societies, largely driven by improvements of agents with social bonds, when compared to groups with no adaptation of stress tolerance. We also found support for our hypothesis

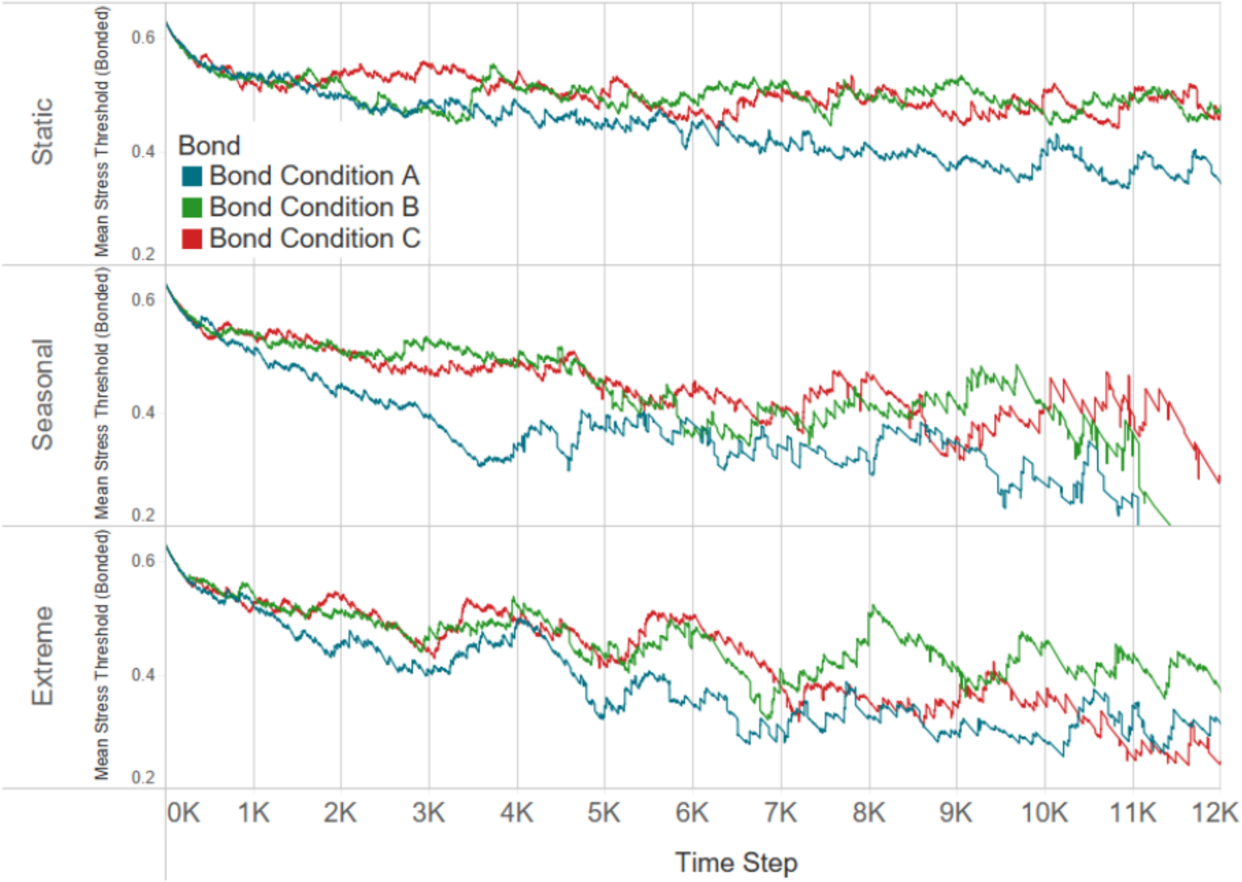

Figure 8.11 Mean trends of internal Stress Thresholds for bonded agents. Results are aggregated across all experimental runs and for the three bonded agents in each condition (Bond Condition A: A1-A2-A6. Bond Condition B: A3-A4-A5. Bond Condition C: A4-A5-A6. Though simulations were run for 15,000 time steps, results after 12,000 were not significant and not included.

that this adaptation of the stress tolerance would result in lower overall aggression rates for bonded agents.

We observed significantly-reduced levels of stress in these groups where two of the stress-reducing effects of social support were accounted for—a reduction of cortisol release, and an adaptation of stress tolerance—compared to when only the former was accounted for (Chapter 7). These results show support for the hypothesis that the stress-reducing effects of social support are diverse and multi-faceted [16], with some potential interdependence of numerous stress-regulating mechanisms. We also note that these effects on stress-reduction were observed in absence of the additional proposed effects of other reward-based systems [131], [228].

We highlight several interesting outcomes. Firstly, adaptation of the stress tolerance for half of the agents in the society resulted in absolute stress levels (cortisol levels) of all agents—both with and without

World Condition	Bond Combination	A1	A2	A3	A4	A5	A6	Mean of Bonded Agents
Static	A	0.31	0.33	0.50	0.50	0.50	0.61	**0.42**
Static	B	0.50	0.50	0.6	0.54	0.53	0.50	**0.56**
Static	C	0.50	0.50	0.50	0.54	0.55	0.52	**0.54**
Seasonal	A	0.29	0.31	0.50	0.50	0.50	0.44	**0.35**
Seasonal	B	0.50	0.50	0.43	0.41	0.42	0.50	**0.42**
Seasonal	C	0.50	0.50	0.50	0.44	0.41	0.4	**0.42**
Extreme	A	0.29	0.31	0.50	0.50	0.50	0.4	**0.33**
Extreme	B	0.50	0.50	0.48	0.52	0.52	0.50	**0.51**
Extreme	C	0.50	0.50	0.50	0.43	0.43	0.40	**0.42**

Table 8.6 Mean values of the *Stress Threshold* for each agent, aggregated across all simulation runs. When agents were not bonded, their default *Stress Threshold* was 0.5. Grey cells indicate agents that shared an affective bond in those conditions

affective bonds—to be reduced. As a result, overall stress levels and socio-negative behaviours of the whole society was also reduced in some conditions, providing benefits to group survival and viability. Our results and observations correspond to our findings from previous chapters (Chapter 6 and 7), where we observed similar types of emotional contagion to emerge in the society as a result of smaller, dyadic interactions. We discuss reasons for this below.

Secondly, groups where oxytocin adapted the stress tolerance reported higher mean levels of oxytocin. Here, the stress-reducing effects of oxytocin-mediated social support resulted in increased oxytocin levels for supported agents, which promoted further adaptation of the stress response and prosocial behaviours: releasing additional oxytocin. As well as corresponding to our earlier findings (Chapter 7), this positive feedback loop of oxytocin mirrors suggestions from biology and neuroscience, suggesting that oxytocin's effects on stress-reduction may be facilitated by its positive feedback loop on seeking prosocial interaction [30]: likely as a result of its effects on other reward-based systems [131], [228]. However, in our model, we found this positive feedback loop to emerge even in absence of other hypothesised reward-based mechanisms, and propose this as an area of focus for further investigation. Though these (oxytocinergic and opioid-based) mechanisms cannot be separated in the real-world, our model can be used to address this challenge and investigate these effects further.

Additionally, we identified that in conditions where dominant agents had an affective bond with the lowest-ranked agent—a relationship akin to a parent-offspring relationship—dominant agents adapted to a slightly lower tolerance to stress compared to the other bonded agents. Though contrary to expectations, these agents nevertheless experienced improved survival and wellbeing. We posit that this may be an outcome with some biological plausibility. Specifically, that higher-ranked individuals who share an affective bond with weaker individuals (i.e. a parent-offspring relationship) must remain more "vigilant" towards stressful stimuli and potentially remain "less tolerant" to threats, particularly in challenging environmental conditions. This result may give insight into the diversity of stress regulation in biological societies, with further implications on how individuals adapt their stress coping strategies (for example, through performing socio-negative behaviours).

We find further support for our suggestion from the previous chapter, that a small society such as the one we have modelled may experience the wider benefits of social support between a smaller group of agents. All individuals in the society may not require the same type, or level, of adaptive or coping mechanisms in order for long-term wellbeing to be improved. With respect to a model of social adaptation for affective agents, we mirror our suggestions from Chapter 7 and again propose that not all agents in a society require the same (types of) adaptive mechanisms in order to experience survival-related benefits. Future work, especially those with physical agents, should take this into consideration.

Overall, these results suggest an important interplay between the adaptation of stress tolerance through oxytocin-mediated social support, its effects on intra-bond social interactions and wider social dynamics. In terms of using our results in the context of our adaptive model of social adaptation, the stress-buffering effects that we have investigated in this chapter propose significant benefits to agent viability through physiological and behavioural adaptation. We discuss a number of our findings in detail below.

8.5.1 Adaptation of Stress Tolerance Promotes "Anticipatory" Behaviours and Survival

From our results, we found that groups where adaptation of the *StressThreshold* (VB+) reported significant improvements in viability compared to groups where this adaptation did not occur (VB). We suggest that this is a result of behavioural adaptation which emerged as a result of adapting tolerances to stress.

Intra-bond *Grooming* during early periods of good food availability maintained high oxytocin levels in bonded agents (Figure 8.10) and adapted their *StressThreshold* accordingly (Figure 8.11). Despite increases in cortisol levels when food availability was challenged (*Seasonal* and *Extreme* environments), *StressThreshold* would not be exceeded and agents would not become stressed (i.e. the "sympathetic nervous system" would not be activated). As a result, intra-bond *Aggression* rates were significantly reduced (Table 8.3), resulting in stronger affective social bonds which continued to provide the stress-reducing effects for bonded agents.

This is contrary to what we had observed in Chapter 7, where an absence of stress tolerance adaptation resulted in higher rates of intra-bond *Aggression* (Table 7.9), limiting the efficacy of social support on stress-reduction which ultimately impacted group viability (Table 7.7). The adaptation of the *StressThreshold* in this investigation created an internal stress buffer during non-challenging situations, and the increased rates of intra-bond *Grooming* created a positive feedback loop that promoted future affective behaviours.

In terms of our agent model, we suggest that these mechanisms underpin low-level "anticipatory" adaptation in the face of upcoming stressful conditions which facilitated individual agent and group viability. These mechanisms, underpinned by oxytocin, can be considered a type of affective adaptation through social allostasis.

8.5.2 Anticipation of Stressful Conditions May Facilitate "Tend-and-Befriend" Behaviours

In our previous investigation (Chapter 7), we had observed higher rates of intra-bond *Aggression* and avoidance—the "fight-or-flight" response [21]—as a stress coping mechanism when affective bond strength was not fixed and subject to change over time (Experiment 3). This was as a combined result of lower bond strength (DSI) and lower oxytocin levels in bonded agents. Conversely, we had seen more socio-positive, "tend-and-befriend" [30] behaviours between affective bond partners when these bond strengths were fixed (Experiment 2, Chapter 7). In those conditions, persistently strong bonds promoted prosocial behaviours between bonded agents. These improved rates of prosocial behaviours between agents with fixed bonds drove increases in oxytocin and bond strength, promoting a positive feedback loop of positive affective interaction throughout the course of the experiment 7.9. With these results, we therefore considered that affective bonds needed to be persistently strong between bond partners in order to promote this stress-coping mechanism, and that bond stability played a significant role in the efficacy of the stress-reducing effects of social support.

However, despite the variable bond type used in this investigation (the VB+ group), we find the emergence of these "tend-and-befriend" behaviours provided an adaptive, anticipatory response for future stressful conditions (Table 8.8). These increased rates of intra-bond *Grooming* provided a stress buffer for bonded agents in future stressful conditions, resulting in lower rates of intra-bond *Aggression*, and retaining the efficacy of the stress-reducing effects of social support in the long-term that was "lost" in Chapter 7.

As a result, the affective relationship—along with the survival-related benefits—between bonded individuals remained strong, despite the challenge presented by the environment. Considering that the "tend-and-befriend" hypothesis of stress-reduction has been considered to be a more female-oriented and maternal response [30], these observed behaviours may be a form of early-stage nurturing [232] found in real-life societies. Given oxytocin's implications in maternal, nurturing-type behaviours [30], [233], [234], we suggest that our model, with its abstractions of oxytocin mechanisms, may present an opportunity to further investigate these hypothesised different gender-based coping mechanisms.

We therefore propose that "fight-or-flight" behaviours seen in VB groups (Table 7.9) may not be a function of poor bond strength as we had previously hypothesised, but a coping strategy performed in the absence of adequate anticipation of stressful conditions. For instance, it may be the case that when future stressful situations have not been anticipated or accounted for (i.e. through a lack of appropriate physiological or behavioural adaptation), "fight-or-flight" behaviours are an immediate behavioural "reflex" response to stressful conditions in an attempt to alleviate stress. Part of our future work may consider investigating the relationship between "early-stage" prosocial interactions and subsequent viability between conditions where adaptation of *StressTolerance* is and is not accounted for, respectively.

In terms of our agent model, these findings suggest that the modelling of affective bonds of variable strength—a more biologically-plausible approach than modelling "fixed" bonds—may not be detrimental to viability as we had previously found (Chapter 7). On the contrary, this type of bond provides significant benefits over its fixed counterpart when additional adaptive mechanisms are considered, and therefore can be used in future development of our model.

8.5.3 Adaptation of Stress Tolerance Results in Lower Stress Contagion

Despite the fact that oxytocin did not have a direct effect on cortisol secretion in our model (Figure 9.3), we still observed lower overall stress levels in the groups where the *StressThreshold* was adapted (VB+) through oxytocin, compared to groups without this mechanism (VB). We found this to be a surprising result, and we considered this to be the emergence of wider physiological and behavioural adaptation which we suggest may works as follows.

As oxytocin "buffers" the *StressThreshold*, higher levels of oxytocin (Table 8.5) resulted in agents being able to tolerate more of the stress hormone (cortisol). As a result of this additional tolerance, bonded agents were less likely to become stressed and perform *Aggression* towards others (Table 8.4). The reduced *Aggression* rates of the society results in reduced levels of cortisol for both bonded and, in some cases, unbonded agent groups (Figure 8.9). This resulted in improved viability for both bonded and unbonded groups, compared to groups where adaptation of *Stress Threshold* did not occur (FB and VB groups, Figure 8.6).

Similar to the contagion of socio-negative behaviours and stress that we saw in previous investigations (Chapters 6 and 7), we again find that, in a small society of agents, the wider social dynamics were significantly affected by the interactions of smaller groups of (bonded) individuals. Taking all of these findings into consideration (Chapter 6, 7 and the current chapter), we highlight the consistent emergence of affective contagion through interactions of affective bond partners.

We suggest that these micro-level interactions and affective contagion may play a role in explaining the diversity of social dynamics seen in biological societies [117], and may also have further consequences on long-term physiological health.

In terms of overall stress levels, we also find that the combination of the two stress-reducing effects of social support that we accounted for in this investigation provided a significant benefit to the reduction of overall social stress 8.9, compared to conditions when only one effect was considered (Chapter 7).

Taking these results into consideration, we also propose that—perhaps counter-intuitively—the stress-reducing effects of the presence of social support may not provide the most benefit to the reduction of stress in the long term. Instead, long-term stress-reduction may be a multi-faceted adaptive process: facilitated only in part by the direct effects on reducing cortisol secretion (Chapter 7), but a result of additional behavioural and physiological adaptation driven by other stress-regulating effects (i.e. such as the adaptation of stress tolerance).

This suggestion also mirrors those from neuroscience [235]–[237] suggesting that the stress-reducing effects can be diverse, long-lasting and multi-circuited. In addition, these socially-based, stress-reducing effects may even extend beyond tactile or visual interaction with bond partners [238]–[241], suggesting that affective bond presence or physical interaction may not always be necessary to regulate stress in the short-term.

8.5.4 Parent-Offspring Bond Were Less Tolerant to Stress

Contrary to our expectation that higher-ranked agents would report higher levels of oxytocin and adapt their stress response to cope with higher levels of stress (i.e. be more "relaxed" during stressful situations), our results did not find this to be the case.

Instead, these higher-ranked agents (A1 and A2) in Bond Condition A—which we considered to be a "parent-offspring" bond—reported slightly lower oxytocin levels across all environmental conditions (Table 8.10), as well as a lower mean *StressThreshold* value (Figure 8.11, Section 8.4.4).

Similarly, based on our previous investigation—where lower-ranked bonded agents (Bond Condition B and C) experienced lower rates of intra-bond *Grooming* (Table 7.9)—and the fact that cortisol is released as a function of perceived control and socio-negative interactions, we had predicted that lower-ranked bonded agents (in Bond Condition B and C) would adapt their *StressThreshold* values to lower values than higher-ranked agents. Instead, we found that these subordinate agents adapted a higher tolerance to stress than the higher-ranked agents in these Bond Conditions.

Despite the fact that all agents started with the same default tolerance to stress—which we modelled as a personality trait (as we discussed in Chapter 5)—the adaptation of their *StressThreshold* in this investigation can be considered as an adaptation of personality, emerging as a function of the social and physical interactions with the environment that they are situated in.

As we saw viability increases for bonded agents (Figure 8.6) when compared to previous groups (VB), these different levels of adaptation of individual "personality traits" (their stress tolerance) was nevertheless an adaptive response that facilitated survival for all agents in the bonded group.

Considering that stress is an important adaptive response that can promote survival [25], [134], these findings may shed light on the diversity of stress-regulating physiological and behavioural adaptation across different social and environmental contexts [242]. Specifically, individuals of different ranks may need to adapt their "personality" differently, based on the environment and the identity of their social partners.

The type of affective relationship an individual has with another may affect their regulation of stress-related behaviours. For instance, higher-ranked individuals who share an affective bond with a much lower-ranked individual (as is case in a parent-offspring relationship) may need to remain more vigilant or "protective" of lower-ranked partners in certain (social and environmental) situations. Though speculative, this dynamic of "protective" or "hyper"-vigilance from parental figures may have some biological plausibility [243]–[245] and these oxytocin-mediated effects may play a role in such coping (or parental) strategies.

However, we propose that this nevertheless proposes an avenue for further investigation. Future work may use our model to investigate additional affective bond combinations (for instance, affective bonds between higher-ranked and middle-ranked agents) to investigate these adaptive effects further.

8.6 Summary

Continuing our investigation into the "social buffering" phenomenon associated with social support, we have investigated a second hypothesised stress-regulating effect of social support. Here, we consider that oxytocin—released as a function of prosocial interactions with bond partners—has an adaptive effect on individual stress tolerance. This stress tolerance can be considered a personality trait of an individual, and associated with the "resilience" or activation of their stress response system.

By also accounting for the effect of social support on ameliorating cortisol release from Chapter 7, we used our simulated model to investigate how the oxytocin-mediated regulation of stress tolerance affected the survival-related performance (viability), as well as the social interactions and dynamics, of a small society of agents.

Based on our previous investigations, we had hypothesised that the adaptation of the stress response via oxytocin mechanisms would result in a decrease in the rates of socio-negative behaviours between agents with affect-based social support, and that this behavioural adaptation would provide viability-related advantages to these agents when compared to conditions without this adaptative effect. We also predicted that these viability performance improvements would be seen in agents without social support.

We conducted our experiments using an identical experimental set up to Chapter 7: testing our hypothesis in a society of six agents across three different bond partner combinations related to social rank. We studied these bonds across three different environmental conditions related to food resource availability, posing different degrees of challenges.

Our results found support for the hypothesis that the stress-reducing effects of social support provides adaptive, survival-related benefits to individuals through oxytocin's effect on the stress system. We saw significantly improved viabilty performance compared to groups where we had only accounted for the stress-reduction effects of social support (Chapter 7). We also found significantly lower rates of socio-negative behaviours between bond partners, as well as some viability-related improvements for agents without social bonds, suggesting widespread viability benefits associated with (some agents being endowed with) such stress-reducing, adaptive mechanisms.

Given the improved performance benefits compared to the investigation in Chapter 7, and also mirroring the findings from Chapter 3, we use this opportunity to highlight the diverse results associated with incremental changes to the agent model. Though an artificial life approach has allowed us to minimise complexities, the (sometimes-paradoxical) results that emerge as a result when accounting for additional effects in our model may give insight into some of the diverse, complex findings from real-world societies.

Much like our findings from the previous chapter (Chapter 7), we find that the behaviours and affective states of the wider society was impacted by interactions between bonded agents. This affective contagion had widespread viability-related benefits for individuals with and without affective bonds. Once again, we highlight that affective social agents may not need to be endowed with the same type, or level, of adaptive mechanisms in order to benefit from the performance advantages and behavioural adaptation associated with adaptive mechanisms. This finding may have significant implications in the development of future models, particularly with physical agents.

In our last two chapters (Chapter 7 and 8), our investigations have been focused on investigating the stress-regulating effects of social support: by either reducing the level, or secretion, of a stress hormone (cortisol) (Chapter 7), or by adapting an individual's tolerance of this stress hormone (Chapter 8).

We found how individuals in this chapter adapted their stress tolerances differently based on the environmental challenge and identity of bond partners. Although our results found indirect physiological adaptation of the hormonal systems as a result, our observations made us consider how adaptation of stress tolerances may also play a more direct role in regulating homeostatic processes. Specifically, we considered that adaptation of the internal stress system may also provide a more adaptive role on individual physiology by "buffering" some of the deleterious consequences associated with elevated cortisol levels.

Despite extensive literature finding correlations, the direct mechanisms and pathways that tie the presence of social support and improvements in health and physiological outcomes are still unclear. Regardless, this socially affect-based physiological adaptation is a fundamental aspect of social allostasis, and further investigating these mechanisms through our model may provide insight into the diverse physiological effects of stress in real-world individuals.

In our final chapter, we continue our investigation into the role of social support on promoting adaptation and long-term wellbeing via its stres-reducing effects. Building on our existing findings and addressing methodological challenges of real-world models, we propose and test a hypothesis for how affect-based adaptation of the stress response can play a role in ameliorating the physiological "cost" associated with stress.

Chapter Nine

Investigating the Role of Oxytocin-Mediated Social Support on Adapting the Physiological Cost of Stress

9.1 Introduction

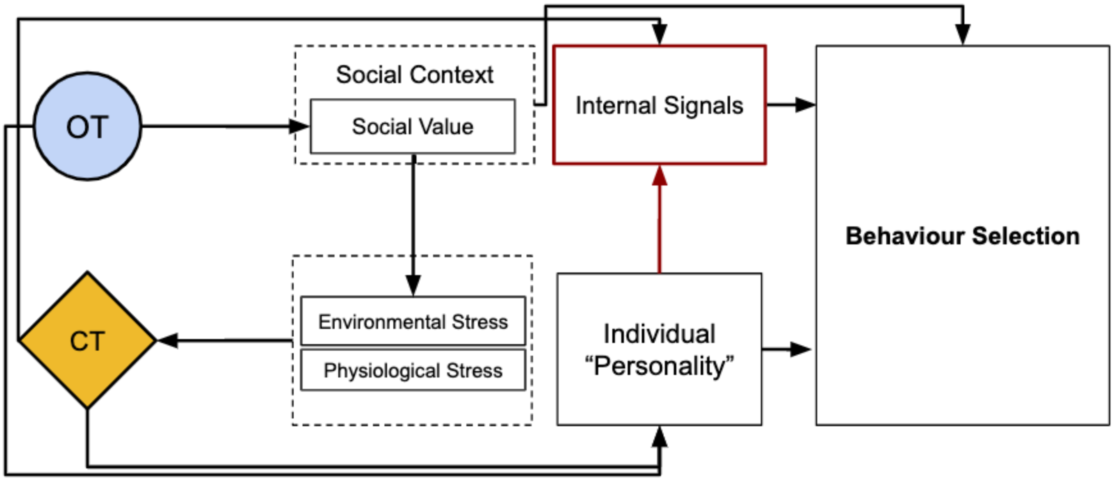

Figure 9.1 A high-level view of the effects investigated in this chapter. Red lines indicate extensions from the previous chapter. In this investigation, oxytocin-mediated "buffering" of individual stress responses ("personality") also dampens the physiological "cost" of stress (energy metabolism). A more granular view of this model can be seen in Figure 9.3

.

In Chapter 8, we found how accounting for two of the hypothesised "stress-buffering" effects of social support—by mediating cortisol release, and the adaptation of stress tolerance—provided significant performance advantages compared to groups where just the former effect was accounted for (Chapter 7). We also found how individuals adapted their stress tolerances differently, related to the challenge of the physical environment and the identity of their bond partners. As a result, we considered how these individual differences in adapting to, or being able to cope with, higher levels of stress may also play a role in different physiological responses to the stress hormone. In this chapter, we investigate these individual differences further, and account for social support's potential buffering effects on one stress-related physiological process.

Extensive literature has found relationships between the presence of social support, stress-reduction, and improvements in physical health markers of socially-supported (human and non-human) animals. Though early research has considered the stress-buffering effects of social support (as we have investigated in Chapters 7 and 8), the presence of (affect-based) social support has also been hypothesised to have additional, direct effects on physiological processes. This includes heart rate and blood pressure changes (cardiovascular reactivity) as well as immune system functionality and energy metabolism, all of which can be considered a physiological cost associated with stress.

Despite the significant correlations found in current literature, the direct relationship between affective social support and benefits on physiological responses has been difficult to investigate in the real world, and is crucially limited by several methodological challenges (Section 9.1.1). Understanding these effects further can provide insight into the additional mechanisms that tie social support to positive health-related outcomes in biological agents. Equally, for a model of adaptation through affect-based social relationships, understanding and abstracting these mechanisms of socially-supported physiological adaptation may provide significant advantages to model performance, as we have seen in previous chapters (Chapters 7 and 8), and underpin the principles of social allostatic adaptation.

Building on biological literature and our previous results, we propose and investigate one additional mechanism through which social support can provide positive, adaptive effects on physiological processes: through its mediation of the "cost" of stress (energy metabolism). We propose a potential pathway between oxytocin-mediated, affective social support and energy metabolism, via oxytocin's "buffering" effects on the stress system—abstractions of the sympathetic nervous system and HPA-axis, respectively—and hypothesise that this adaptive mechanism would provide significant survival-related benefits for agents with social bonds.

We use our simulated agent model to investigate this hypothesis and address methodological limitations from real-world studies. We conduct our experiments using the same experimental set up as Chapters 7 and 8—with three different world conditions of varying food availability, and three different affective bond combinations.

Our results from this investigation find that the adaptation of energy metabolism via oxytocin's adaptive effects on the stress system provided some significant advantages to the viability of agents with affective social bonds, compared to groups where this adaptation was not present. This improved performance was statistically-significant across all experiments in the most challenging physical environment, with mixed results of improvement noted in other world conditions.

Our results also find the emergence of an "anticipatory" feed-forward mechanism of oxytocin effects, where oxytocin secretion facilitates future oxytocin secretion via its promotion of prosocial behaviours and stress-regulation. This interplay between oxytocin, stress-reduction, behavioural and physiological adaptation mirrors biological hypotheses of these systems, providing further insight into the complex, recursive nature of hormonal mechanisms that underpin the formation and maintenance of affect-based relationships.

We also find the emergence of two affect-based behaviours which may provide insight into some of the behavioural adaptation that facilitates wellbeing of social individuals in real-world societies. Firstly, we find that bonded agents would sometimes regulate their energy metabolism by forming small "clusters" between themselves for short periods of time. We propose that this behaviour has some relation to that of "social thermoregulation" seen in biological societies, where social animals form small groups as a cost-effective way to regulate homeostasis.

Secondly, we find that stressed agents without social bonds consumed higher amounts of food to alleviate the stress associated with physiological deficits. After performing some post-hoc analysis, we propose that our results may give insight into the potential underlying mechanisms associated with the "stress-eating" phenomenon seen in the real world.

Overall, our model of adaptation through affect-based social support demonstrated improved performance across a number of dynamic environmental conditions, facilitated by anticipatory, adaptive regulation of homeostatically-controlled processes and social interactions. We consider these to be adaptive mechanisms of social allostasis. This chapter also presents our final iteration of our (affect-based) social allostasis-inspired model of adaptation for embodied agents, and we present it as a framework for future investigations. We summarise by proposing a number of future directions to investigate our findings further.

9.1.1 Affect-Based Adaptation of Energy Metabolism

As a hormone of adaptation, cortisol facilitates the mobilisation of energy [14], [29], [133]. This increase in energetic cost (energy metabolism), along with changes in cardiovascular function [246]–[248] can be seen as a physiological “cost” associated with stress and adaptation [249], [250]. While adaptive in the short-term, prolonged changes of such physiological function underpins the deleterious pathological consequences associated with prolonged stress—including dysregulated metabolic and cardiovascular function [14], [200], [251], [252]—which may threaten the long-term wellbeing of individuals.

However, it is not axiomatic that energy metabolism is directly correlated with stress levels themselves [253]. The modulation of internal receptors and systems may play a role in determining the energetic cost associated with stress. In biological entities, the HPA-axis and the autonomic nervous system [231] are found to play a role in the metabolic effects and deleterious consequences of elevated stress levels [200], [254]. These internal differences, loosely corresponding to the ability to cope with stress, may play a role in the range of physiological and behavioural responses associated with elevated cortisol levels [201], [255], [256].

In addition to adapting the release of cortisol (Chapter 7) and stress threshold (Chapter 8), social support may play a further stress-reducing role by mediating the physiological “cost” of stress.

As [247] have suggested in their conceptual model of social support and physical health (Figure 1 in [247]), alterations to metabolic and cardiovascular activity may be determined not only by individual stress (cortisol) levels, but also by the presence of social support. Other direct physiological effects, such as reduced cardiovascular activity [208], [246], [257], [258] and blood pressure [259] have been correlated with the presence of social support. These effects have been termed the “cardiovascular reactivity” hypothesis of social support [204], [260], [261].

Despite the correlations, current literature is yet to find a direct physiological mechanism linking social support and the observed benefits on underlying physiological measures [207], [262]. As [262] states, rather than direct physiological adaptation, it may be the case that the positive health consequences associated with social support are linked to its psychological stress-reducing, or affect-based effects. This suggestion has therefore made it difficult to accurately investigate the physiological and psychological health consequences of social support independently. The experimental challenges of separating the psychological and physiological effects of social support also includes the dissociation between (perceived or received) affect-based social support [262], psychological appraisals of a situation [263], methodological challenges regarding social or environmental cues [264] or temporal interference of physiological effects [265]. These challenges have resulted in limited studies accurately investigating the precise mechanisms linking affect-based social support and

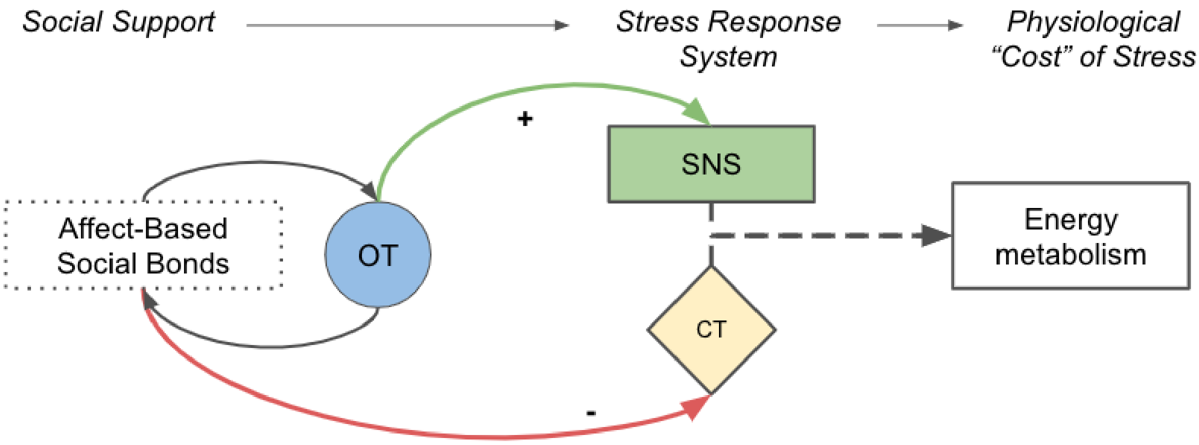

Figure 9.2 A high-level illustration of our hypothesised downstream effect linking affect-based social support and the metabolic consequences of stress (energy metabolism) as discussed in this section. Interactions with social support promotes oxytocin (OT) release, which provides a buffer to the activation of the sympathetic nervous system (SNS) and reduces cortisol (CT) secretion. A combined effect of this stress response then affects energy metabolism. We show a detailed implementation of these mechanisms in our model in Figure 9.4.

positive physiological outcomes: including the regulation of stress-induced energy metabolism [251], [262]. One potential mechanism that may link social support and the regulation of energy metabolism may be via oxytocinergic pathways. Independent of social support, treatment with oxytocin has been found to play a role in the regulation of energy metabolism [266] as well as a range of positive effects on physiological health in animals (including humans): including anti-inflammatory effects [233], [267], the preservation of energy expenditure [268], and the regulation of glucose homeostasis and rates of energy metabolism [269].

In addition, energy metabolic rates have also been found to be determined as a function of the total stress response (i.e. stress levels and autonomic nervous system activation) [200], [231], [251], [253]. These responses to stress—both the release and the tolerance to cortisol—are adapted via oxytocin (as discussed in Chapters 7 and 8), which is itself secreted as a function of social support and interaction.

Combining this literature, we suggest a potential mechanism through which social support can play a role in adapting the physiological effects of stress: that the "buffering" of the internal stress tolerance can also potentially dampen the physiological "cost" associated with elevated stress (cortisol) levels.

This potential pathway linking oxytocin-mediated buffering of the stress system and the dampening of physiological responses to stress may give insight into how social support can provide direct physiological benefits and promote long-term wellbeing, and may give insight into diverse, individual physiological responses associated with stress [201], [256].

The adaptation of energy metabolism can be considered an important type of adaptation for adaptive systems [270] and has been investigated in previous adaptive models of artificial systems: including physical agents [271] and simulated cellular systems [272], and minimally-cognitive agent models in the context of a two-resource [273] and three-resource [274] problem. However, as models of social allostasis are still limited in the literature, potential mechanisms underpinning social interactions and this adaptation of metabolism have not yet been investigated.

9.2 Aims & Hypothesis

Continuing our investigation into the stress-buffering effects of social support, we aim to investigate the role of social support on adapting one of the physiological costs associated with stress: energy metabolism. Addressing limitations from biological studies and building on our previous results, we propose a potential pathway linking social support and wellbeing benefits via the regulation of energy metabolism.

Specifically, we propose that the adaptation of energy metabolism—driven as a function of the stress system—is achieved through oxytocin's "buffering" effects of stress tolerance (the "sympathetic nervous system").

We study the effects of this proposed mechanism on the viability-related performance and social interactions of a small group of agents across several physically-challenging conditions using several different affect-based social bond combinations.

We hypothesise that this adaptation of energy metabolism will significantly benefit the viability of individuals with social support in our model. We also hypothesise that this adaptation of one physiological variable will result in an increase in socio-positive behaviours between agents with social bonds.

9.3 Extensions to the Agent Model

Building upon the two previous stress-reducing effects of social support (presented in Chapters 7 and 8), we present a single extension to our agent model to capture the hypothesised mechanism of adaptation of energy metabolism and expenditure as discussed in Section 9.1.1 (Figure 9.4). In this investigation, we define energy "metabolism" and "expenditure" as the rate at which energy is depleted, and we consider it an abstraction of one type of physiological response associated with stress.

As a reminder, we retain the relationship between an agent's movement speed and their rate of energy depletion: the *EnergeticCost*. In previous models, this rate of energy cost (*EnergeticCost*) was determined

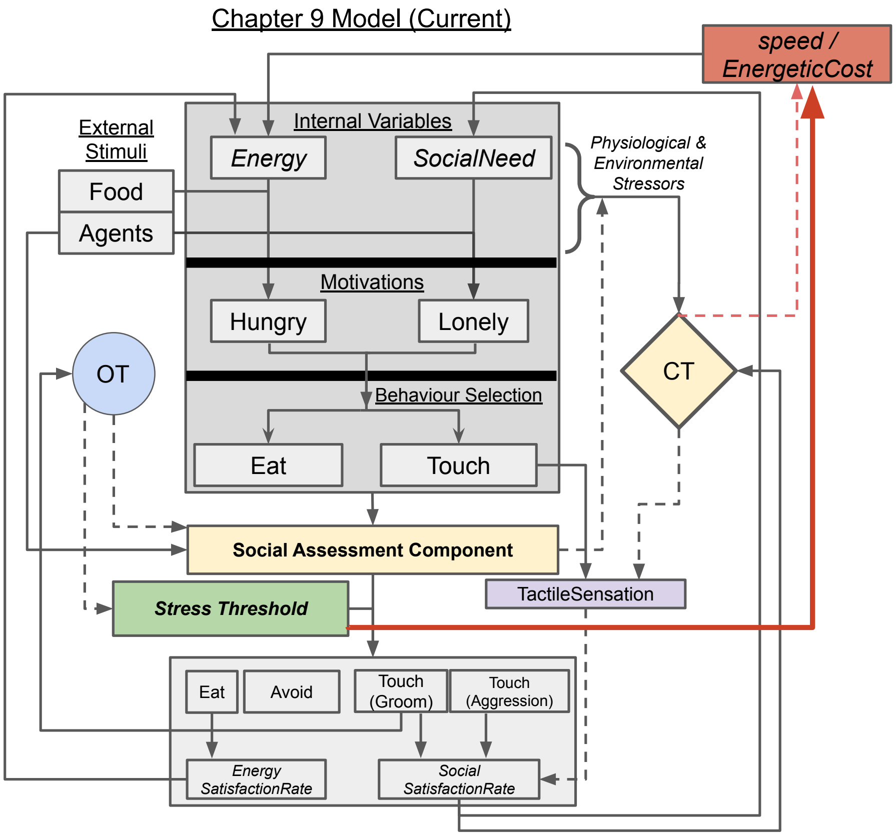

Figure 9.3 Illustration of the Action-Selection Architecture model used in this investigation. Extensions to the previous model are denoted in red. In this model, agent speed and energy metabolism (Energetic Cost) is driven through cortisol (CT) levels and buffered by the oxytocin-mediated *StressThreshold* (an abstraction of the sympathetic nervous system activation). This updated calculation can be seen in Equation 9.1. A more granular view of this mechanism can be seen in Figure 9.4.

as a direct function of internal cortisol (stress) levels (Equation 6.2.2). In this model, we account for agents' individual *Stress Threshold*, and determine the rate of its movement (and energy metabolism) as a function between absolute stress levels and individual tolerances to stress:

$$speed_t = speed_0 \times (1 + (\mathbf{(CT\ -\ StressThreshold)} \times c)) \tag{9.1}$$

Where speed_0 is the default speed (0.5) and speed_t is the movement speed at time step t and c is a constant to regulate values: where c is 0.75 if CT $<$ *StressTolerance*, and 1.2 if CT $>$ *StressTolerance*. We highlight changes to the previous function (Equation 6.2.2) in bold.

As with previous models, the *EnergeticCost* is modulated by the rate of this speed:

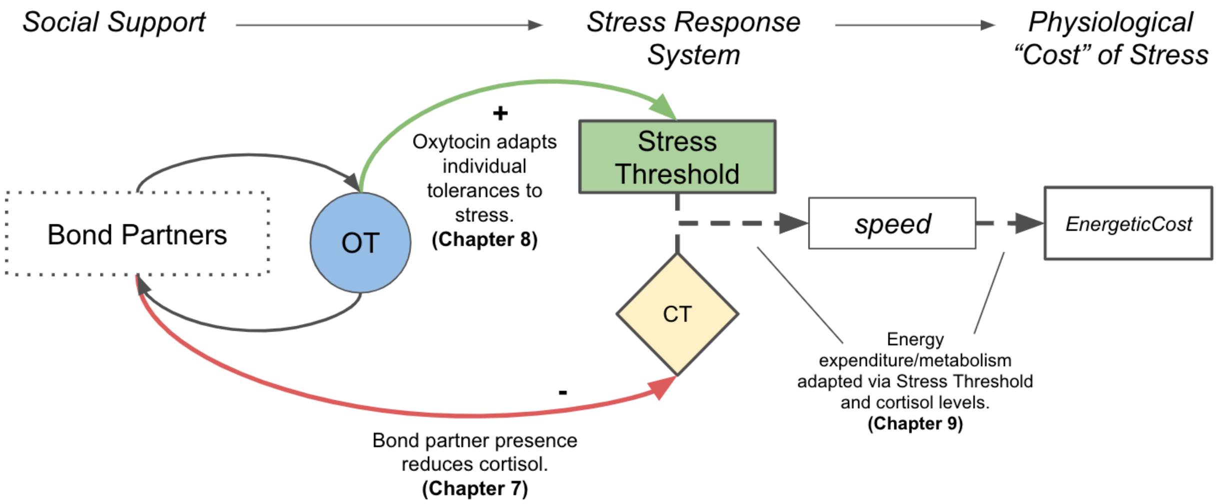

Figure 9.4 Illustration of the hypothesised adaptive effects between affect-based social support and the metabolic consequences of stress (Figure 9.2), as implemented in our agent model. This overview shows the mechanisms investigated in previous chapters (red and green lines) and the extension presented in this chapter (dashed lines from *Stress Threshold* and *CT* to *speed*, Equation 9.1). Each adaptive mechanism is labelled along with the chapter that it was introduced in. OT = Oxytocin. CT = Cortisol.

$$EnergeticCost_t = EnergyLossRate \times (2 \times speed_t) \tag{9.2}$$

where *EnergyLossRate* is 0.003, the default rate of energy loss at each time step. Therefore, *EnergeticCost* is driven by the rate at which an agent expends energy (its *speed*), which is determined by Equation 9.1. As with previous models, the rate at which *Energy* depletes is updated by the updated *EnergeticCost*:

$$Energy_t = Energy_{t-1} - EnergeticCost \tag{9.3}$$

As per Chapters 7 and 8, the *StressThreshold* value is initialised at 0.5. For agents with no affective bonds (and no oxytocin present within their physiology), this value remains fixed. For agents with affective social bonds and who are endowed with oxytocin mechanisms, the *StressThreshold* value is adapted between 0.25 and 0.75, modulated by oxytocin levels (Equation 8.1.1). All other effects from our model remain unchanged.

9.4 Experimental Set Up

As a direct extension to the investigation in Chapter 8, we retain the three experimental parameters and set up for our current investigation. We describe those parameters again here for clarity.

Time Step	**Start**	0	2001	3001	4001	5001	6001	7001	8001	9001	10001	11001	12001	13001	14001
	End	2000	3000	4000	5000	6000	7000	8000	9000	10000	11000	12000	13000	14000	15000
% of total simulation time at start		7%	14%	21%	28%	34%	41%	48%	54%	61%	68%	74%	81%	88%	94%
World Condition	**Phase**	**1A/1B**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
Static		4													
Seasonal		4	3	2	1	2	3	4	3	2	1	2	3	4	3
Extreme		4	1	4	1	4	1	4	1	4	1	4	1	4	1

Table 9.1 Availability of food across different phases in each of the world conditions. This is also illustrated in Figure 9.5. Green cells highlight periods of the best food availability (4). Yellow (3) and orange (2) highlight poorer periods of food availability, and red cells highlight periods of the worst food availability (1). The percentage describes the time elapsed at the start of each phase. Note that the initial period (1A/1B) elapses 2000 time steps. Though the simulation ran for 14 phases, results are reported for phases where agents remained viable (1-12).

Condition Label	**Bond Type**	**Description**	**Effects of Social Support**	**Chapter of Investigation**
FB	Fixed	3 agents bonded, 3 unbonded. Affective bond strength remain fixed values (DSI=2).	Reduce CT Release	Chapter 7 (Experiment 2)
VB	Variable	3 agents bonded, 3 unbonded. Affective bond strength changes over time.	Reduce CT Release	Chapter 7 (Experiment 3)
VB+	Variable	3 agents bonded, 3 unbonded. Affective bond strength changes over time.	Reduce CT Release OT Adapts Stress Tolerance	Chapter 8 (Current)
VB++	Variable	3 agents bonded, 3 unbonded. Affective bond strength changes over time.	Reduce CT Release OT Adapts Stress Tolerance + adapts energy metabolism.	Chapter 9 (Current)

Table 9.2 Table describing the different experimental groups discussed in our results. We highlight the group from this chapter in grey. Experimental results from FB and VB groups are reused from Chapter 7. Results for the VB+ group are reused from Chapter 8. OT = Oxytocin. CT = Cortisol. "Reduce CT Release" relates to the social buffering effects discussed in Section 7.2.3. "Adapt Stress Tolerance" relates to the modulation of the *StressThreshold* presented in chapter Section 8.2.1). "Adapts energy metabolism" relates to the adaptation of energetic cost, described in Section 9.3.

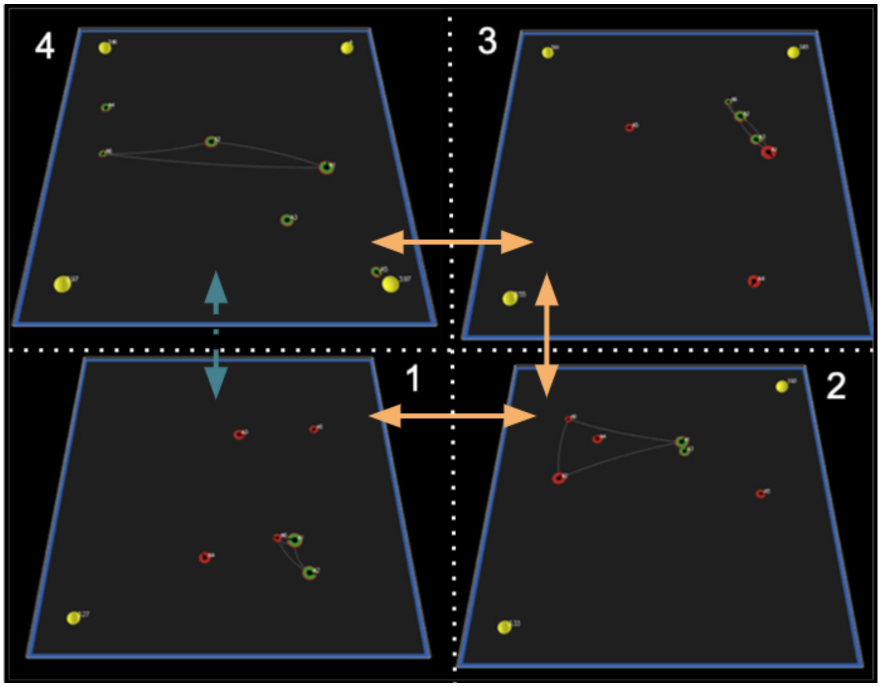

Figure 9.5 Screenshots of different food availability in each of the world conditions. *Static* environments remain fixed with four food resources (yellow discs, top left). The orange arrow denotes changes in the *Seasonal* condition. The blue arrow denotes changes in the *Extreme* condition. Details of each food phase can be seen in Table 7.2. Agents are rendered as doughnut-type shapes. Colour indicates their "stressed" state (green = not stressed, red = stressed). A dashed line between agents denotes an existing affective bond between them.

The first experimental parameter is the affective bond type. In previous investigations, we have investigated using "Fixed" and "Variable" bond types, with respect to the dynamic strength of the affective bond between agents (Chapters 7 and 8), respectively. Here, we only model the "Variable" bond type from Chapter 8 (the VB+ group).

Bond strength (DSI) is initialised at its maximum value of 2, but is susceptible to a small decay rate at each time step ($\mu = 0.9997$) and subject to change as a function of (pro and anti-)social interaction from bond partners (Equation 7.7).

The second experimental parameter relates to the affective bond partner combination. As with Chapter 8, we investigated three different bond partner combinations related to agent rank. Bond Condition A consists of A1-A2-A6 (two most-dominants and most-subordinate agents) with an affective bond between them, Bond Condition B consists of A3-A4-A5 (middle-ranked agents), and Bond Condition C consists of A4-A5-A6 (lowest-ranked agents). These combinations can be seen in Figure 9.6.

Finally, experiments were performed using the same three world conditions related to different food

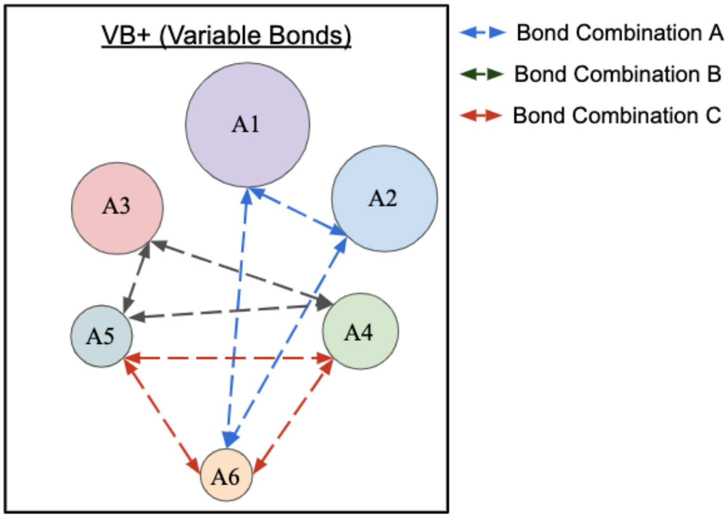

Figure 9.6 Illustrations of the three different affective bond combinations used in this experiment. These bonded agent combinations remain the same as those investigated in Chapters 7 and 8. Dashed lines illustrate variable bond strength (DSI initialised at its maximum value of 2 and changes over time). Agent number denotes rank, from A1 (highest-ranked) to A6 (lowest-ranked).

challenges as in Chapters 7 and 8. These relate to one static world (*Static*) of a fixed amount of food resources (4) throughout the simulation, and two dynamic environments. The first dynamic world, *Seasonal* (Figure 9.5, orange arrows), has steady changes in food resources every 1000 time steps (Table 9.1). Our second dynamic world, *Extreme* (Figure 9.5, blue arrows), has more extreme changes in food resources, changing from very good to very poor food availability every 1000 time steps. Figure 9.5 shows screenshots of these world conditions, and Table 9.1 describes the food availability across all phases for all three world conditions.

We performed 20 simulation runs for each of the 9 experimental combinations (Table 9.2, totalling 180 simulation runs. Each experiment was run for 15,000 time steps, resulting in approximately 12 hours of run time. For the VB and VB+ groups, we reuse the data from previous investigations during comparative analysis of results. We refer to results from this experiment as "VB++" groups.

We report results using the three viability-related metrics (*Life Length, Average Comfort, Physiological Wellbeing*) as described in Section (2.4.1. We recorded internal hormone (cortisol and oxytocin) levels at each time step, as well as both the mean and trended *StressThreshold* value for each agent. We report the rates and distribution of social behaviours (*Grooming* and *Aggression*) at an individual and group level. In line with preceding chapters, we further complement these results with qualitative analysis of simulation runs.

Much like Chapter 8, we compare our results from this investigation to the results of previous chapters. In our comparative analysis, we include the results from the FB and VB groups from Chapter 7 and the VB+ group from Chapter 8. Results from this investigation are denoted as the VB++ group. Table 9.2 describes these different experimental groups.

Statistical significance was performed between-groups using one-way ANOVA testing, with significance declared at p = 0.05. Correlation strength between metrics was calculated using Pearson's Correlation Coefficient.

9.5 Results

9.5.1 Viability Indicators

At an aggregated group level, we find small improvements across all viability indicators, across all world conditions, when compared to previous groups (VB and VB+). Mean *Life Length* for all agents saw improvements across all conditions when compared to the VB+ groups: 66% (+11% vs. VB+) in *Static* environments, 59% (+17% vs. VB+) in *Seasonal* environments and 57% (+15% vs. VB+) in *Extreme* environments.

Average Comfort saw improvements of +4% vs. VB+ groups in the *Static* environments, +2% in the *Seasonal* environments, and 6% in the *Extreme* environments

In terms of *Physiological Wellbeing*, we found no improvements in aggregated performance in *Static* world conditions, a +4% improvement in *Seasonal* world conditions, and an 11% improvement in *Extreme* world conditions when compared to the VB+ group.

For bonded agents, we see statistically-significant improvements in *Life Length* between VB++ and VB groups across all conditions (Figure 9.8). Agents in Bond Condition A reported a 77±4%, 68±3% and 66±3% survival rate for in *Static*, *Seasonal*, and *Extreme* conditions; Bond Condition B reported 69±3%, 65±4% and 63±3%; and Bond Condition C reported 67±4%, 60±3% and 54±4% for the three environments respectively.

These results were statistically-significant compared to the VB group (p values ranging from 0.01–0.023), but only statistically-significant in three conditions compared to the VB+ group (Bond Condition A, *Seasonal*, p = 0.02; Bond Condition B, *Seasonal*, p = 0.01; Bond Condition C, *Static*, p = 0.04). Nevertheless, we note that the consistent improvement across all conditions warrants further investigation and statistical analysis, and may have been limited by the relative challenges of our environments. We discuss this further in the next section.

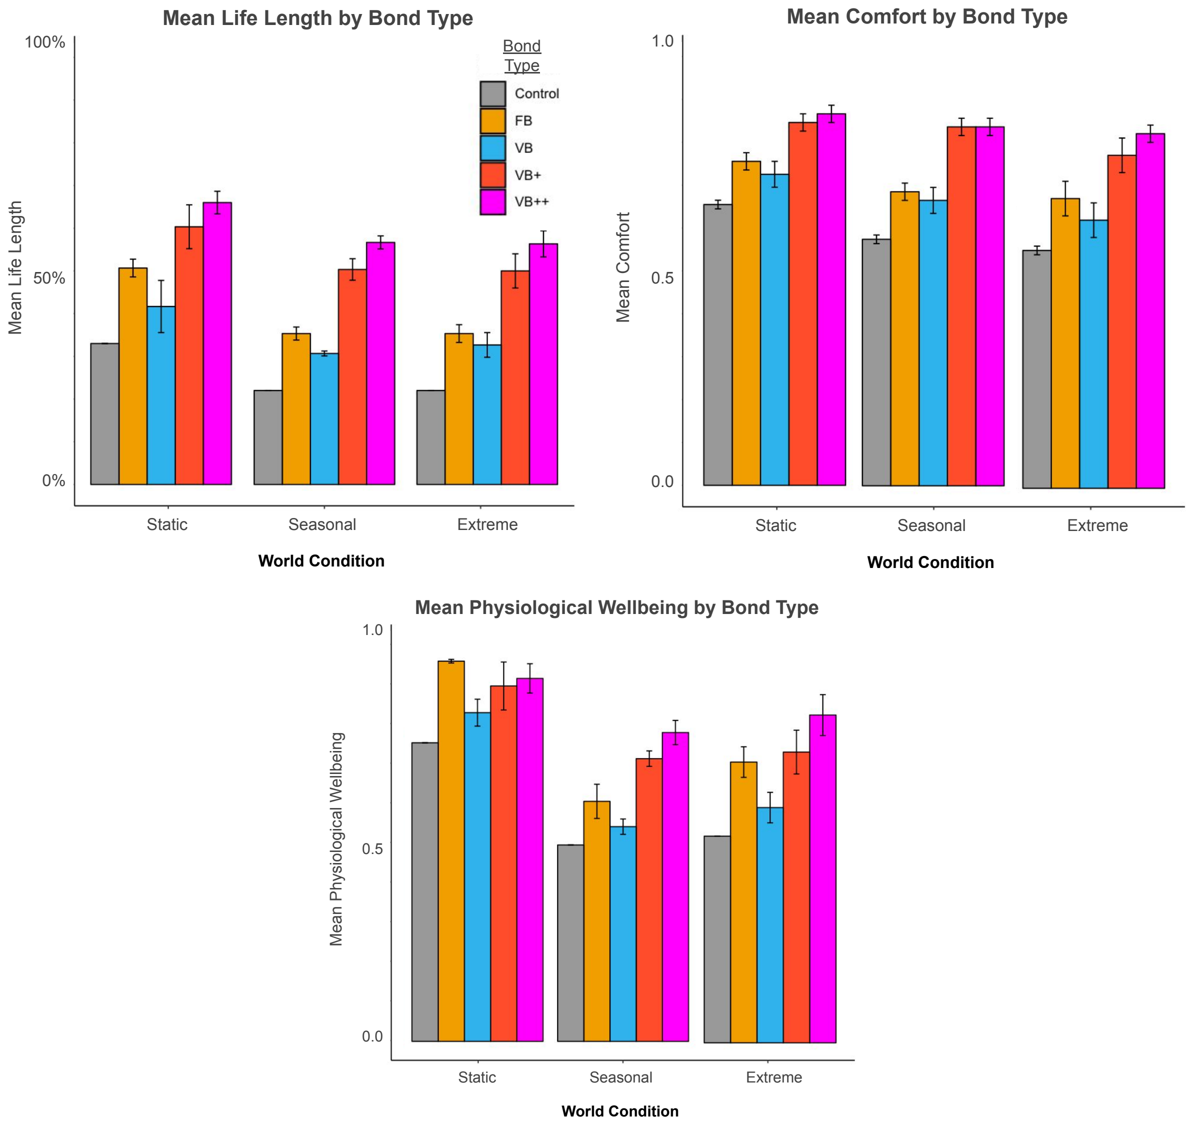

Figure 9.7 Mean values of all three Viability Indicator metrics, aggregated across all simulation runs. Results from Control, FB and VB groups are taken from Chapter 7. VB+ results are taken from Chapter 8. VB++ results are new results in this chapter. Control: Conditions with all agents unbonded (Experiment 1, Chapter 7). FB: Conditions of Fixed Bonds (Experiment 2, Chapter 7). VB: Conditions of Variable Bonds (Experiment 3, Chapter 7). VB+: Conditions of Variable Bonds with Stress Tolerance adaptation (Chapter 8). VB++: As VB+, with the adaptation of energy expenditure (current chapter). Error bars show standard error of the mean (SEM) for each group.

Average Comfort for bonded agents also saw improvements across all-but-one experimental condition (Table 9.3). In *Static* world conditions, bonded agents experienced an *Average Comfort* improvement of 1%–6% (with no change seen in Bond Condition B), 1%–6% improvement in *Seasonal* world conditions, and

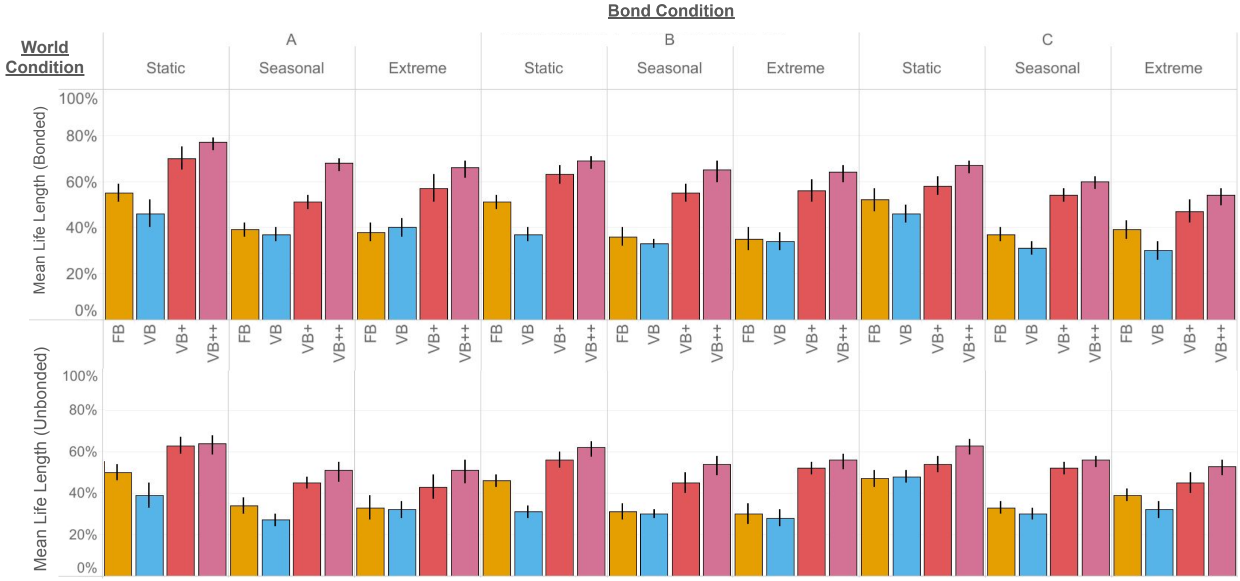

Figure 9.8 Mean Life Length values for bonded agents (top) and unbonded agents (bottom) across all world conditions. Error bars denote standard error of the mean (SEM) for each group. FB: Fixed Bonds (Experiment 2, Chapter 7). VB: Variable Bonds (Experiment 3, Chapter 7). VB+ Variable Bonds with Stress Tolerance adaptation (Chapter 8. VB++: As VB+, with the adaptation of energy expenditure (current chapter). Error bars show standard error of the mean (SEM) for each group.

5%–11% in *Extreme* world conditions, compared to VB+ groups. All results were statistically-significant (p < 0.05) compared to the VB groups, but only significant differences in *Extreme* world conditions and one experiment in the *Static* environment vs. VB+ groups.

We observed small improvements in the *Physiological Wellbeing* of bonded agents vs. VB+ group (Table 9.4). Much like *Average Comfort* trends, these differences became larger over increasing world challenges: agents in *Static* world conditions saw small changes of −1%–2% vs. VB+ groups; 1%–5% in *Seasonal* conditions, and 3%–6% in *Extreme* world conditions (Table 9.4). All results were statistically-significant improvements vs. VB groups (p < 0.05), but did not reach uniform significance vs. VB+ groups (p = 0.03–0.29).

In terms of unbonded agents, we saw improved *Life Lengths* for those agents vs. VB and VB++ groups. Unbonded agents in Bond Condition A experienced *Life Lengths* between 51±3%–64±3%; 54±5%–62±4% in Bond Condition B and 53±4%–63±5% in Bond Condition C. Differences vs. VB groups were statistically-significant across all conditions (p = 0.01–0.035). Comparisons with VB+ groups were non-significant, with p values between 0.03–0.34, and Bond Condition A in *Static* world conditions resulting in p = 0.74 (Figure 9.8, bottom row, left-most result) .

Average Comfort levels for unbonded agents saw small improvements, with two conditions (Bond Condi-

Bond Condition	World Condition	Comfort (Bonded)	vs. VB	vs. VB+	Comfort (Unbonded)	vs. VB	vs. VB+
A	Static	0.86	19%	1%	0.83	19%	1%
B	Static	0.86	19%	0%	0.82	17%	-1%
C	Static	0.84	17%	6%*	0.80	14%	-1%
A	Seasonal	0.85	29%	2%	0.77	22%	4%
B	Seasonal	0.84	27%	1%	0.75	19%	4%
C	Seasonal	0.84	27%	4%*	0.80	27%	1%
A	Extreme	0.82	32%	11%*	0.73	18%	6%*
B	Extreme	0.81	31%	5%*	0.74	19%	1%
C	Extreme	0.77	24%	10%*	0.73	18%	7%*

Table 9.3 Table describing the *Average Comfort* results, broken down by bonded and unbonded agents, across all world conditions. Results are compared to these same conditions in the VB group (Chapter 7, Table 7.6) and VB+ group (Chapter 8, Table 8.2). Grey cells highlight differences to VB+ groups. * denotes statistical significance ($p < 0.05$.).

tion B and C, in *Static* world conditions), which experienced a small (non-significant) detriment to *Average Comfort* (Table 9.3, right column). Much like the results for bonded agents, these changes were non-significant across most conditions, but did achieve significance ($p < 0.05$) in two out of the three bond combinations (A and C) in *Extreme* world conditions. It is worth noting that changes in *Seasonal* conditions for Bond Condition A ($p = 0.055$) and B (p=0.063) were close to reaching significance.

Physiological Wellbeing for unbonded agents followed similar trends to bonded agents across the *Static* (−2%–1%) and *Seasonal* (4%–6%) world conditions, with larger changes in the *Extreme* conditions (11%–16%). We note mixed results from the statistical-significance testing compared to the VB+ group, with significant changes only being observed across all bond conditions in the *Extreme* environment, and Bond Condition A in the *Seasonal* environment.

In summary, viability of agents—both bonded and unbonded—saw small improvements across all bond and world conditions vs. previous groups. These improvements were most notable in the two dynamic environments of limited food availability (*Seasonal* and *Extreme*). While these were significant improvements compared to the VB group (Chapter 7), statistical results were mixed when compared to the VB+ group (Chapter 8). Though the magnitude of changes may appear to be small in these environments, this often resulted in agents surviving through an additional "phase" of food availability and may have been limited by the challenge of the environment. We discuss this more in Section 9.6.

Bond Condition	World Condition	Physiological Wellbeing (Bonded)	vs. VB	vs. VB+	Physiological Wellbeing (Unbonded)	vs. VB	vs. VB+
A	Static	0.9	11%	1%	0.84	10%	1%
B	Static	0.96	13%	-1%	0.92	22%	-1%
C	Static	0.89	20%	2%	0.96	-1%	-2%
A	Seasonal	0.75	42%	4%	0.70	70%	6%*
B	Seasonal	0.78	76%	1%	0.72	49%	4%
C	Seasonal	0.77	96%	5%*	0.73	19%	5%*
A	Extreme	0.76	28%	6%*	0.73	68%	11%*
B	Extreme	0.85	66%	3%	0.83	52%	11%*
C	Extreme	0.75	64%	3%	0.85	20%	16%*

Table 9.4 Table describing the *Physiological Wellbeing* results, broken down by bonded and unbonded agents, across all world conditions. Results are compared to these same conditions in the VB group (Chapter 7, Table 7.9) and VB+ group (Chapter 8, Table 8.2). Grey cells highlight differences to VB+ groups. * denotes statistical significance ($p < 0.05$).

9.5.2 Hormonal Levels

Across all conditions, we found lower mean cortisol levels for both bonded and unbonded agents vs. VB groups, with comparative levels to the VB+ group (Figure 9.9). Changes in cortisol levels were statistically-significant vs. VB groups ($p < 0.05$ for all conditions) except for two conditions: bonded agents in Bond Condition B, *Extreme* conditions; and unbonded agents in Bond Condition C, *Extreme* conditions. Results were non-significant across the majority of comparisons with the VB+ group. We note any statistically-significant results with an asterisk in Figure 9.9. Overall, societies experienced comparable to slightly lower overall stress vs. previous experiments, and this reduced stress was primarily seen in agents with affective social bonds.

Along with reduced stress levels, we found increased oxytocin levels for bonded agents across all-but-one condition (Bond Condition B, in the *Seasonal* environment), vs. VB+ groups from Chapter 8 (Table 9.5).

Bond Condition A saw 6%–40% increases in oxytocin, with smaller changes (−1%–7%) seen in Bond Condition B and Bond Condition C (2%–14%). Although the magnitude of change in oxytocin levels was smaller than the change between VB and VB+ groups (Table 8.10), these were nevertheless statistically-significant differences in most conditions (marked with an asterisk in Table 9.5). For all bond conditions, the largest increases in oxytocin levels were seen in *Seasonal* world conditions: +40% for Bond Condition A, +6% for Bond Condition B, and +14% for Bond Condition C.

We reported a moderate-to-strong correlation between mean oxytocin values for bonded agents and *Life*

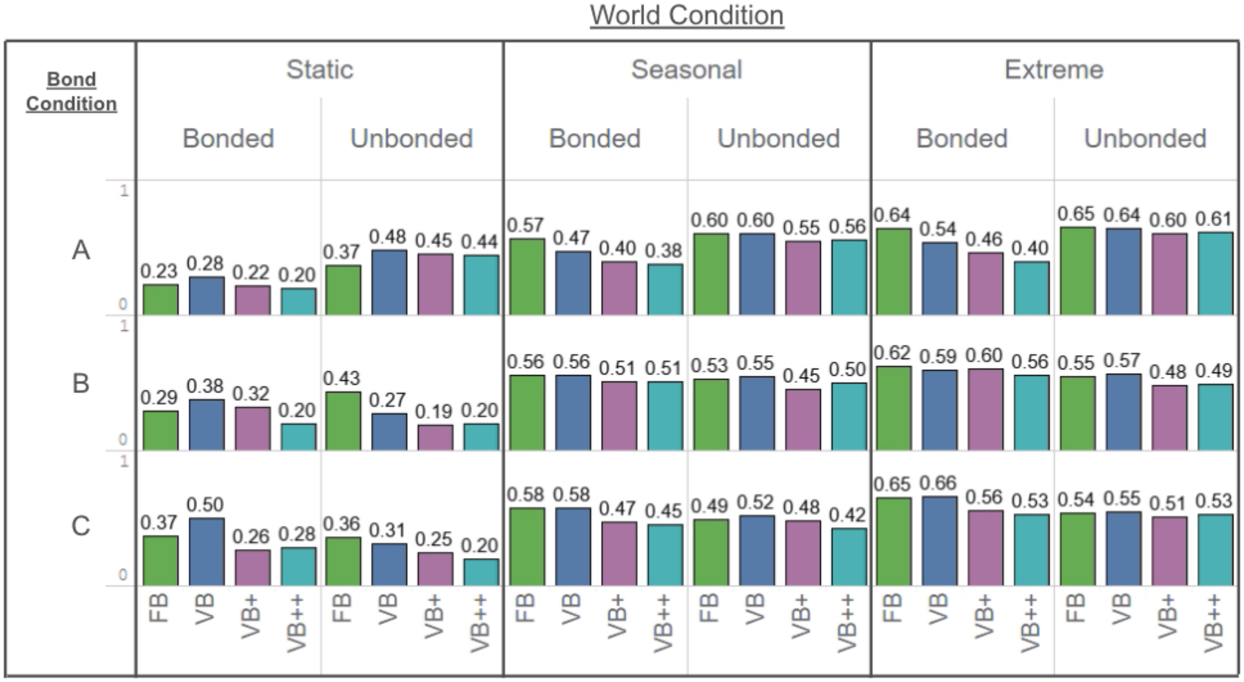

Figure 9.9 Comparison of the mean cortisol levels for bonded and unbonded agents across all bond types, aggregated across all simulation runs. FB: Fixed Bonds (Experiment 2, Chapter 7. VB: Variable Bonds group (Experiment 3, Chapter 7). VB+ Variable Bonds with tolerance adaptation (Chapter 8). VB++: As VB+, As VB+, and the adaptation of energy metabolism (current chapter). Cortisol levels exist in the range 0 to 1 and are only recorded while agents are still alive (*Energy* > 0).

Length for bonded agents across all experimental conditions. These correlations were stronger as the world challenge increased (from *Static* ($r^2 = 0.49$–0.56), *Seasonal* ($r^2 = 0.70$–0.84) and *Extreme* ($r^2 = 0.87$–0.89)). We report the results of our statistical test in Appendix H.

In sum, the mean stress (cortisol) level were comparable to VB+ groups (with some decreases observed in the VB++ group) though they were not statistically-significant. However, we found statistically-significant increases in mean oxytocin levels for bonded agents (Table 9.5), compared to both VB and VB++ groups.

9.5.3 Social Interactions

For all bond conditions, we observed higher rates of Intra-Bond *Grooming* compared to VB+ groups (Table 9.6). For Bond Condition A, Intra-Bond *Grooming* saw increases of 4%–12%; Bond Condition B reported a 7%–9% increase, and Bond Condition C saw a 6%–13% increase.

Intra-Bond *Aggression* rates were also lower across all bond combinations and world environments, apart from one (Bond Condition A, *Seasonal* world condition, +4%). Absolute numbers of *Aggression* (Table 9.6)

Bond Condition	**World Condition**	**VB**	**VB+**	**VB++**	**Δ**
A	Static	0.49	0.64	0.68	+6%
	Seasonal	0.37	0.40	0.56	+40%
	Extreme	0.31	0.55	0.60	+9%
B	Static	0.51	0.74	0.79	+7%
	Seasonal	0.44	0.69	0.68	-1%
	Extreme	0.45	0.67	0.71	+6%
C	Static	0.47	0.74	0.78	+5%
	Seasonal	0.46	0.65	0.74	+14%
	Extreme	0.42	0.65	0.66	+2%

Table 9.5 Mean oxytocin levels for bonded agents across all conditions, aggregated across all simulation runs. VB: Variable Bond groups with no adaptation of the *Stress Threshold* (Experiment 3, Chapter 7). VB+: Variable Bond groups with adaptation of *Stress Threshold* (Chapter 8). VB++: As VB+, and the adaptation of energy metabolism (current chapter). Δ = percentage difference between VB++ and VB+ oxytocin levels. Red values indicates a negative Δ. Green values indicate positive Δ.

Bond Condition	**World Condition**	**Intra-Bond Grooming %**	**vs. VB+**	**Intra-Bond Grooming #**	**Intra-Bond Aggression %**	**vs. VB+**	**Intra-Bond Aggression #**
A	Static	71%	4%	733	17%	-32%	10
A	Seasonal	68%	7%	508	28%	4%	11
A	Extreme	73%	12%	452	22%	-4%	8
B	Static	58%	7%	1,478	72%	-5%	31
B	Seasonal	49%	9%	792	50%	-4%	24
B	Extreme	51%	9%	988	65%	-11%	27
C	Static	26%	13%	1,226	100%	0%	23
C	Seasonal	30%	7%	638	100%	0%	23
C	Extreme	37%	6%	554	100%	0%	29

Table 9.6 Table showing the total amount of Intra-Bond *Grooming* and *Aggression* across each of the Bond Conditions and World Conditions in this investigation. Intra-Bond Grooming/Aggression values describes what percentage of total grooming/aggression performed by bonded agents was performed towards their affective bond partners. Notable differences to results of VB+ groups (Chapter 8, Table 8.3) are highlighted in grey.

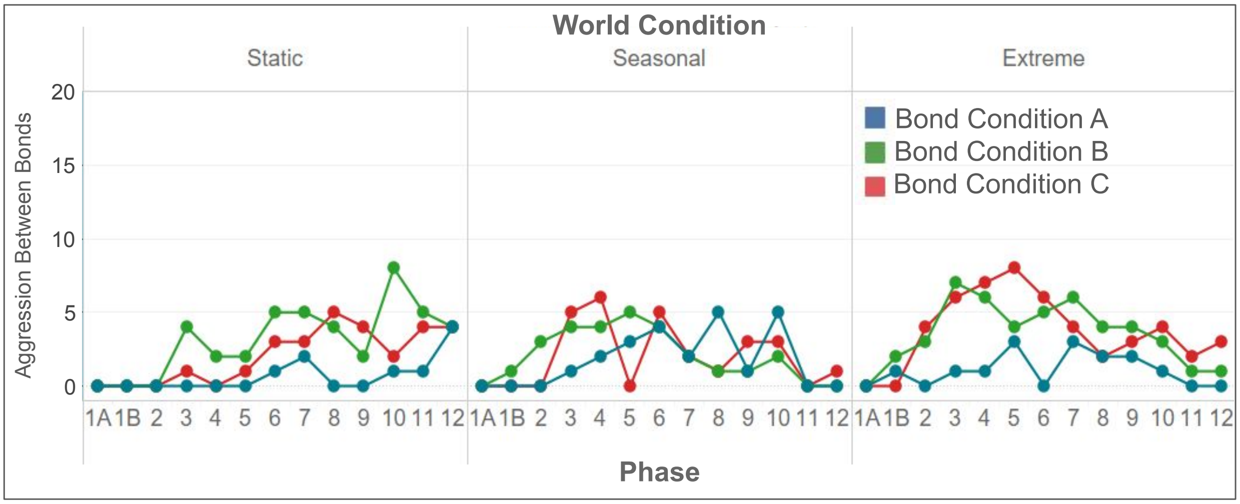

Figure 9.10 Total amount of intra-bond *Aggression* (y-axis) (Aggression performed between the three bonded agents) during each phase of food availability (x-axis), aggregated across all simulation runs. Bond Condition A: A1-A2-A6 bonded. Bond Condition B: A3-A4-A5 bonded. Bond Condition C: A4-A5-A6 bonded. Further data can be found in Table 9.6.

were also lower compared to the VB+ groups, with 186 total Intra-Bond *Aggression* encounters recorded for VB++ groups and 296 recorded for VB+ groups (Table 8.3).

In sum, the results from bonded agents saw a shift in the types of social behaviours being performed: skewing towards higher rates of socio-positive *Grooming* between bonded agents, with less *Aggression* (socio-negative) behaviours between these agents vs. previous groups.

Figure 9.10 shows the trended *Aggression* interactions between bonded agents across each of the different phases (Table 9.1). Much like the results from the VB+ condition (Figure 8.8), mean amounts of intra-bond *Aggression* remained similar between all world conditions, despite the increases in relative physical challenge (from *Static*, *Seasonal* and *Extreme* world conditions respectively).

Despite these changes in *Aggression* during each of the phases, we reiterate that the absolute number of *Aggression* encounters between bonded agents was significantly lower in this investigation (186) compared to previous groups (FB: 594, VB: 810, VB+: 296).

9.5.4 Stress Threshold Trends

Figure 9.11 shows the mean trended values of the internal *Stress Threshold* for bonded agents across all conditions. As a reminder, Table 9.1 describes the different amount of food availability during each of the phases in each of the world conditions.

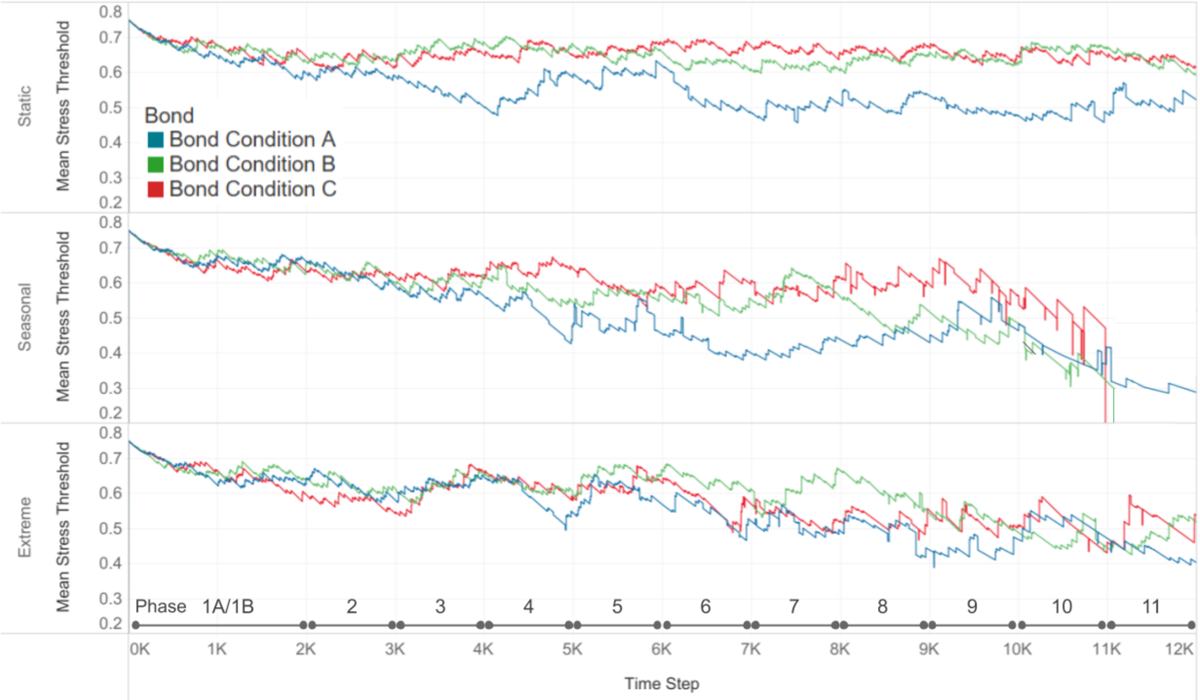

Figure 9.11 Mean trends of internal Stress Thresholds for bonded agents. Results are aggregated across all experimental runs and for the three bonded agents in each condition (Bond Condition A: A1-A2-A6. Bond Condition B: A3-A4-A5. Bond Condition C: A4-A5-A6. Though simulations were run for 15,000 time steps, results after 12,000 were not significant and not included. Table 9.1 describes the food availability during each of the food phases.

In *Static* world conditions, agents in Bond Conditions B and C adapted and maintained their *Stress Threshold* to similar values (0.65) throughout all experimental runs. Like the results from Chapter 8 (Figure 8.4.4), agents in Bond Condition A adapted their *Stress Threshold* value to a slightly lower value (0.55) compared to the other two bonded conditions. For all bond combinations, these values were significantly higher compared to VB+ groups: +32% ($p < 0.01$) for agents in Bond Condition A, +16% ($p < 0.01$) for Bond Condition B, and +21% ($p < 0.01$) in Bond Condition C.

In *Seasonal* world conditions, we found similar values for both Bond Conditions B and C throughout all experiments (0.58 and 0.62 respectively), with a lower adapted threshold for Bond Condition A (0.52). Bond Condition A maintained a similar *Stress Threshold* value (0.61) compared to the other two groups in earlier phases (until Time Step = 5K). However, after Phase 4 (4000-5000 time steps)—where one food resource is available—the mean *Stress Threshold* for this group reduced to 0.43. This change in *Stress Threshold* adaptation was not observed for the other two bond groups, suggesting that this group adapted differently compared to other bond groups. These values were significant changes compared to the VB+ group: +11%

Bond Condition	**World Condition**	**A1**	**A2**	**A3**	**A4**	**A5**	**A6**	**Mean of Bonded Agents**
A	Static	0.61	0.54	0.50	0.50	0.50	0.52	**0.55**
B	Static	0.50	0.50	0.67	0.66	0.64	0.50	**0.65**
C	Static	0.50	0.50	0.50	0.64	0.67	0.68	**0.65**
A	Seasonal	0.70	0.55	0.50	0.50	0.50	0.48	**0.52**
B	Seasonal	0.50	0.50	0.60	0.58	0.56	0.50	**0.57**
C	Seasonal	0.50	0.50	0.50	0.60	0.59	0.63	**0.61**
A	Extreme	0.72	0.57	0.50	0.50	0.50	0.52	**0.57**
B	Extreme	0.50	0.50	0.63	0.59	0.61	0.50	**0.60**
C	Extreme	0.50	0.50	0.50	0.58	0.57	0.60	**0.59**

Table 9.7 Mean values of the *Stress Threshold* for each agent, aggregated across all simulation runs. When agents were not bonded, their default *Stress Threshold* was 0.5. Grey cells indicate agents that shared an affective bond in those conditions.

($p < 0.01$) for Bond Condition A, +41% ($p < 0.01$) for Bond Condition B, and +51% ($p < 0.01$) for Bond Condition C.

In *Extreme* world conditions, *Stress Threshold* trends were similar for all bond conditions, with Bond Condition A, B and C reporting mean values of 0.56 (+57% vs. VB), 0.60 (+46%), and 0.59 (+44%) respectively. In line with the other two world conditions, these were statistically-significant changes ($p < 0.01$ for all effects) compared to the VB+ group.

Table 9.7 shows the mean *Stress Threshold* values for individual agents across all conditions. With one exception (Bond Condition A, *Seasonal* world condition, agent A6), bonded agents across all conditions adapted their *Stress Threshold* values to above the default value (of 0.50). Particularly in Bond Condition A, we saw how, unlike in the VB+ group (Table 8.6), the higher-ranked A1 and A2 adapted their *Stress Thresholds* to higher values than their bonded A6. We note how accounting for an additional effect of social support on physiological adaptation (the adaptation of energy metabolism) resulted in a different type of adaptation of this value for these higher-ranking agents: adapting to a higher tolerance to stress compared to the previous investigation (Table 8.6).

Overall, compared to the previous investigation, we find that agents adapted their *Stress Threshold* values to significantly higher values—resulting in an increased tolerance to stress—which may have contributed to some improvements in agent performance seen in our investigation. We posit that this higher rate of stress tolerance was facilitated, in part, by an overall improvement in oxytocin levels.

9.6 Discussion

Overall, we observed small improvements across all three viability-related metrics in groups who adapted their energy metabolism (VB++, present chapter), when compared to groups who were not endowed with this adaptation (VB+, Chapter 8). Despite these consistent increases, the results from our statistical tests reported a range of both significant and non-significant results. Non-significant differences were seen largely in the *Static* world environments, whereas statistically-significant differences were seen across our most challenging dynamic environment related to food availability (*Extreme*). Some significant results were also found in the second-most challenging environment (*Seasonal* world conditions).

We therefore found some contextual support for our hypothesis—that affect-based adaptation of energy metabolism would promote the viability of bonded agents—across dynamic, challenging conditions. Here, we posit that these adaptive mechanisms were not necessary to further promote viability in comparatively "easier" environments (*Static*), but did provide adaptive advantages in more challenging environments (*Seasonal, Extreme*). We propose that performing additional simulation runs could strengthen the statistical confidence across these conditions.

Though the magnitude of change in viability indicators (in particular, *Life Length*) was smaller compared to the VB+ groups (Figure 9.8), these differences nevertheless had practical implications within our world conditions. Across several conditions, small improvements in *Life Length* resulted in agents surviving through an additional "phase" of food availability in challenging, dynamic environments (Table 9.1). We discuss this in more detail below.

Despite the fact that our model did not make any direct changes to hormonal secretion or affective bond strength, we find significant increases to both oxytocin levels and aggregated bond strength between bonded agents as a result of the adaptation of energy metabolism that we had included in our model. We posit that these emergent effects relate to the positive feed-forward mechanisms of oxytocin on the maintenance of affect-based relationships and affective states, also proposed by [159], [175], [275], and suggest that our findings can be used to support this hypothesised phenomena in biological systems.

From a qualitative perspective, we see the emergence of behaviour akin to the "social thermoregulation" hypothesis of social support's effects on physiological adaptation. Social thermoregulation describes how body temperature, and in turn the energetic cost associated with maintaining homeostasis, can be regulated through close proximity with social partners. Although we do not claim that this was alone in driving increases in performance, we highlight the emergence of this behaviour and consider its plausibility as a mechanism for how affect-based social support has physiological benefits in biological agents. We discuss this more in detail below.

In addition, we observed behaviours mirroring “stress-eating” for “stressed” agents without social bonds in our model. After performing post-hoc analysis, we found quantitative support for our observations and highlighted this as an area for further investigation. These observations mirror those from the natural world, describing how social support can also provide advantages to the management of homeostasis through indirect health benefits driven by behavioural adaptation. We discuss this further below.

9.6.1 Small Viability Improvements Results in Additional Phase Survival

As we had reported (Section 9.5), we saw viability improvements across all groups with adaptation of energy metabolism (VB++ groups) compared to previous groups (VB+), though these results did not achieve significance to the extent that previous investigations did. Despite these mixed statistical results for viability performance (specifically, for *Life Length*) compared to previous groups (Figure 9.8), we take into account the practical implications of these improvements in our current world conditions.

Each phase of food availability seen in Table 9.1 corresponds to approximately 7% of the total simulation run time. Therefore, a 7% improvement in *Life Length* corresponds to agents surviving through an additional phase of food availability, in each of the (*Seasonal* and *Extreme*) world conditions.

For instance, in Bond Condition B, in the *Seasonal* world condition, we see a statistically-significant improvement vs. previous groups for Bond Condition B (VB+: 55±4%, VB++: 65±4%) and Bond Condition C (VB+: 56±3%, VB++: 60±2%). These results corresponded to bonded agents surviving through an extra phase of challenging food availability. A similar performance was seen in *Extreme* conditions for Bond Condition B (VB+: 56±5%, VB++: 64±3%).

In terms of survival in these environments, these results were the difference between agents either succumbing to, or surviving through, a period of poor food availability (phase 8 in *Extreme* world conditions, Table 9.1).

Overall, despite the fact that these effect sizes were small compared to previous experiments, the implications of this increase in the world conditions that we had investigated were still a notable result. We propose that these small effect sizes may be due to the relative challenges of our current world conditions. To further investigate the long-term adaptive effects of these mechanisms and to correct for any limitations presented by

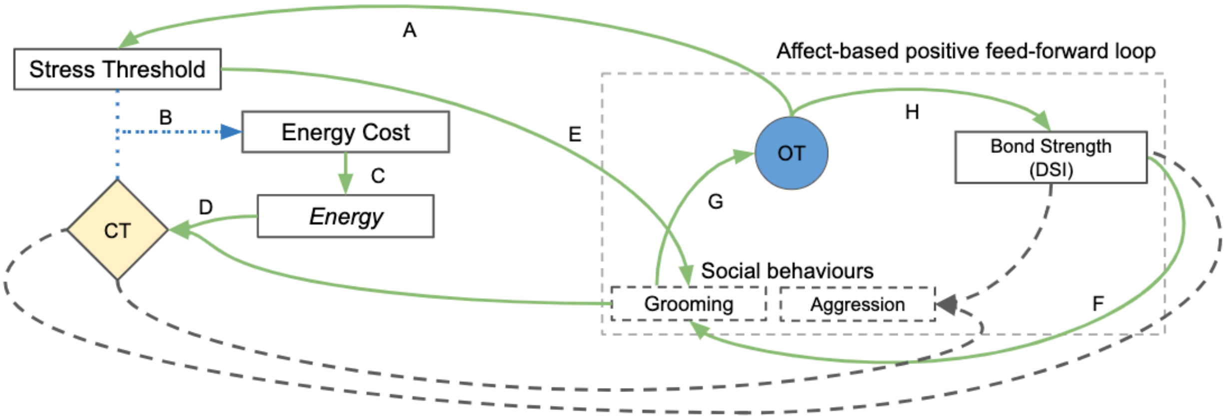

Figure 9.12 Illustration of the underlying mechanisms resulting in the anticipatory adaptation of physiology. The dotted blue line indicates the change made in this current model. Green arrow indicate a positive (modulatory) effect in that direction. Dotted black lines indicate a reduced likelihood of that physiological/behavioural loop. (A) High levels of oxytocin buffers stress tolerance (the *StressThreshold*). (B) Energy cost is adapted as a function of current cortisol levels and *StressThreshold* value. (C) Reduced energetic cost reduces the rate of *Energy* loss, which (D) reduces the rate of cortisol release. This then further adapts energetic cost (B). (E) Increased tolerance to stress reduces the likelihood of agents being stressed, and prioritises *Grooming* with bond partners (F). (G) *Grooming* releases further OT, which further buffers the *StressThreshold* (A) and further strengthens bond strength with social bond partners (H). This interplay of physiological and behavioural adaptation results in an affect-based feedback loop (A→E→G→H→F→D.)

the current environmental conditions, we suggest that future experiments should consider testing this model across a number of different world conditions related to resource availability to correct for any potential "overfitting" of our model to our existing world conditions.

9.6.2 Anticipatory Adaptation via the Emergence of an Affect-Based Positive Feed-Forward Loop

Despite the fact that we had made no direct changes to oxytocin's function in our model, we find how oxytocin levels for bonded agents saw a significant increase in this investigation compared to previous groups (Table 9.5). Bonded agents also saw increased rates of intra-bond *Grooming* (Table 9.6) and tolerances to stress (Figure 9.11). We also found reductions in intra-bond *Aggression* rates (Table 9.6), allowing bond strength and efficacy of social support on stress-reduction to remain strong.

Even in our model—with its incremental approach towards modelling hormonal mechanisms—there appears to be a significant interplay between oxytocin levels, the strength and stress-reducing effects of affect-based social support, the types of social interactions, and the amount (and behavioural and physiological consequences) of stress experienced by agents: all of which impact the resultant social dynamics and interactions.

In line with our previous investigations (Chapters 7 and 8), we find how changing just one of these mechanisms (in this case, the adaptation of *EnergeticCost*) had a significant downstream effect on future behaviours and hormonal balance. We briefly describe the interplay of these hormonal mechanisms and behaviours below. Figure 9.12 shows the interplay between these hormones and social interactions in our model. We justify this description with our observed quantitative results from this investigation.

The adaptation of energy metabolism (*EnergeticCost*)—as a function of the *Stress Threshold* and cortisol levels—reduces the rate at which the *Energy* is depleted. This regulation of the amount of the *Energy* deficit reduces cortisol release. This then reduces the likelihood of *Aggression*, promoting *Grooming* and further increasing oxytocin and bond partner preference (via the Social Assessment Component, Equation 5.7).

The combined reduction in stress and (initially-)high oxytocin levels results in the emergence of an affect-based positive feedback loop between oxytocin, stress regulation and social interactions: promoting prosocial behaviours, reducing stress levels and maintaining strong bonds between bond partners prior to the world conditions becoming challenging.

We propose that, in line with the literature from neuroscience and biology [159], [175], [275], this positive feedback loop of hormonal mechanisms may underpin the emergence of affect-based relationships and behaviours via a positive feed-forward mechanism on future behaviours. Unlike the observations from some biological literature [159], [175], [276], our model saw these positive feed-forward loops emerge without any mechanisms related to reward-based systems. We suggest that these results may challenge existing hypotheses towards the underlying mechanisms of the maintenance and formation of affect-based relationships, and that these results can be used as motivation to investigate such mechanisms in biological systems further.

We propose that future work on models of social adaptation should take such mechanisms and effects into consideration, by also accounting for the potential downstream effects on incremental changes to their model in their analysis and observations. Such an approach may provide additional value to using artificial life models as a tool to understand these mechanisms in biological agents.

9.6.3 Affect-Based Regulation of Energy Metabolism Results in Emergence of Social Thermoregulation

In our model, the presence and interaction with bonded agents provided a stress-reducing response: lowering cortisol levels (Equation 7.2) and adapting energy metabolism via agent movement speed (Equation 9.1).

This regulation of energetic cost and speed resulted in periods where bond partners formed small "clusters" as they moved through the environment. These clusters emerged as agents perceived bond partners and, through a reduction in cortisol levels, would reduce their movement speed. This in turn resulted in a reduction in the rate at which *Energy* was depleted (Equation 9.3), providing an adaptive role on *Energy* regulation.

We consider that the emergence of these small clusters, which facilitated a reduction in energy expenditure through reduced speed and lower rates of *EnergeticCost*, can be considered as an adaptive behaviour akin to "social thermoregulation" in biological agents. We highlight some similarities between this hypothesised effect of social support and the behaviours emerging from our model.

The theory of social thermoregulation describes how social animals cooperate to preserve body heat, either through tactile contact [277] or non-tactile huddling between social partners [278]. This reduction in temperature then has downstream physiological benefits, particularly on energetic costs (as regulating temperature is metabolically-costly [277], [279] and a reduction in physiological stress [278]. Therefore, social thermoregulation has been considered to be a key social behaviour that directly links sociability to survival benefits [277], [278] as a cost-effective mechanism to regulate homeostasis.

Though we do not explicitly model the regulation of internal temperature to completely satisfy the definition of thermoregulation, our model does accounts for its underlying effects on energetic cost and metabolism (the *EnergeticCost* in our model) and the underlying principle of homeostatic regulation. The huddling behaviour described above resulted in reduced energetic cost (metabolism) for some agents in that group. This had the adaptive benefit of reducing their (critical) homeostatic drive for food intake.

Additionally, being in close proximity with other bonded agents as a result of this behaviour allowed agents more opportunities to perform *Grooming* with bond partners. This behaviour may be one reason as to why our results found increased rates of intra-bond *Grooming* compared to previous groups (Table 9.6). Figure 9.13 shows screenshots of one way in how this huddling behaviour emerged, along with our theoretical mechanisms for how this behaviour affected internal physiology.

Similar suggestions have been made by [277], who posits a potential relationship between the proximity required for social thermoregulation, and the emergence of higher-order cognitive functions—including affective regulation—which may have a causative link between physiological adaptation and the formation of affect-based relationships [280].

Taken together, our findings suggest that this socially-affective regulation of energetic cost—through an "economy of energy expenditure and action" [281]—may provide insight into a biologically-plausible

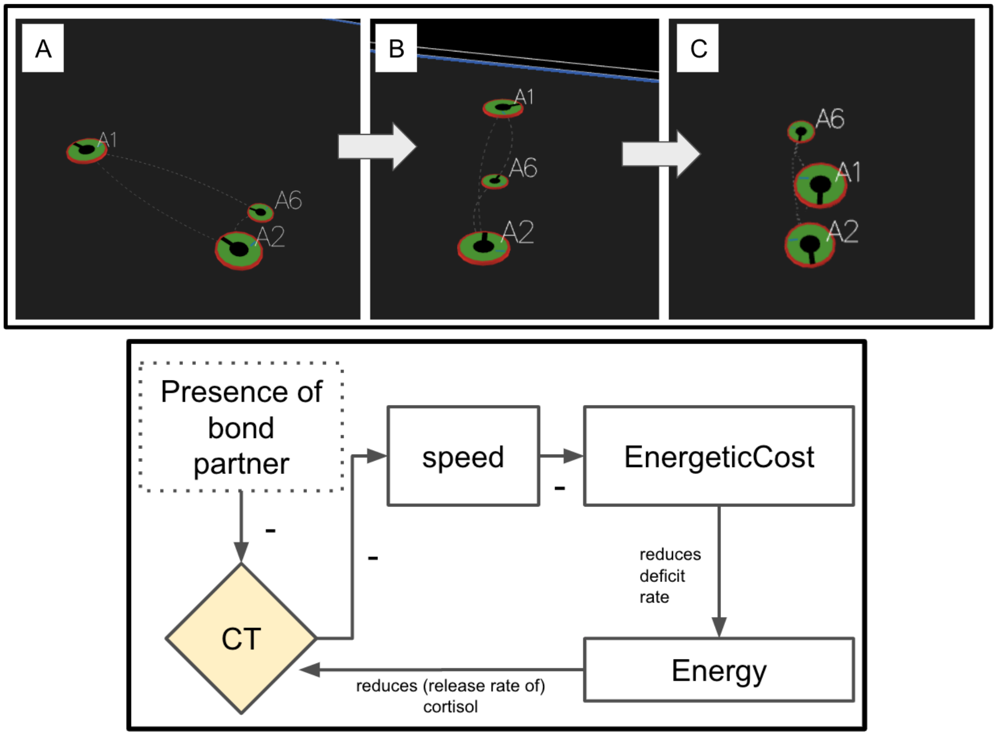

Figure 9.13 Top: An example of how "clusters" form between bonded agents. A&B: Bond partners agents A6 and A2 perceive A1 (bond partner) and begin to follow it, reducing cortisol levels and reducing energetic expenditure (speed) (Equation 9.1) and cost (Equation 9.3). As a result, proximity between bond partners is reduced. C: Bond partners continue to follow each other, resulting in a "cluster" moving together, further reducing cortisol in A1 and A6 and adapting the *EnergeticCost*. Bottom: An illustration of the theoretical interplay of mechanisms that underpin this adaptation. The presence of a bond partner reduces cortisol (CT) levels. Reduced levels of CT reduces agent speed, which then reduces its *EnergeticCost*. A lower *EnergeticCost* then reduces the rate at which *Energy* depletes, reducing its *Energy* deficit. The smaller physiological deficit then further reduces the release rate of cortisol (Equation 7.2), which then further reduces agent speed and *EnergeticCost*, resulting in ongoing adaptation of *Energy* depletion.

mechanism through which affect-based social support can have a direct physiological effect on homeostatic regulation in biological individuals. For both biological and artificial agents, using such a low-cost mechanism to adapt energy metabolism or to maintain homeostatic balance may allow energy or attention to be directed towards other life-critical actions.

In terms of a model of social allostasis, this (socially-)affective regulation of homeostatic variables may be an important adaptive mechanism that can provide significant improvement for future models, and therefore warrants further investigation.

However, we do not claim that this behaviour was solely responsible for driving performance of our model. We highlight that significant additional work needs to be undertaken in order to understand such physiological adaptation further. For instance, future work with our model may look to identify the specific temporal dynamics of these behaviours and its subsequent effects during periods of poor food availability (i.e. whether social thermoregulation is observed more during periods of poor food availability, and its impact on subsequent viability). Additionally, analysing positional or proximal data for agents may further investigate the frequency and precise benefits of this social thermoregulation behaviour).

Our findings may also be limited by the size of our bonded agent groups (which is currently three agents) which may limit the potential emergence of this behaviour in our current model. Future work should consider using larger groups of socially-bonded agents to evaluate whether this behaviour is still observed in larger bonded groups.

9.6.4 Higher-Ranking Bond Partners "Shield" Subordinates from Aggression

The proximity caused by the social thermoregulation behaviours described in the previous section had an additional adaptive benefit for bonded agents. We observed a number of cases (in Bond Condition A, A1-A2-A6 bonded) where higher-ranked agents (A1 and A2) provided a type of "shield" around the weaker A6 (Figure 9.14) from potential aggressors.

In situations where stressed, unbonded agents in the society could displace their stress by performing *Aggression* on the subordinate agent (A6), the presence of higher-ranking agents (A1 or A2) would "scare" those agents away (Figure 9.14). As a result of not receiving aggression, cortisol levels were not elevated for the subordinate A6, allowing their energetic cost and expenditure to remain low, which we propose may have impacted the viability of both the subordinate and the potential-aggressor. Specifically, as a result of not being able to displace stress, the *EnergeticCost* of unbonded, stressed agents would remain elevated, causing their *Energy* to deplete at a faster rate. For subordinate agents who did not receive *Aggression*, cortisol levels remained lower, resulting in lower rates of *EnergeticCost* and thus slower depletion of *Energy*.

Though our observations focused on how such interactions affected the viability and behaviours of bonded agents, we also recognise that these interactions may also have had wider impacts on unbonded agents. Future work may look to investigate these wider effects further. Considering our findings from previous chapters (Chapter 7 and 8), where we saw broader implications on the wider society emerging from intra-bond interactions, we predict that there may be a relationship between the frequency of (*Aggression*) performed (as a stress-coping mechanism) and the viability for unbonded agents.

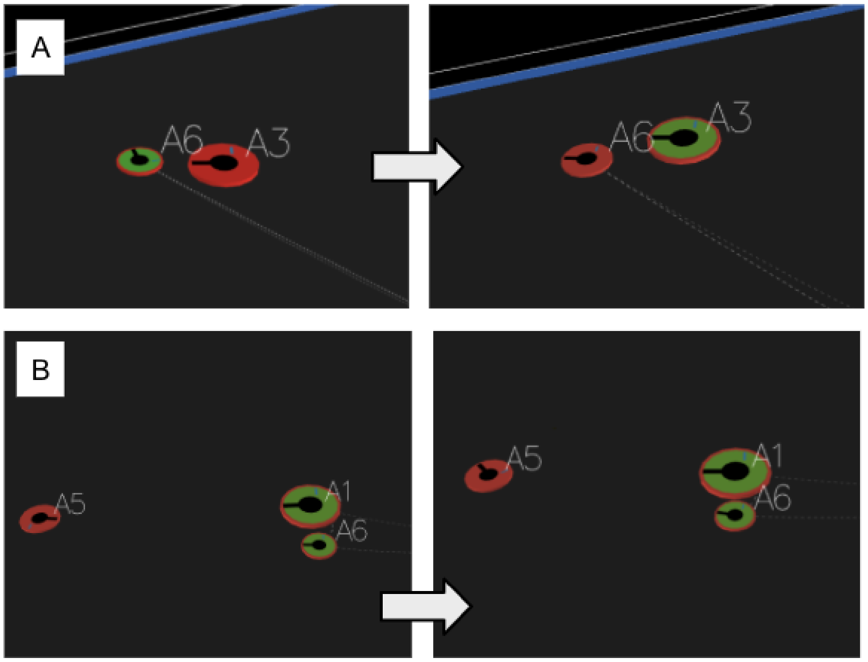

Figure 9.14 Screenshots of the "shield" against external *Aggression* provided by higher-ranking agents. Agents are facing the direction of the black line. Green agents are not stressed. Red agents are stressed. A: (Left) Agent A3 is stressed and looks to perform *Aggression* on A6. (Right) As *Aggression* is performed, cortisol in A3 is reduced and is increased in A6, resulting in it becoming stressed. B: (Left) A5 is stressed and looks to perform *Aggression* on A6. (Right) A5 perceives the higher-ranking A1 and avoids both A1 and A6 as a result.

9.6.5 Emergence of "Stress Eating" Even in Absence of Reward-Based Mechanisms

During our observations of a number of experimental runs, we had anecdotally observed that agents without social bond partners (unbonded agents) spent more time consuming food resources than those agents with social bonds. In absence of the stress-reducing effects provided by (the presence of) social bond partners, unbonded agents would rely on consummatory behaviours (in this case, *Eat*) to reduce their cortisol levels by reducing the physiological stress associated with internal physiological deficits (Equation 7.2).

We proposed that these behaviours—where unbonded agents were seen to eat more food than unbonded agents—may have some relevance to stress-induced eating behaviours ("stress eating") seen in biological animals. We find that these behaviours emerge even in absence of the reward-based mechanisms which have been hypothesised to play a role in biological animals [282].

To validate our initial anecdotal observations, we performed a preliminary post-hoc analysis using our

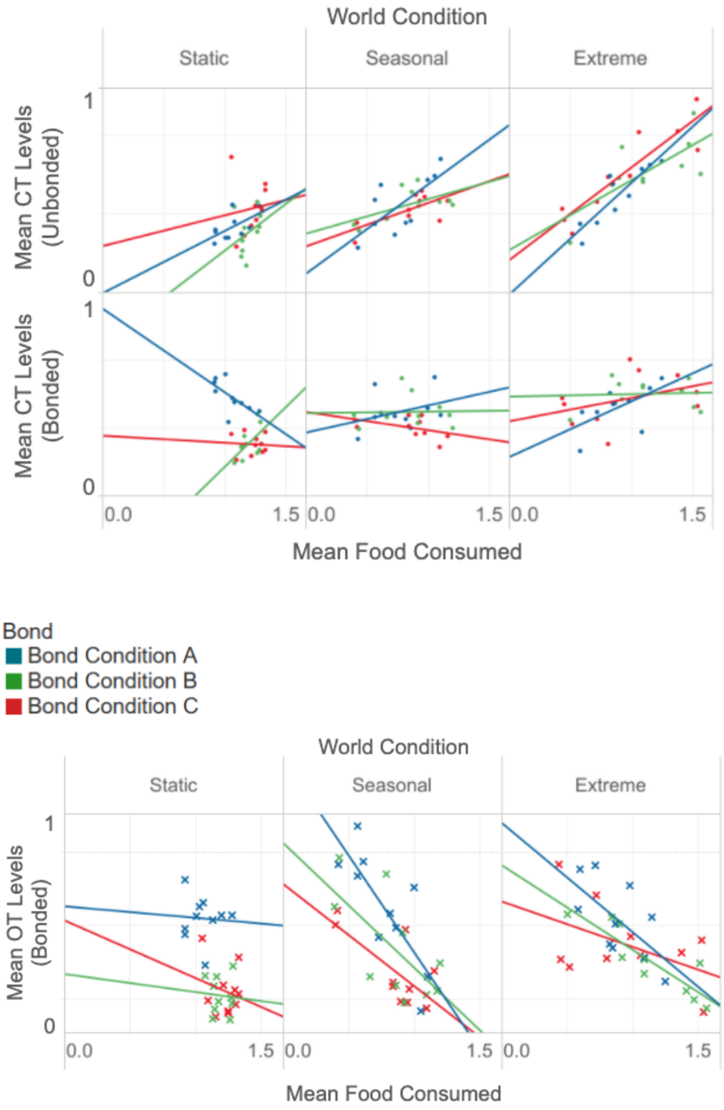

Figure 9.15 Top: Results of the correlation analysis between mean food consumption (x-axis) and mean cortisol levels for unbonded (top) and bonded (bottom) agents. Bottom: Results of the correlation analysis between mean food consumption (x-axis) and mean oxytocin levels for the three bonded agents. Lines denote trends for each of the respective bond combinations. Results from these statistical tests can be seen in Table 9.8.

data set to investigate three additional metrics: the mean amount of food consumed by each agent, the correlations between oxytocin levels and mean food consumption, and the correlations between cortisol levels (for both bonded and unbonded agents) and mean food consumption.

In line with our qualitative observations, we found that unbonded agents, on average, consumed larger

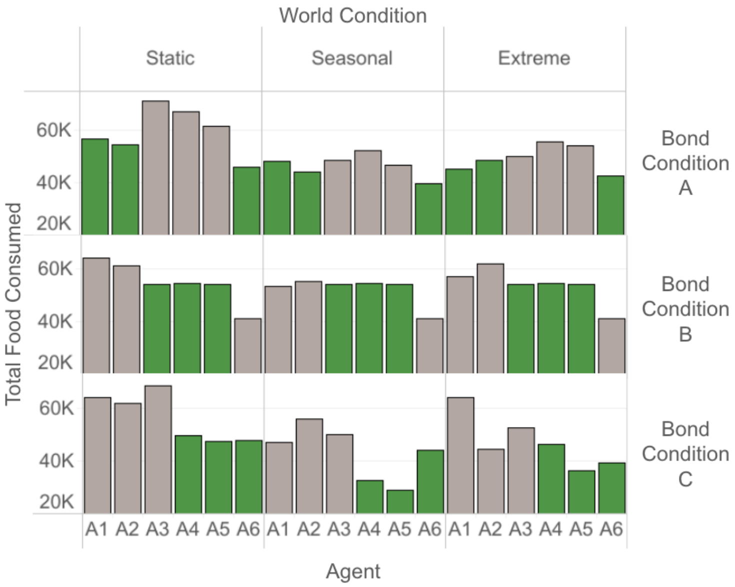

Figure 9.16 Total "units" of food consumed by bonded (green bars) and unbonded (grey bars) agents across each world condition, aggregated across all simulation runs.

amounts of food across all experimental conditions (Figure 9.16, grey bars), than bonded agents (Figure 9.16, green bars). Despite *Energy* being a critical internal variable and where food intake has been directly correlated with survival in our previous chapters, bonded agents remained comparatively (and, in most cases, more) viable compared to unbonded agents despite consuming less food overall.

Additionally, we found strong correlations between the mean cortisol levels of unbonded agents and their mean food consumption, suggesting a potential relationship between these two variables (Figure 9.15, top). This correlation became stronger as the relative world challenge increased (from *Static*, *Seasonal* and *Extreme*). Conversely, these correlations were weaker for agents with affective social bonds, suggesting no direct relationship between mean stress (cortisol) levels and the amount of food consumed (Figure 9.15, bottom left).

For bonded agents, we observed a moderately-strong correlation between mean oxytocin levels and the amount of food consumed (Figure 9.15, bottom) in the *Seasonal* and *Extreme* world conditions, with weak correlations observed in the *Static* world condition. In other words, the higher the amount of oxytocin for bonded agents in these challenging worlds, the less food they consumed on average across all experimental runs. Results from our statistical tests are seen in Table 9.8.

This relationship between increased oxytocin and reduced food intake (including reward-driven food intake) has also been seen in real-world animals (including humans) [268], [283]–[285]

World Condition	**Bond Condition**	r^2	p
Static	A	0.004	0.86
Static	B	0.005	0.53
Static	C	0.005	0.84
Seasonal	A	0.629	0.01*
Seasonal	B	0.662	0.03*
Seasonal	C	0.525	0.01*
Extreme	A	0.603	0.04*
Extreme	B	0.769	0.01*
Extreme	C	0.334	0.12

Table 9.8 Results of the statistical tests for correlations between mean oxytocin levels for bonded agents and mean food consumed. Results are calculated using the mean values of all three agents in each bond condition.

In sum, higher levels of cortisol levels were associated with increased food intake for unbonded agents, and higher levels of oxytocin for bonded agents was found to be associated with lower food intake. Though the causal relationship between these variables warrants further investigation, these initial findings still mirror those seen in biological literature. For instance, cortisol (including exogenous treatment), and exposure to high levels of stressors, have been highly correlated with an increased intake of food [282], [286]–[288]. This behaviour has also been considered to emerge through a ”lack of cognitive control” of these consummatory behaviours [289].

From the findings of the current literature [268], [284], [285], [289] and our own results, we suggest that the presence (or absence) of social support—with its positive feed-forward effects on oxytocin—may play a role in regulating the physiological processes that precede stress-eating behaviours. In our model, we find that behaviours akin to stress eating emerged even in absence of any of the reward-based mechanisms that has been implicated in such behaviours in biological animals [282]. We therefore suggest that the mechanisms in our model may provide an alternative explanation into some of the underlying mechanisms in stress-induced eating behaviours.

However, the relationship between increased oxytocin and cortisol levels, and the specific effects on food consumption is unclear from our current results. Furthermore, our model does not currently account for “overconsumption” of resources associated with the “stress-eating” behaviour [282]. As this did not form part of the initial hypothesis, we limit this preliminary analysis and retain this avenue of investigation for potential future work. Nevertheless, we highlight these early findings may mirror studies from ethology and neuroscience that considers the relationship between (the lack of) social support and (over)consumption of food [290]–[292].

9.7 Summary

Continuing our investigation into the "social buffering" effects of social support, we have proposed and investigated a potential mechanism through which social support can play an adaptive role on physiological health and facilitate long-term survival in socially-supported individuals.

By accounting for the previously-investigated relationship between oxytocin and social support (Chapters 6 to 9), individual "personality" associated with stress tolerance (an abstraction of the autonomic nervous system, Chapters 5 and 8), and observed effects from biological systems, we consider that the presence of social support, via oxytocin's "buffering" effects on the stress system, can provide a further adaptive effect on one type of physiological cost (energy metabolism) associated with elevated stress levels.

Based on our observations from previous chapters and findings from biological systems, we hypothesised that this mechanism of physiological adaptation would provide survival-related benefits to socially-bonded individuals. We considered this potential pathway as a type of "affect-based regulation" of a physiological process.

We accounted for this hypothesised pathway in our agent model and conducted our experiments using an identical set up to the experiments in (Chapters 7 and 8): testing our hypothesis in the same society of six agents across three different bond partner combinations related to social rank. We studied these different bond combinations across three different environmental conditions of varying challenge related to food resource availability. Through our approach of a simulated model, we aimed to address several methodological challenges from real-world studies that have previously attempted to investigate the relationship between (affect-based) social support and its effects on physiological health markers.

Our results showed that this proposed mechanism of physiological adaptation (the adaptation of energy metabolism) through social support's buffering of the stress system provided significant advantages in terms of viability-related performance for agents with social bonds—particularly in challenging, dynamic world environments—when compared to groups where this adaptative mechanism was not present. We also observed increased rates of socio-positive behaviours between bonded agents when this adaptive mechanism was accounted for. Our results suggest that such a mechanism may have biological plausibility into explaining one potential pathway linking the effects of affect-based social support and positive health outcomes.

Mirroring suggestions from biology and neuroscience, we highlight the emergence of an affect-based feed-forward loop between social support, the secretion of hormones and the adaptation of physiological processes when additional stress-reducing mechanisms are accounted for in our model. These feed-forward affective loops result in "anticipatory" behavioural and physiological adaptation: adapting behaviours and physiology prior to homeostasis being threatened. This anticipatory adaptation of homeostatic processes and behaviours—underpinned by affect-based social relationships—can be contextualised as a mechanism of social allostasis.

Our observations also saw the emergence of behaviours akin to social thermoregulation—a cost-effective behaviour to socially regulate homeostasis—between socially-bonded agents in some conditions. We propose that the emergence of this behaviour may shed some light into how social support provides physiological health benefits for individuals with social bonds, as a cost-effective mechanism to regulate energetic or physiological costs—and propose that this is an avenue for further investigation in both natural and artificial systems. Though this initial work appears to provide some adaptive benefits, further work should look to investigate these behaviours and mechanisms further as a potentially low-cost mechanism of adaptation for embodied social agents. We suggest that this low-cost adaptation of one physiological process through social relationships may free up attentional, computational or "physiological" resource to prioritise other viability-related goals in autonomous agents. We highlight this as an avenue for future investigations.

Our qualitative analysis found the emergence of behaviours akin to "stress-eating" by unbonded agents, and we performed preliminary post-hoc analysis to further investigate this. We found that these behaviours emerged even in absence of other (reward-based) mechanisms or systems that have been implicated in biological systems. We therefore proposed that the mechanisms that have been investigated in our model may provide an alternative explanation into some of the mechanisms linking (a lack of) social support and stress-induced eating in biological agents. From our results, we hypothesise that affect-based social support may have a downstream effect on the reduction of stress-eating, via oxytocin's buffering effects on the stress system and energy metabolism. However, significant future work is necessary to further investigate this, and we retain this avenue of research for further analysis and research.

Finally, the embodied model of adaptation—underpinned by affect-based social relationships and mechanisms of social allostasis—that we have developed and systematically investigated has demonstrated significant performance advantages with respect to viability across several dynamic environmental and social conditions, at both an individual and group level. We have presented our final agent model in this investigation, and propose that future models of social adaptation build on the framework that we have presented.

Chapter Ten

Conclusion

10.1 Summary

In this Ph. D. , we aimed to investigate the role of affect-based social relationships, underpinned by social allostasis principles, as potential mechanisms of adaptation for a homeostatically-controlled embodied agent model, to promote the stability (or “viability”) of this model across dynamic conditions. As mechanisms of social allostasis are yet to be systematically investigated in current approaches, we addressed this limitation by developing a model of social adaptation inspired by the principles of social allostasis.

We focused our investigation on the hypothesised mechanisms of two biological hormones—oxytocin and cortisol—that underpin affect-based social relationships and social allostatic adaptation.

Inspired by the findings of social species from ethology, along with findings from previous models of adaptation for homeostatically-controlled systems, we hypothesised that adaptation through affect-based social relationships and social allostatic principles would significantly improve the viability-related performance of our homeostatically-controlled embodied agent model across dynamic physical and social conditions.

We have used an artificial life approach to systematically investigate numerous hypothesised hormonal mechanisms that underpin social adaptation and social allostasis from biological systems. We developed computational abstractions of these hormonal effects in our agent model and systematically tested existing hypotheses of their effects across physically and socially-challenging environments. Different physical challenges were tested using different worlds related to food availability, and social challenges were tested in the form of a dominance rank-based society (from Chapter 5 onwards).

We found how an embodied agent model of adaptation, underpinned by affective social relationships, interactions and social allostasis mechanisms, significantly improved the viability performance of our agent model across several dynamic social and physical environments. We contribute our social allostasis-inspired model as a framework to address the absence of this approach in the current literature, and propose that this approach towards adaptation of agent models should be built on and investigated further. We propose several avenues for further investigation below.

In the process of our investigation, we presented a quantitative measure and modelling of dynamic, affective bond strength between dyads, called the "Dyadic Strength Index" in Chapter 5. This measure to quantify affective bond strength capitalises on our artificial life system, addressing limitations of social bond measuring and modelling in current approaches in artificial agents. Given its flexibility of implementation, we propose that future artificial life models investigating or modelling social bonds continue to use the framework provided by the Dyadic Strength Index to measure and model affective social relationships.

We also developed a mathematical model to describe the "trade-off" associated with partner selection in (rank-based) biological systems, which we have called the "Social Assessment Component" and introduced it in our agent model. This model accounts for social rank differences and the presence (and strength) of an existing social bond to contextualise social behaviours and partner selection. Though our implementation uses mechanisms unique to our current model (the Dyadic Strength Index and oxytocin), its framework nevertheless proposes flexible integration into future social simulation models.

Across all of our investigations with social bond partners (Chapters 6, 7, 8, 9), we find how endowing half of the society with social bonds and mechanisms of adaptation also provided viability benefits for the wider social group. We suggest that these findings may have implications for the development of models in the future. Specifically, that not all agents in a social group require (the same type of) adaptive mechanisms in order to achieve improved viability.

During the course of these investigations, we have used our biologically-inspired agent model as a scientific tool to address methodological challenges or limitations of real-world studies, and tested numerous hypothesised effects of hormonal mechanisms from biological literature. In the process, we have observed the emergence of numerous individual and group interactions and dynamics, and discussed their potential relevance to similar behaviours seen in natural systems. As this analysis was beyond the scope of the current research, we propose that the underlying mechanisms that we have investigated should provide potential avenues for investigation in, and comparisons with, behaviours in natural systems. We summarise these in Section 10.2.

Below, we present a summary of the work undertaken in each of the chapters.

In our first investigation **(Chapter 3)**, we investigated two contrasting hypotheses of oxytocin (a hormone found to play a role in social interaction) on the processing of social cues. The effects of the prosocial hypothesis (which posits that oxytocin increases prosocial behaviours through increased social salience) and the interactionist hypothesis (which proposes that oxytocin's effects should be more context-dependent) were investigated with respect to agent viability. Those results found viability improvements when oxytocin increased social salience (the prosocial hypothesis), though we considered the limited context of those findings.

In **Chapter 4**, we looked to address one aspect of the limited context by accounting for oxytocin's hypothesised effects on internal signals. These internal signals occurred when individuals performed a prosocial behaviour (grooming). We considered this to be a type of "affective touch". Contrary to our earlier chapter, we found that the beneficial effects of the prosocial hypothesis of oxytocin were not universal. Rather, from a viability perspective, oxytocin effects depended on several contextual factors. We therefore found support for the "interactionist" hypothesis of oxytocin's effects which warranted further investigation.

One notable limitation from these earlier chapters was the lack of social context to our findings. From our observations, we considered that oxytocin's prosocial effects should also account for the social context. The investigation in Chapter 5 account for this social context through a rank-based social structure and affective social bond partners. We investigated one hypothesised, contextual role of oxytocin—that it specifically promotes partner preference for social interaction in a rank-based society—and found that this effect provided survival-related benefits for agents with these affective social bonds in a challenging physical environment.

From our observations, we considered that social dynamics may be affected by the interactions and behaviours performed by dominant individuals in a society, and that hierarchical societies can also be a source of stress for some individuals. In these stressful environments, the presence of social support can also play a role in regulating the stress response—known as the "social buffering" hypothesis of social support—which may facilitate their long-term survival in challenging conditions.

In **Chapter 6**, we presented our initial work into the "social buffering" hypothesis, investigating how different stress-related "personalities"—related to individual tolerances to stress—affect rank-based social dynamics. We had hypothesised that rank-based social dynamics would be determined by the interactions of higher-ranked agents, but our results found that this was not always the case. Instead, these interactions and subsequent dynamics were dependent upon both individual personality traits, related to stress tolerance, and environmental conditions. In some conditions, wider social dynamics and interactions were also determined by the behaviours and stress tolerances of subordinates.

In **Chapter 7**, we extended beyond these strictly rank-based interactions, and investigated the role of social support, through affective social bonds, on the reduction of stress. We investigated how affect-based interactions affected social dynamics when the physical environment was challenged. Despite our hypothesis that social bonds would reduce stress and therefore improve viability, our results only found limited support for this. Instead, we found that, unless these affect-based social bonds remained persistently strong, these affective bonds became susceptible to "damage beyond repair" if the social and physical environment posed significant challenges. This supported the hypothesis that stress reduction is not the only way in which social support "buffers" the stress response.

We then investigated a second hypothesised stress-regulating role of social support—the "buffering" of stress tolerances—in **Chapter 8**. This adaptation of the stress tolerance, an abstraction of the autonomic nervous system, can be seen as a mechanism of social allostasis. Our results found how, over time, different individuals adapted their "personalities" (tolerances to stress) and social interactions based on their dominance rank and the types of affective bonds that they maintained. We found how accounting for two of these stress-reducing effects resulted in lower overall stress levels and improved viability-related performance for the overall society. Contrary to expectations, our investigations showed how higher-ranked individuals adapted a lower stress tolerance in some contexts, where lower-ranked individuals maintained a higher tolerance to stress in order to maintain viability across challenging social and physical environments.

In our final chapter **(Chapter 9)**, we investigated a third hypothesised role of affect-based social support: that it also provides an adaptive effect on the physiological cost associated with stress. Combining previous literature, our previous results, and motivated by methodological challenges in biological systems, we proposed a potential pathway linking social support and its effects on energetic cost associated with stress (energy metabolism) via oxytocin's adaptive effects on the stress system, and used our model to investigate this hypothesised effect on viability and social interactions. Our results found that this mechanism provided some significant advantages to the wellbeing (viability) of agents in some conditions, and we propose that these mechanisms should be investigated in future artificial and animal models. We saw the emergence of stress-regulating behaviours related to "social thermoregulation" and "stress-eating" and proposed that these behaviours warrant further investigation using this model. This chapter also presented our final embodied model of social adaptation inspired by social allostasis, and we demonstrated significant adaptive potential of using this model across a variety of physical and socially-challenging conditions.

10.2 Limitations & Future Work

Given the finite scope of the work undertaken and the nature of the artificial life approach that we have taken, we recognise that our current model and experimental methodology still have a number of limitations and that future work should look to address these. We discuss a number of limitations here and propose a number of potential avenues in order to guide future work.

As discussed in Section 2.1, we recognise that our observations regarding the emergence of social interactions and dynamics throughout this experiment may suffer from subjectivity, and affected by the multi-faceted position we hold across these investigations. To mitigate this, we have capitalised on the capabilities of our artificial life model and provided quantatitive context to our qualitative analysis of interactions and behaviours as best as we could. Such qualitative narrative will always suffer from subjectivity: what is deemed "interesting" behaviour in these models to one researcher may be deemed uninteresting by another. That is not an issue that we have looked to overcome during this research, and so we have contributed our model to the public domain in the hope that future work contributes to our observations. By either using our model or replicating our investigations in a different (artificial) system, researchers may provide further insight into, or even challenge, the observations that we have made using our model.

We believe that our results were driven, in part, by the initial hormonal concentrations—with maximum oxytocin levels and no cortisol—and affective bond strength (initialised at their maximum strength) used in each investigation. While this approach minimised complexity and allowed us to systematically investigate each of our respective hypotheses, it may present limitations when looking to extrapolate our findings to biological systems.

Using our model, we propose that future work should consider moving beyond the biological specificity of our current approach. Specifically, investigating the effects of different parameter settings with respect to hormonal concentrations, or different initial strengths of affective bonds, may provide a more holistic understanding of the adaptive effects of our socially-adaptive model. Furthermore, we propose that future work capitalises on the capabilities of the artificial life approach: manipulating these parameters in ways that cannot be achieved in natural systems, to further investigate the emergence of social dynamics.

As we suggested in Chapter 9, we also suggest that the world conditions that we had conducted our investigations in may have contributed to its limited performance improvements in latter stages of the investigation. To investigate whether this is the case, and in order to avoid creating a model of adaptation that "overfits" our environments, future work should consider investigating this model across additional environmental conditions related to resource availability.

We also propose that future work should also consider the dynamic nature of social hierarchy and affect-based social bonds. To minimise complexity, the work undertaken here has focused on fixed dominance ranks and several fixed bond combinations in the investigations. For a more biologically-viable model, and similar to the approach taken by [72], [75], [77], future work should account for the notion that dominance rank can change over time, or that the notion of "social rank" can extend beyond our current "dominance" approach. For instance, social rank may be determined as a function of social bond formation and maintenance (i.e. being more sociable results in a higher social rank).

Our model can also be extended to study the formation of additional affect-based bonds, or accounting for changing affective appraisal of one's own rank before performing social behaviours. Accounting for changes in dominance rank, the formation of new social bonds, and the perception of rank differences in this way may have significant effects on the social dynamics that emerge from our society.

We would like to see the performance of our adaptive model compared against other previously proposed homeostatically-controlled models in the same environment, such as the epigenetic, allostatic-type model proposed by [40], and the allostatic-type framework proposed by [103]. This can be investigated by replicating these models and endowing either the whole society of artificial agents with the same model, or different agents in the same environment with different adaptive mechanisms. Considering that adaptation through allostasis has a cost (i.e. the "allostatic load" [122]), evaluating the most cost-effective mechanism for allostatic-inspired adaptation can have a significant impact in the development of future adaptive models.

Further work may also look at closely investigating the specific temporal relationship with respect to social interactions, hormonal concentrations and model performance. Understanding if there are any potential key features—such as the distribution of different interactions, or hormonal concentrations—during the early stages of agent "life" may provide an opportunity to develop a predictive model to predict the emergent social dynamics, evolution of relationships and even the long-term wellbeing based on these key features. Based on the different affective interactions that we have observed during our investigations, such a predictive model may also have implications in understanding affective development for biological agents: further utilising our simulated model as a tool for hypothesis generation.

Future investigations modelling affect-based relationships as mechanisms for adaptation using our framework may also consider accounting for additional hormonal mechanisms in the model. Oxytocin's effects on promoting social behaviours are widely considered to be a result of its effects on the reward-driven circuitry in biological agents. Future work with our model may also choose to extend this system to include abstractions of these additional mechanisms. Such reward-based mechanisms have been investigated in previous embodied models of adaptation [39].

Finally, we would like to see the results that we have presented throughout this investigation, and summarised below, to inspire future work in biological systems. Additionally, we propose that our simulation model—inspired by biological mechanisms—be used as an additional tool for cross-disciplinary research, in order to address some of the methodological challenges of studying biological agents and to contribute and generate future hypotheses of these mechanisms in the real world. Specifically, in contrast to our earlier suggestion (in Chapter 2) that we had no biological target in mind with our modeling approach, we recognise that the later chapters drew inspiration from mechanisms and behaviours observed in non-human primate (particularly chimpanzees) societies. We are therefore supportive of Webb's [121] perspective that "animats" can be analogous to biological agents, and hold that the agents in our model can be seen, in some way, as "artificial primates". Assuming such studies were possible in the natural world, we would like to see these later studies (chapters 7 to 9) replicated in non-human primate societies. Alternatively, part of our future work should look to see whether similar work has already been undertaken, and to validate our existing results with any similar studies from the natural world. Here, we summarise our most interesting findings from our investigation, which we believe may provide avenues for further investigation in both natural and artificial societies:

- Social "ignorance" may be determined by low oxytocin levels or its effects on decreasing attention to social cues (Chapter 3).
- Social amnesia towards social bond partners in later stages may be facilitated by a lack of early-stage bond reinforcement (Chapter 6).
- Social dynamics, with respect to the "aggressiveness" of individuals in a society, may be partially determined by "personality" types of dominant agents, but also by the personalities and interactions of subordinates (Chapter 6).
- High oxytocin levels may promote altruistic behaviours in "parent-offspring" relationships (Chapter 5).
- Affective contagion in a small society may depend upon the interactions and behavioural adaptation of individuals in bonded groups (Chapters 6, 7, 8, 9).
- Oxytocin levels and the quality of affective bonds may play a role in determining whether stress-coping strategies are "tend-and-befriend" or "fight-or-flight" (Chapter 7).

- The efficacy of social buffering effects may be dependent upon the social and environmental contexts (Chapter 8).
- Social support may promote the wellbeing of socially-supported individuals via its downstream effects on energy metabolism and adaptation of the stress response system (Chapter 9).
- Affect-based social regulation of energy metabolism may underpin "social thermoregulation" behaviours (Chapter 9).

REFERENCES

[1] M. Ahmad, O. Mubin, and J. Orlando, “A systematic review of adaptivity in human-robot interaction,” *Multimodal Technologies and Interaction*, vol. 1, no. 3, p. 14, 2017.

[2] R. Kachouie, S. Sedighadeli, R. Khosla, and M.-T. Chu, “Socially assistive robots in elderly care: A mixed-method systematic literature review,” *International Journal of Human-Computer Interaction*, vol. 30, no. 5, pp. 369–393, 2014.

[3] B. Robins, K. Dautenhahn, and J. Nadel, “Kaspar, the social robot and ways it may help children with autism–an overview,” *Enfance*, no. 1, pp. 91–102, 2018.

[4] A. Coninx, P. Baxter, E. Oleari, S. Bellini, B. Bierman, O. Henkemans, L. Cañamero, P. Cosi, V. Enescu, R. Espinoza, *et al.*, “Towards long-term social child-robot interaction: Using multi-activity switching to engage young users,” *Journal of Human-Robot Interaction*, 2016.

[5] A. Tapus *et al.*, “Stress game: The role of motivational robotic assistance in reducing user’s task stress,” *International Journal of Social Robotics*, vol. 7, no. 2, pp. 227–240, 2015.

[6] P. Baxter, E. Ashurst, R. Read, J. Kennedy, and T. Belpaeme, “Robot education peers in a situated primary school study: Personalisation promotes child learning,” *PloS one*, vol. 12, no. 5, e0178126, 2017.

[7] M. Milne, M. H. Luerssen, T. W. Lewis, R. E. Leibbrandt, and D. M. Powers, “Development of a virtual agent based social tutor for children with autism spectrum disorders,” in *the 2010 international joint conference on neural networks (IJCNN)*, IEEE, 2010, pp. 1–9.

[8] F. Pecune, A. Cafaro, M. Ochs, and C. Pelachaud, “Evaluating social attitudes of a virtual tutor,” in *International Conference on Intelligent Virtual Agents*, Springer, 2016, pp. 245–255.

[9] A. Oker, F. Pecune, and C. Declercq, “Virtual tutor and pupil interaction: A study of empathic feedback as extrinsic motivation for learning,” *Education and Information Technologies*, pp. 1–16, 2020.

[10] M. E. Foster, R. Alami, O. Gestranius, O. Lemon, M. Niemelä, J.-M. Odobez, and A. K. Pandey, “The mummer project: Engaging human-robot interaction in real-world public spaces,” in *International Conference on Social Robotics*, Springer, 2016, pp. 753–763.

[11] E. Zalama, J. G. Garcıa-Bermejo, S. Marcos, S. Domınguez, R. Feliz, R. Pinillos, and J. López, “Sacarino, a service robot in a hotel environment,” in *ROBOT2013: First Iberian Robotics Conference*, Springer, 2014, pp. 3–14.

[12] D. Vernon, R. Lowe, S. Thill, and T. Ziemke, “Embodied cognition and circular causality: On the role of constitutive autonomy in the reciprocal coupling of perception and action,” *Frontiers in psychology*, vol. 6, p. 1660, 2015.

[13] C. Darwin, A. R. Wallace, *et al.*, “Evolution by natural selection.,” *Evolution by natural selection.*, 1958.

[14] J. Schulkin *et al.*, *Allostasis, homeostasis, and the costs of physiological adaptation.* Cambridge University Press, 2004.

[15] J. Panksepp, *Affective neuroscience: The foundations of human and animal emotions.* Oxford university press, 2004.

[16] R. M. Wittig, C. Crockford, A. Weltring, K. E. Langergraber, T. Deschner, and K. Zuberbühler, “Social support reduces stress hormone levels in wild chimpanzees across stressful events and everyday affiliations,” *Nature communications*, vol. 7, no. 1, pp. 1–8, 2016.

[17] A. C. DeVries, E. R. Glasper, and C. E. Detillion, “Social modulation of stress responses,” *Physiology & behavior*, vol. 79, no. 3, pp. 399–407, 2003.

[18] O. Schülke, J. Bhagavatula, L. Vigilant, and J. Ostner, “Social bonds enhance reproductive success in male macaques,” *Current Biology*, vol. 20, no. 24, pp. 2207–2210, 2010.

[19] E. Z. Cameron, T. H. Setsaas, and W. L. Linklater, “Social bonds between unrelated females increase reproductive success in feral horses,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 33, pp. 13 850–13 853, 2009.

[20] J. Holt-Lunstad, T. B. Smith, and J. B. Layton, “Social relationships and mortality risk: A meta-analytic review,” *PLoS med*, vol. 7, no. 7, e1000316, 2010.

[21] W. B. Cannon, “Organization for physiological homeostasis,” *Physiological reviews*, vol. 9, no. 3, pp. 399–431, 1929.

[22] W. R. Ashby, *Design for a Brain.* Chapman & Hall London, 1954.

[23] R. Carpenter, “Homeostasis: A plea for a unified approach,” *Advances in Physiology Education*, vol. 28, no. 4, pp. 180–187, 2004.

[24] S. C. Woods and D. S. Ramsay, “Homeostasis: Beyond curt richter,” *Appetite*, vol. 49, no. 2, pp. 388–398, 2007.

[25] P. Sterling, “Allostasis: A model of predictive regulation,” *Physiology & behavior*, vol. 106, no. 1, pp. 5–15, 2012.

[26] B. S. McEwen, “Stressed or stressed out: What is the difference?” *Journal of Psychiatry and Neuroscience*, vol. 30, no. 5, p. 315, 2005.

[27] T. A. Day, “Defining stress as a prelude to mapping its neurocircuitry: No help from allostasis,” *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, vol. 29, no. 8, pp. 1195–1200, 2005.

[28] D. S. Ramsay and S. C. Woods, “Clarifying the roles of homeostasis and allostasis in physiological regulation.,” *Psychological review*, vol. 121, no. 2, p. 225, 2014.

[29] J. Schulkin, *Rethinking homeostasis: Allostatic regulation in physiology and pathophysiology.* MIT Press, 2003.

[30] S. E. Taylor, “Tend and befriend: Biobehavioral bases of affiliation under stress,” *Current directions in psychological science*, vol. 15, no. 6, pp. 273–277, 2006.

[31] R. M. Wittig, C. Crockford, T. Deschner, K. E. Langergraber, T. E. Ziegler, and K. Zuberbühler, "Food sharing is linked to urinary oxytocin levels and bonding in related and unrelated wild chimpanzees," *Proceedings of the Royal Society B: Biological Sciences*, vol. 281, no. 1778, p. 20 133 096, 2014.

[32] A. K. Seth, *The cybernetic Bayesian brain.* Open MIND. Frankfurt am Main: MIND Group, 2014.

[33] I. Muntean and C. D. Wright, "Autonomous agency, ai, and allostasis: A biomimetic perspective," *Pragmatics & cognition*, vol. 15, no. 3, pp. 485–513, 2007.

[34] D. Canamero, "Modeling motivations and emotions as a basis for intelligent behavior," *In Proceedings of the first international conference on Autonomous agents*, vol. 8, pp. 148–155, 1997.

[35] O. Avila-Garcia and L. Cañamero, "Using hormonal feedback to modulate action selection in a competitive scenario," *From animals to animats*, vol. 8, pp. 243–252, 2004.

[36] R. L. French and L. Canamero, "Introducing neuromodulation to a braitenberg vehicle," in *Proceedings of the 2005 ieee international conference on robotics and automation*, IEEE, 2005, pp. 4188–4193.

[37] L. Cañamero and O. Avila-Garcıa, "A bottom-up investigation of emotional modulation in competitive scenarios," in *International Conference on Affective Computing and Intelligent Interaction*, Springer, 2007, pp. 398–409.

[38] M. Egbert and L. Canamero, "Habit-based regulation of essential variables," in *Artificial Life Conference Proceedings 14*, MIT Press, 2014, pp. 168–175.

[39] M. Lewis and L. Canamero, "Hedonic quality or reward? a study of basic pleasure in homeostasis and decision making of a motivated autonomous robot," *Adaptive Behavior*, vol. 24, no. 5, pp. 267–291, 2016.

[40] J. Lones, M. Lewis, and L. Cañamero, "A hormone-driven epigenetic mechanism for adaptation in autonomous robots," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 10, no. 2, pp. 445–454, 2017.

[41] J. Lones, L. Cañamero, and M. Lewis, "Epigenetic adaptation in action selection environments with temporal dynamics," in *Artificial Life Conference Proceedings 13*, MIT Press, 2013, pp. 505–512.

[42] L. Pitonakova, "Ultrastable neuroendocrine robot controller," *Adaptive Behavior*, vol. 21, no. 1, pp. 47–63, 2013.

[43] C. Sauzé, M. Neal, *et al.*, "A neuro-endocrine inspired approach to long term energy autonomy in sailing robots," *Proceedings of TAROS (Towards Autonomous Robotic Systems)*, pp. 255–262, 2010.

[44] J. Timmis, M. Neal, and J. Thorniley, "An adaptive neuro-endocrine system for robotic systems," in *2009 IEEE Workshop on Robotic Intelligence in Informationally Structured Space*, IEEE, 2009, pp. 129–136.

[45] H. Hamann, J. Stradner, T. Schmickl, and K. Crailsheim, "A hormone-based controller for evolutionary multi-modular robotics: From single modules to gait learning," in *IEEE Congress on Evolutionary Computation*, IEEE, 2010, pp. 1–8.

[46] D. Moser, R. Thenius, and T. Schmickl, "First investigations into artificial emotions in cognitive robotics," in *International Workshop on Medical and Service Robots*, Springer, 2016, pp. 213–227.

[47] A. Montebelli, R. Lowe, and T. Ziemke, “The cognitive body: From dynamic modulation to anticipation,” in *Workshop on Anticipatory Behavior in Adaptive Learning Systems*, Springer, 2008, pp. 132–151.

[48] S. Schaal, A. J. Ijspeert, A. Billard, S. Vijayakumar, and J.-A. Meyer, “Homeostatic plasticity in recurrent neural networks,” 2004.

[49] H. Iizuka and E. A. Di Paolo, “Extended homeostatic adaptation: Improving the link between internal and behavioural stability,” in *International Conference on Simulation of Adaptive Behavior*, Springer, 2008, pp. 1–11.

[50] P. Zahadat and T. Schmickl, “Evolving controllers for programmable robots to influence non-programmable lifeforms: A casy study,” in *European Conference on the Applications of Evolutionary Computation*, Springer, 2015, pp. 831–841.

[51] O. C. Görür, B. Rosman, F. Sivrikaya, and S. Albayrak, “Social cobots: Anticipatory decision-making for collaborative robots incorporating unexpected human behaviors,” in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI ’18, Chicago, IL, USA: Association for Computing Machinery, 2018, pp. 398–406, ISBN: 9781450349536. DOI: 10.1145/3171221.3171256. [Online]. Available: https://doi.org/10.1145/3171221.3171256.

[52] G. Hoffman and C. Breazeal, “Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team,” in *Proceedings of the ACM/IEEE international conference on Human-robot interaction*, 2007, pp. 1–8.

[53] A. F. Winfield, V. V. Hafner, and R. Poli, “Anticipation in robotics,” *Handbook of Anticipation: Theoretical and Applied Aspects of the Use of Future in Decision Making*, pp. 1–30, 2018.

[54] H. Mellmann, B. Schlotter, L. Musiolek, and V. V. Hafner, “Anticipation as a mechanism for complex behavior in artificial agents,” in *Artificial Life Conference Proceedings*, MIT Press, 2020, pp. 157–159.

[55] M. Ghafurian, N. Budnarain, and J. Hoey, “Role of emotions in perception of humanness of virtual agents,” in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, International Foundation for Autonomous Agents and Multiagent Systems, 2019, pp. 1979–1981.

[56] A. B. Youssef, M. Chollet, H. Jones, N. Sabouret, C. Pelachaud, and M. Ochs, “Towards a socially adaptive virtual agent,” in *International Conference on Intelligent Virtual Agents*, Springer, 2015, pp. 3–16.

[57] M. G. Bedia, M. Heras-Escribano, D. Cajal, M. Aguilera, and X. E. Barandiaran, “Towards modelling social habits: An organismically inspired evolutionary robotics approach,” in *Artificial Life Conference Proceedings*, MIT Press, 2019, pp. 341–348.

[58] A. Acerbi and S. Nolfi, “Social learning and cultural evolution in embodied and situated agents,” in *2007 IEEE Symposium on Artificial Life*, IEEE, 2007, pp. 333–340.

[59] L. Berthouze, C. Prince, M. Littman, H. Kozima, and C. Balkenius, “Adaptation and social facilitation in a population of autonomous robots,”

[60] N. Le, A. Brabazon, and M. O’Neill, “Social learning vs self-teaching in a multi-agent neural network system,” in *Applications of Evolutionary Computation*, P. A. Castillo, J. L. Jiménez Laredo, and F. Fernández de Vega, Eds., Springer International Publishing, 2020, pp. 354–368.

[61] N. Le, M. O'Neill, and A. Brabazon, "Adaptive advantage of learning strategies: A study through dynamic landscape," in *Parallel Problem Solving from Nature – PPSN XV*, A. Auger, C. M. Fonseca, N. Lourenço, P. Machado, L. Paquete, and D. Whitley, Eds., Springer International Publishing, 2018, pp. 387–398.

[62] H. Yonenoh, R. Suzuki, and T. Arita, "The effects of individual and social learning on the evolution of cognitive and communicative aspects of language abilities," in *Artificial Life Conference Proceedings*, MIT Press, 2019, pp. 541–542.

[63] C. Marriott and J. Chebib, "The effect of social learning on individual learning," in *Artificial Life Conference Proceedings 14*, MIT Press, 2014, pp. 736–743.

[64] J. Heinerman, B. Bussmann, R. Groenendijk, E. Van Krieken, J. Slik, A. Tezza, E. Haasdijk, and A. Eiben, "Benefits of social learning in physical robots," in *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, IEEE, 2018, pp. 851–858.

[65] A. Bartoli, M. Catto, A. De Lorenzo, E. Medvet, and J. Talamini, "Mechanisms of social learning in evolved artificial life," in *Artificial Life Conference Proceedings*, MIT Press, 2020, pp. 190–198.

[66] S. Cao, N. G. MacLaren, Y. Cao, Y. Dong, H. Sayama, F. J. Yammarino, S. D. Dionne, M. D. Mumford, S. Connelly, R. Martin, *et al.*, "An agent-based model of leader emergence and leadership perception within a collective," *Complexity*, vol. 2020, 2020.

[67] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Advances in Complex Systems*, vol. 3, no. 01n04, pp. 87–98, 2000.

[68] R. Hegselmann, U. Krause, *et al.*, "Opinion dynamics and bounded confidence models, analysis, and simulation," *Journal of artificial societies and social simulation*, vol. 5, no. 3, 2002.

[69] H. Sayama, "Extreme ideas emerging from social conformity and homophily: An adaptive social network model," in *Artificial Life Conference Proceedings*, MIT Press, 2020, pp. 113–120.

[70] B. E. Juel, R. Comolatti, G. Tononi, and L. Albantakis, "When is an action caused from within? quantifying the causal chain leading to actions in simulated agents," in *Artificial Life Conference Proceedings*, MIT Press, 2019, pp. 477–484.

[71] A. Covert III, S. McFetridge, and E. DeLord, "Structured populations with limited resources exhibit higher rates of complex function evolution," in *Artificial Life Conference Proceedings 14*, MIT Press, 2014, pp. 129–134.

[72] C. K. Hemelrijk, P. M. Kappeler, and I. Puga-Gonzalez, *The self-organization of social complexity in group-living animals: Lessons from the DomWorld model.* Elsevier Academic Press, 2017.

[73] S. Attwood, L. Canamero, and R. te Boekhorst, "Simianworld-a study of social organisation using an artificial life model," in *Artificial Life Conference Proceedings 13*, MIT Press, 2013, pp. 633–640.

[74] O. Miramontes, D. Boyer, and G. Ramos-Fernandez, "Understanding fission-fusion dynamics in social animals through agent-based modelling," in *Proceedings of the Artificial Life Conference 2016 13*, MIT Press, 2016, pp. 338–339.

[75] I. Puga-Gonzalez, H. Hildenbrandt, and C. K. Hemelrijk, "Emergent patterns of social affiliation in primates, a model," *PLoS Comput Biol*, vol. 5, no. 12, e1000630, 2009.

[76] C. Sueur, J.-L. Deneubourg, O. Petit, and I. D. Couzin, "Group size, grooming and fission in primates: A modeling approach based on group structure," *Journal of Theoretical Biology*, vol. 273, no. 1, pp. 156–166, 2011.

[77] R. Dolado, F. S. Beltran, and V. Quera, "Modeling the emergence of social structure from a phylogenetic point of view," *Journal of Artificial Societies and Social Simulation*, vol. 17, no. 1, p. 8, 2014.

[78] P. Rajagopalan, K. E. Holekamp, and R. Miikkulainen, "Factors that affect the evolution of complex cooperative behavior," in *The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, MIT Press, 2019, pp. 333–340.

[79] P. Andras, "Social learning, environmental adversity and the evolution of cooperation," in *Artificial Life Conference Proceedings 13*, MIT Press, 2016, pp. 290–297.

[80] M. Solomon, T. Soule, and R. B. Heckendorn, "A comparison of a communication strategies in cooperative learning," in *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, 2012, pp. 153–160.

[81] J. Fairey and T. Soule, "Evolution of communication and cooperation," in *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, 2014, pp. 169–176.

[82] C. Adam, R. Canal, B. Gaudou, H. T. Vinh, P. Taillandier, *et al.*, "Simulation of the emotion dynamics in a group of agents in an evacuation situation," in *International Conference on Principles and Practice of Multi-Agent Systems*, Springer, 2010, pp. 604–619.

[83] L. Luo, S. Zhou, W. Cai, M. Y. H. Low, F. Tian, Y. Wang, X. Xiao, and D. Chen, "Agent-based human behavior modeling for crowd simulation," *Computer Animation and Virtual Worlds*, vol. 19, no. 3-4, pp. 271–281, 2008.

[84] A. Zoumpoulaki, N. Avradinis, and S. Vosinakis, "A multi-agent simulation framework for emergency evacuations incorporating personality and emotions," in *Hellenic conference on artificial intelligence*, Springer, 2010, pp. 423–428.

[85] E. Oliveira and L. Sarmento, "Emotional advantage for adaptability and autonomy," in *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, 2003, pp. 305–312.

[86] R. Santos, G. Marreiros, C. Ramos, J. Neves, and J. Bulas-Cruz, "Multi-agent approach for ubiquitous group decision support involving emotions," in *International Conference on Ubiquitous Intelligence and Computing*, Springer, 2006, pp. 1174–1185.

[87] G. Marreiros, R. Santos, C. Ramos, and J. Neves, "Context-aware emotion-based model for group decision making," *IEEE Intelligent Systems*, vol. 25, no. 2, pp. 31–39, 2010.

[88] R. F. Antunes and N. Magnenat-Thalmann, "Bio-inspired virtual populations: Adaptive behavior with affective feedback," in *Proceedings of the 29th International Conference on Computer Animation and Social Agents*, 2016, pp. 101–110.

[89] R. Santos, G. Marreiros, C. Ramos, J. Neves, and J. Bulas-Cruz, "Personality, emotion, and mood in agent-based group decision making," *IEEE Intelligent Systems*, vol. 26, no. 6, pp. 58–66, 2011.

[90] J. Collenette, K. Atkinson, D. Bloembergen, and K. Tuyls, "Stability of cooperation in societies of emotional and moody agents," in *Artificial Life Conference Proceedings*, MIT Press, 2019, pp. 467–474.

[91] M. Lloyd-Kelly, K. Atkinson, and T. Bench-Capon, "Fostering co-operative behaviour through social intervention," in *2014 4th International Conference On Simulation And Modeling Methodologies, Technologies And Applications (SIMULTECH)*, IEEE, 2014, pp. 578–585.

[92] J. Taverner, B. Alfonso, E. Vivancos, and V. J. Botti, "Modeling personality in the affective agent architecture genia3.," in *ICAART (1)*, 2018, pp. 236–243.

[93] J. D. Jung, J. Hoey, J. H. Morgan, T. Schröder, and I. Wolf, "Grounding social interaction with affective intelligence," in *Canadian Conference on Artificial Intelligence*, Springer, 2016, pp. 52–57.

[94] M. Y. Lim, J. Dias, R. Aylett, and A. Paiva, "Creating adaptive affective autonomous npcs," *Autonomous Agents and Multi-Agent Systems*, vol. 24, no. 2, pp. 287–311, 2012.

[95] J. Dias, S. Mascarenhas, and A. Paiva, "Fatima modular: Towards an agent architecture with a generic appraisal framework," in *Emotion modeling*, Springer, 2014, pp. 44–56.

[96] I. Rausch, P. Simoens, and Y. Khaluf, "Adaptive foraging in dynamic environments using scale-free interaction networks," *Frontiers in Robotics and AI*, vol. 7, p. 86, 2020.

[97] X. Gu and T. FitzGerald, "Interoceptive inference: Homeostasis and decision-making," *Trends Cogn Sci*, vol. 18, no. 6, pp. 269–70, 2014.

[98] D. Vernon, "Reconciling constitutive and behavioural autonomy. the challenge of modelling development in enactive cognition," *Intellectica*, vol. 65, no. 1, pp. 63–79, 2016.

[99] R. Lowe, "The role of allostasis in sense-making: A better fit for interactivity than cybernetic-enactivism?," 2016.

[100] A. K. Seth and K. J. Friston, "Active interoceptive inference and the emotional brain," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 371, no. 1708, p. 20 160 007, 2016.

[101] M. Sanchez-Fibla, U. Bernardet, E. Wasserman, T. Pelc, M. Mintz, J. C. Jackson, C. Lansink, C. Pennartz, and P. F. Verschure, "Allostatic control for robot behavior regulation: A comparative rodent-robot study," *Advances in Complex Systems*, vol. 13, no. 03, pp. 377–403, 2010.

[102] V. Vouloutsi, S. Lallée, and P. F. Verschure, "Modulating behaviors using allostatic control," in *Conference on Biomimetic and Biohybrid Systems*, Springer, 2013, pp. 287–298.

[103] C. Moulin-Frier, J.-Y. Puigbo, X. D. Arsiwalla, M. Sanchez-Fibla, and P. F. Verschure, "Embodied artificial intelligence through distributed adaptive control: An integrated framework," in *2017 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, IEEE, 2017, pp. 324–330.

[104] J. Hoey, T. Schröder, J. Morgan, K. B. Rogers, D. Rishi, and M. Nagappan, "Artificial intelligence and social simulation: Studying group dynamics on a massive scale," *Small Group Research*, vol. 49, no. 6, pp. 647–683, 2018.

[105] A. J. Blanchard and L. Canamero, "Anticipating rewards in continuous time and space: A case study in developmental robotics," in *Workshop on Anticipatory Behavior in Adaptive Learning Systems*, Springer, 2006, pp. 267–284.

[106] S. G. Shamay-Tsoory and A. Abu-Akel, “The social salience hypothesis of oxytocin,” *Biological psychiatry*, vol. 79, no. 3, pp. 194–202, 2016.

[107] R. B. Ebitz, K. K. Watson, and M. L. Platt, “Oxytocin blunts social vigilance in the rhesus macaque,” *Proceedings of the National Academy of Sciences*, vol. 110, no. 28, pp. 11 630–11 635, 2013.

[108] R. M. Sapolsky, “The influence of social hierarchy on primate health,” *Science*, vol. 308, no. 5722, pp. 648–652, 2005.

[109] C. Crockford, T. Deschner, T. E. Ziegler, and R. M. Wittig, “Endogenous peripheral oxytocin measures can give insight into the dynamics of social relationships: A review,” *Frontiers in behavioral neuroscience*, vol. 8, p. 68, 2014.

[110] R. A. Bethlehem, S. Baron-Cohen, J. van Honk, B. Auyeung, and P. A. Bos, “The oxytocin paradox,” *Frontiers in behavioral neuroscience*, vol. 8, p. 48, 2014.

[111] E. T. Terris, L. E. Beavin, J. A. Barraza, J. Schloss, and P. J. Zak, “Endogenous oxytocin release eliminates in-group bias in monetary transfers with perspective-taking,” *Frontiers in behavioral neuroscience*, vol. 12, p. 35, 2018.

[112] A. Mielke, C. Crockford, and R. M. Wittig, “Predictability and variability of association patterns in sooty mangabeys,” *Behavioral ecology and sociobiology*, vol. 74, no. 4, pp. 1–13, 2020.

[113] K. Macdonald and D. Feifel, “Helping oxytocin deliver: Considerations in the development of oxytocin-based therapeutics for brain disorders,” *Frontiers in neuroscience*, vol. 7, p. 35, 2013.

[114] U. Wilensky, “Netlogo 5.3.1,” 2016.

[115] R. M. Sapolsky, “Individual differences and the stress response,” *Seminars in Neuroscience*, vol. 6, no. 4, pp. 261–269, 1994.

[116] R. M. Sapolsky, L. M. Romero, and A. U. Munck, “How do glucocorticoids influence stress responses? integrating permissive, suppressive, stimulatory, and preparative actions,” *Endocrine reviews*, vol. 21, no. 1, pp. 55–89, 2000.

[117] D. H. Abbott, E. B. Keverne, F. B. Bercovitch, C. A. Shively, S. P. Mendoza, W. Saltzman, C. T. Snowdon, T. E. Ziegler, M. Banjevic, T. Garland Jr, *et al.*, “Are subordinates always stressed? a comparative analysis of rank differences in cortisol levels among primates,” *Hormones and behavior*, vol. 43, no. 1, pp. 67–82, 2003.

[118] M. Porfiri, Y. Yang, R. J. Clement, and S. Ghirlanda, “A comparison of individual learning and social learning in zebrafish through an ethorobotics approach,” *Frontiers in Robotics and AI*, vol. 6, p. 71, 2019.

[119] Á. Miklósi, P. Korondi, V. Matellán, and M. Gácsi, “Ethorobotics: A new approach to human-robot relationship,” *Frontiers in psychology*, vol. 8, p. 958, 2017.

[120] R. K. Sawyer, “Interpreting and understanding simulations: The philosophy of social simulation,” in *Simulating social complexity*, Springer, 2013, pp. 273–289.

[121] B. Webb, “Animals versus animats: Or why not model the real iguana?” *Adaptive Behavior*, vol. 17, no. 4, pp. 269–286, 2009.

[122] P. Sterling, "Principles of allostasis: Optimal design, predictive regulation, pathophysiology, and rational," *Allostasis, homeostasis, and the costs of physiological adaptation*, vol. 17, 2004.

[123] C. A. Pedersen and A. J. Prange, "Induction of maternal behavior in virgin rats after intracerebroventricular administration of oxytocin," *Proceedings of the National Academy of Sciences*, vol. 76, no. 12, pp. 6661–6665, 1979.

[124] K. Kendrick, E. Keverne, M. Hinton, and J. Goode, "Cerebrospinal fluid and plasma concentrations of oxytocin and vasopressin during parturition and vaginocervical stimulation in the sheep," *Brain research bulletin*, vol. 26, no. 5, pp. 803–807, 1991.

[125] M. Kosfeld, M. Heinrichs, P. J. Zak, U. Fischbacher, and E. Fehr, "Oxytocin increases trust in humans," *Nature*, vol. 435, no. 7042, pp. 673–676, 2005.

[126] S. W. Chang, J. W. Barter, R. B. Ebitz, K. K. Watson, and M. L. Platt, "Inhaled oxytocin amplifies both vicarious reinforcement and self reinforcement in rhesus macaques (macaca mulatta)," *Proceedings of the National Academy of Sciences*, vol. 109, no. 3, pp. 959–964, 2012.

[127] C. K. De Dreu and M. E. Kret, "Oxytocin conditions intergroup relations through upregulated in-group empathy, cooperation, conformity, and defense," *Biological psychiatry*, vol. 79, no. 3, pp. 165–173, 2016.

[128] J. R. Williams, T. R. Insel, C. R. Harbaugh, and C. S. Carter, "Oxytocin administered centrally facilitates formation of a partner preference in female prairie voles (microtus ochrogaster)," *Journal of neuroendocrinology*, vol. 6, no. 3, pp. 247–250, 1994.

[129] T. R. Insel and T. J. Hulihan, "A gender-specific mechanism for pair bonding: Oxytocin and partner preference formation in monogamous voles.," *Behavioral neuroscience*, vol. 109, no. 4, p. 782, 1995.

[130] C. K. De Dreu, L. L. Greer, M. J. Handgraaf, S. Shalvi, G. A. Van Kleef, M. Baas, F. S. Ten Velden, E. Van Dijk, and S. W. Feith, "The neuropeptide oxytocin regulates parochial altruism in intergroup conflict among humans," *Science*, vol. 328, no. 5984, pp. 1408–1411, 2010.

[131] I. F. Bielsky and L. J. Young, "Oxytocin, vasopressin, and social recognition in mammals," *Peptides*, vol. 25, no. 9, pp. 1565–1574, 2004.

[132] J. A. Bartz, J. Zaki, N. Bolger, and K. N. Ochsner, "Social effects of oxytocin in humans: Context and person matter," *Trends in cognitive sciences*, vol. 15, no. 7, pp. 301–309, 2011.

[133] D. Rooney, R. Neely, C. Cullen, C. Ennis, B. Sheridan, A. Atkinson, E. Trimble, and P. Bell, "The effect of cortisol on glucose/glucose-6-phosphate cycle activity and insulin action," *The Journal of Clinical Endocrinology & Metabolism*, vol. 77, no. 5, pp. 1180–1183, 1993.

[134] M. Akinola and W. B. Mendes, "Stress-induced cortisol facilitates threat-related decision making among police officers.," *Behavioral neuroscience*, vol. 126, no. 1, p. 167, 2012.

[135] T. W. Buchanan and W. R. Lovallo, "Enhanced memory for emotional material following stress-level cortisol treatment in humans," *Psychoneuroendocrinology*, vol. 26, no. 3, pp. 307–317, 2001.

[136] K. E. Hannibal and M. D. Bishop, "Chronic stress, cortisol dysfunction, and pain: A psychoneuroendocrine rationale for stress management in pain rehabilitation," *Physical therapy*, vol. 94, no. 12, pp. 1816–1825, 2014.

[137] J. Robyn, C. Koch, J. Montalto, A. Yong, G. Warne, and J. Batch, “Cushing’s syndrome in childhood and adolescence,” *Journal of paediatrics and child health*, vol. 33, no. 6, pp. 522–527, 1997.

[138] J. Nerup, “Addison’s disease–clinical studies. a report of 108 cases,” *European Journal of Endocrinology*, vol. 76, no. 1, pp. 127–141, 1974.

[139] M. Lewis and L. Canamero, “A robot model of stress-induced compulsive behavior,” in *Proc. 8th International Conference on Affective Computing & Intelligent Interaction (ACII 2019)*, IEEE, 2019.

[140] C. Bernard, *The c c on ph ’e name ‘e nes of life common to animals and v é g ’e rates*. Bailli ‘e re, 1879, vol. 2.

[141] R. D. Beer, “The dynamics of active categorical perception in an evolved model agent,” *Adaptive behavior*, vol. 11, no. 4, pp. 209–243, 2003.

[142] A. Mielke, A. Preis, L. Samuni, J. F. Gogarten, R. M. Wittig, and C. Crockford, “Flexible decision-making in grooming partner choice in sooty mangabeys and chimpanzees,” *Royal Society open science*, vol. 5, no. 7, p. 172 143, 2018.

[143] J. A. Bartz, J. Zaki, N. Bolger, and K. N. Ochsner, “Social effects of oxytocin in humans: Context and person matter,” *Trends in cognitive sciences*, vol. 15, no. 7, pp. 301–309, 2011.

[144] C. L. Nehaniv, K. Dautenhahn, and K. Dautenhahn, *Imitation in animals and artifacts*. MIT press, 2002.

[145] J. A. Bartz, J. Zaki, N. Bolger, and K. N. Ochsner, “Social effects of oxytocin in humans: Context and person matter,” *Trends in cognitive sciences*, vol. 15, no. 7, pp. 301–309, 2011.

[146] J. N. Ferguson, L. J. Young, E. F. Hearn, M. M. Matzuk, T. R. Insel, and J. T. Winslow, “Social amnesia in mice lacking the oxytocin gene,” *Nature genetics*, vol. 25, no. 3, pp. 284–288, 2000.

[147] R. Wood, P. Baxter, and T. Belpaeme, “A review of long-term memory in natural and synthetic systems,” *Adaptive Behavior*, vol. 20, no. 2, pp. 81–103, 2012.

[148] M. Y. Lim, R. Aylett, W. C. Ho, S. Enz, and P. Vargas, “A socially-aware memory for companion agents,” in *International Workshop on Intelligent Virtual Agents*, Springer, 2009, pp. 20–26.

[149] W. C. Ho, K. Dautenhahn, M. Y. Lim, P. A. Vargas, R. Aylett, and S. Enz, “An initial memory model for virtual and robot companions supporting migration and long-term interaction,” in *RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication*, IEEE, 2009, pp. 277–284.

[150] M. I. Ahmad, “An emotion and memory model for social robots: A long-term interaction,” PhD thesis, Western Sydney University (Australia), 2018.

[151] R. Feldman, A. Weller, O. Zagoory-Sharon, and A. Levine, “Evidence for a neuroendocrinological foundation of human affiliation: Plasma oxytocin levels across pregnancy and the postpartum period predict mother-infant bonding,” *Psychological Science*, vol. 18, no. 11, pp. 965–970, 2007.

[152] S. Yohanan and K. E. MacLean, “The role of affective touch in human-robot interaction: Human intent and expectations in touching the haptic creature,” *International Journal of Social Robotics*, vol. 4, no. 2, pp. 163–180, 2012.

[153] A. Gallace and C. Spence, "The science of interpersonal touch: An overview," *Neuroscience & Biobehavioral Reviews*, vol. 34, no. 2, pp. 246–259, 2010.

[154] M. J. Hertenstein, J. M. Verkamp, A. M. Kerestes, and R. M. Holmes, "The communicative functions of touch in humans, nonhuman primates, and rats: A review and synthesis of the empirical research," *Genetic, social, and general psychology monographs*, vol. 132, no. 1, pp. 5–94, 2006.

[155] F. G. Hall, D. E. Knapp, and K. F. Huemmrich, "Physically based classification and satellite mapping of biophysical characteristics in the southern boreal forest," *Journal of Geophysical Research: Atmospheres*, vol. 102, no. D24, pp. 29 567–29 580, 1997.

[156] S. E. Jones and A. E. Yarbrough, "A naturalistic study of the meanings of touch," *Communications Monographs*, vol. 52, no. 1, pp. 19–56, 1985.

[157] J. S. Odendaal and R. A. Meintjes, "Neurophysiological correlates of affiliative behaviour between humans and dogs," *The Veterinary Journal*, vol. 165, no. 3, pp. 296–301, 2003.

[158] T. Onaka, K. Ikeda, T. Yamashita, and K. Honda, "Facilitative role of endogenous oxytocin in noradrenaline release in the rat supraoptic nucleus," *European Journal of Neuroscience*, vol. 18, no. 11, pp. 3018–3026, 2003.

[159] D. Scheele, K. M. Kendrick, C. Khouri, E. Kretzer, T. E. Schläpfer, B. Stoffel-Wagner, O. Güntürkün, W. Maier, and R. Hurlemann, "An oxytocin-induced facilitation of neural and emotional responses to social touch correlates inversely with autism traits," *Neuropsychopharmacology*, vol. 39, no. 9, pp. 2078–2085, 2014.

[160] D.-M. Ellingsen, J. Wessberg, O. Chelnokova, H. Olausson, B. Laeng, and S. Leknes, "In touch with your emotions: Oxytocin and touch change social impressions while others' facial expressions can alter touch," *Psychoneuroendocrinology*, vol. 39, pp. 11–20, 2014.

[161] I. Morrison, "Keep calm and cuddle on: Social touch as a stress buffer," *Adaptive Human Behavior and Physiology*, vol. 2, pp. 344–362, 2016.

[162] N. Striepens, D. Scheele, K. M. Kendrick, B. Becker, L. Schäfer, K. Schwalba, J. Reul, W. Maier, and R. Hurlemann, "Oxytocin facilitates protective responses to aversive social stimuli in males," *Proceedings of the National Academy of Sciences*, vol. 109, no. 44, pp. 18 144–18 149, 2012.

[163] N. Striepens, A. Matusch, K. M. Kendrick, Y. Mihov, D. Elmenhorst, B. Becker, M. Lang, H. H. Coenen, W. Maier, R. Hurlemann, *et al.*, "Oxytocin enhances attractiveness of unfamiliar female faces independent of the dopamine reward system," *Psychoneuroendocrinology*, vol. 39, pp. 74–87, 2014.

[164] D. Scheele, N. Striepens, O. Güntürkün, S. Deutschländer, W. Maier, K. M. Kendrick, and R. Hurlemann, "Oxytocin modulates social distance between males and females," *Journal of Neuroscience*, vol. 32, no. 46, pp. 16 074–16 079, 2012.

[165] S. Yao, B. Becker, W. Zhao, Z. Zhao, J. Kou, X. Ma, Y. Geng, P. Ren, and K. M. Kendrick, "Oxytocin modulates attention switching between interoceptive signals and external social cues," *Neuropsychopharmacology*, vol. 43, no. 2, pp. 294–301, 2018.

[166] E. Quattrocki and K. Friston, "Autism, oxytocin and interoception," *Neuroscience & Biobehavioral Reviews*, vol. 47, pp. 410–430, 2014.

[167] E. L. Knight and P. H. Mehta, "Hormones and hierarchies," in *The psychology of social status*, Springer, 2014, pp. 269–301.

[168] R. B. Ebitz and M. M. Platt, "An evolutionary perspective on the behavioral consequences of exogenous oxytocin application," *Frontiers in behavioral neuroscience*, vol. 7, p. 225, 2014.

[169] S. E. Muroy, K. L. Long, D. Kaufer, and E. D. Kirby, "Moderate stress-induced social bonding and oxytocin signaling are disrupted by predator odor in male rats," *Neuropsychopharmacology*, vol. 41, no. 8, pp. 2160–2170, 2016.

[170] N. Watanabe and M. Yamamoto, "Neural mechanisms of social dominance," *Frontiers in neuroscience*, vol. 9, p. 154, 2015.

[171] M. Marmot and J. J. Allen, *Social determinants of health equity*, 2014.

[172] E. A. Archie, J. Altmann, and S. C. Alberts, "Social status predicts wound healing in wild baboons," *Proceedings of the National Academy of Sciences*, vol. 109, no. 23, pp. 9017–9022, 2012.

[173] V. Roubová, M. Konečná, P. Šmilauer, and B. Wallner, "Whom to groom and for what? patterns of grooming in female barbary macaques (macaca sylvanus)," *PloS one*, vol. 10, no. 2, 2015.

[174] K. D. Broad, J. P. Curley, and E. B. Keverne, "Mother–infant bonding and the evolution of mammalian social relationships," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 361, no. 1476, pp. 2199–2214, 2006.

[175] C. Carter and E. B. Keverne, "The neurobiology of social affiliation and pair bonding," in *Hormones, brain and behavior*, Elsevier, 2002, pp. 299–337.

[176] C. Crockford, T. Deschner, and R. M. Wittig, "The role of oxytocin in social buffering: What do primate studies add?" In *Behavioral Pharmacology of Neuropeptides: Oxytocin*, Springer, 2017, pp. 155–173.

[177] Y. Liu and Z. Wang, "Nucleus accumbens oxytocin and dopamine interact to regulate pair bond formation in female prairie voles," *Neuroscience*, vol. 121, no. 3, pp. 537–544, 2003.

[178] E. Choleris, A. E. Clipperton-Allen, A. Phan, and M. Kavaliers, "Neuroendocrinology of social information processing in rats and mice," *Frontiers in Neuroendocrinology*, vol. 30, no. 4, pp. 442–459, 2009.

[179] T. Kikusui, J. T. Winslow, and Y. Mori, "Social buffering: Relief from stress and anxiety," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 361, no. 1476, pp. 2215–2228, 2006.

[180] L. Samuni, A. Preis, A. Mielke, T. Deschner, R. M. Wittig, and C. Crockford, "Social bonds facilitate cooperative resource sharing in wild chimpanzees," *Proceedings of the Royal Society B*, vol. 285, no. 1888, p. 20 181 643, 2018.

[181] L. Kulik, "Development and consequences of social behavior in rhesus macaques (macaca mulatta)," PhD thesis, Univ. Leipzig, 2015.

[182] J. B. Silk, S. F. Brosnan, J. Henrich, S. P. Lambeth, and S. Shapiro, "Chimpanzees share food for many reasons: The role of kinship, reciprocity, social bonds and harassment on food transfers," *Animal behaviour*, vol. 85, no. 5, pp. 941–947, 2013.

[183] J. Silk, D. Cheney, and R. Seyfarth, “A practical guide to the study of social relationships,” *Evolutionary Anthropology: Issues, News, and Reviews*, vol. 22, no. 5, pp. 213–225, 2013.

[184] M. Takano and K. Nakazato, “Emergence of the tradeoff law of social relationships in artificial societies driven by dual memory mechanisms,” in *Artificial Life Conference Proceedings*, MIT Press, 2020, pp. 87–94.

[185] M. Takano, “Two types of social grooming methods depending on the trade-off between the number and strength of social relationships,” *Royal Society open science*, vol. 5, no. 8, p. 180 148, 2018.

[186] A. S. Goffe, J. Fischer, and H. Sennhenn-Reulen, “Bayesian inference and simulation approaches improve the assessment of elo-scores in the analysis of social behaviour,” *bioRxiv*, p. 160 671, 2017.

[187] N. E. Newton-Fisher, “Modeling social dominance: Elo-ratings, prior history, and the intensity of aggression,” *International journal of primatology*, vol. 38, no. 3, pp. 427–447, 2017.

[188] R. M. Sapolsky, S. C. Alberts, and J. Altmann, “Hypercortisolism associated with social subordinance or social isolation among wild baboons,” *Archives of general psychiatry*, vol. 54, no. 12, pp. 1137–1143, 1997.

[189] M. L. Platt, R. M. Seyfarth, and D. L. Cheney, “Adaptations for social cognition in the primate brain,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 371, no. 1687, p. 20 150 096, 2016.

[190] R. M. Seyfarth and D. L. Cheney, “Social cognition,” *Animal Behaviour*, vol. 103, pp. 191–202, 2015.

[191] C. Crockford, R. M. Wittig, K. Langergraber, T. E. Ziegler, K. Zuberbühler, and T. Deschner, “Urinary oxytocin and social bonding in related and unrelated wild chimpanzees,” *Proceedings of the Royal Society B: Biological Sciences*, vol. 280, no. 1755, p. 20 122 765, 2013.

[192] T. R. Insel and T. J. Hulihan, “A gender-specific mechanism for pair bonding: Oxytocin and partner preference formation in monogamous voles.,” *Behavioral neuroscience*, vol. 109, no. 4, p. 782, 1995.

[193] F. B. Bercovitch and S. C. Strum, “Dominance rank, resource availability, and reproductive maturation in female savanna baboons,” *Behavioral Ecology and Sociobiology*, vol. 33, no. 5, pp. 313–318, 1993.

[194] D. S. Albeck, C. R. McKittrick, D. C. Blanchard, R. J. Blanchard, J. Nikulina, B. S. McEwen, and R. R. Sakai, “Chronic social stress alters levels of corticotropin-releasing factor and arginine vasopressin mrna in rat brain,” *Journal of Neuroscience*, vol. 17, no. 12, pp. 4895–4903, 1997.

[195] C. E. Virgin Jr and R. M. Sapolsky, “Styles of male social behavior and their endocrine correlates among low-ranking baboons,” *American Journal of Primatology*, vol. 42, no. 1, pp. 25–39, 1997.

[196] S. Creel, “Social dominance and stress hormones,” *Trends in ecology & evolution*, vol. 16, no. 9, pp. 491–497, 2001.

[197] B. W. Roberts, N. R. Kuncel, R. Shiner, A. Caspi, and L. R. Goldberg, “The power of personality: The comparative validity of personality traits, socioeconomic status, and cognitive ability for predicting important life outcomes,” *Perspectives on Psychological science*, vol. 2, no. 4, pp. 313–345, 2007.

[198] A. Bartolomucci, P. Palanza, P. Sacerdote, A. E. Panerai, A. Sgoifo, R. Dantzer, and S. Parmigiani, “Social factors and individual vulnerability to chronic stress exposure,” *Neuroscience & Biobehavioral Reviews*, vol. 29, no. 1, pp. 67–81, 2005.

[199] C. T. Wolff, S. Friedman, M. Hofer, and J. Mason, "Relationship between psychological defenses and mean urinary 17-hydroxycorticosteroid excretion rates," *emotion*, vol. 2, p. 6, 1964.

[200] I. Kyrou and C. Tsigos, "Stress hormones: Physiological stress and regulation of metabolism," *Current opinion in pharmacology*, vol. 9, no. 6, pp. 787–793, 2009.

[201] B. S. McEwen, N. P. Bowles, J. D. Gray, M. N. Hill, R. G. Hunter, I. N. Karatsoreos, and C. Nasca, "Mechanisms of stress in the brain," *Nature neuroscience*, vol. 18, no. 10, pp. 1353–1363, 2015.

[202] M. B. Hennessy, S. Kaiser, and N. Sachser, "Social buffering of the stress response: Diversity, mechanisms, and functions," *Frontiers in neuroendocrinology*, vol. 30, no. 4, pp. 470–482, 2009.

[203] L. M. Glynn, N. Christenfeld, and W. Gerin, "Gender, social support, and cardiovascular responses to stress," *Psychosomatic medicine*, vol. 61, no. 2, pp. 234–242, 1999.

[204] T. W. Kamarck, S. B. Manuck, and J. R. Jennings, "Social support reduces cardiovascular reactivity to psychological challenge: A laboratory model.," *Psychosomatic medicine*, 1990.

[205] A. C. DeVries, M. B. DeVries, S. E. Taymans, and C. S. Carter, "The effects of stress on social preferences are sexually dimorphic in prairie voles," *Proceedings of the National Academy of Sciences*, vol. 93, no. 21, pp. 11 980–11 984, 1996.

[206] S. Levine, "Influence of psychological variables on the activity of the hypothalamic–pituitary–adrenal axis," *European journal of pharmacology*, vol. 405, no. 1-3, pp. 149–160, 2000.

[207] M. Heinrichs, T. Baumgartner, C. Kirschbaum, and U. Ehlert, "Social support and oxytocin interact to suppress cortisol and subjective responses to psychosocial stress," *Biological psychiatry*, vol. 54, no. 12, pp. 1389–1398, 2003.

[208] C. Kirschbaum, K.-M. Pirke, and D. H. Hellhammer, "The 'trier social stress test'–a tool for investigating psychobiological stress responses in a laboratory setting," *Neuropsychobiology*, vol. 28, no. 1-2, pp. 76–81, 1993.

[209] I. Neumann, H. Johnstone, M. Hatzinger, G. Liebsch, M. Shipston, J. Russell, R. Landgraf, and A. Douglas, "Attenuated neuroendocrine responses to emotional and physical stressors in pregnant rats involve adenohypophysial changes," *The Journal of physiology*, vol. 508, no. 1, pp. 289–300, 1998.

[210] R. Windle, S. Wood, N. Shanks, P. Perks, G. Conde, A. Da Costa, C. Ingram, and S. Lightman, "Endocrine and behavioural responses to noise stress: Comparison of virgin and lactating female ratsduring non-disrupted maternal activity," *Journal of neuroendocrinology*, vol. 9, no. 6, pp. 407–414, 1997.

[211] J. Born, T. Lange, W. Kern, G. P. McGregor, U. Bickel, and H. L. Fehm, "Sniffing neuropeptides: A transnasal approach to the human brain," *Nature neuroscience*, vol. 5, no. 6, pp. 514–516, 2002.

[212] M. B. Hennessy, C. K. Nigh, M. L. Sims, and S. J. Long, "Plasma cortisol and vocalization responses of postweaning age guinea pigs to maternal and sibling separation: Evidence for filial attachment after weaning," *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology*, vol. 28, no. 2, pp. 103–115, 1995.

[213] K. M. Kendrick, A. P. Da Costa, K. D. Broad, S. Ohkura, R. Guevara, F. Lévy, and E. B. Keverne, "Neural control of maternal behaviour and olfactory recognition of offspring," *Brain research bulletin*, vol. 44, no. 4, pp. 383–395, 1997.

[214] I. D. Neumann, "Involvement of the brain oxytocin system in stress coping: Interactions with the hypothalamo-pituitary-adrenal axis," in *Progress in Brain Research*, vol. 139, Elsevier, 2002, pp. 147–162.

[215] K. Ebner, O. J. Bosch, S. A. Krömer, N. Singewald, and I. D. Neumann, "Release of oxytocin in the rat central amygdala modulates stress-coping behavior and the release of excitatory amino acids," *Neuropsychopharmacology*, vol. 30, no. 2, pp. 223–230, 2005.

[216] R. Fleming, A. Baum, M. M. Gisriel, and R. J. Gatchel, "Mediating influences of social support on stress at three mile island," *Journal of Human stress*, vol. 8, no. 3, pp. 14–23, 1982.

[217] T. E. Seeman, L. F. Berkman, D. Blazer, and J. W. Rowe, "Social ties and support and neuroendocrine function: The macarthur studies of successful aging," *Annals of Behavioral Medicine*, vol. 16, no. 2, pp. 95–106, 1994.

[218] B. Ditzen and M. Heinrichs, "Psychobiology of social support: The social dimension of stress buffering," *Restorative neurology and neuroscience*, vol. 32, no. 1, pp. 149–162, 2014.

[219] D. Uno, B. N. Uchino, and T. W. Smith, "Relationship quality moderates the effect of social support given by close friends on cardiovascular reactivity in women," *International Journal of Behavioral Medicine*, vol. 9, no. 3, pp. 243–262, 2002.

[220] T. K. Inagaki and N. I. Eisenberger, "Giving support to others reduces sympathetic nervous system-related responses to stress," *Psychophysiology*, vol. 53, no. 4, pp. 427–435, 2016.

[221] R. L. Piferi and K. A. Lawler, "Social support and ambulatory blood pressure: An examination of both receiving and giving," *International Journal of Psychophysiology*, vol. 62, no. 2, pp. 328–336, 2006.

[222] F. Ozbay, D. C. Johnson, E. Dimoulas, C. Morgan III, D. Charney, and S. Southwick, "Social support and resilience to stress: From neurobiology to clinical practice," *Psychiatry (Edgmont)*, vol. 4, no. 5, p. 35, 2007.

[223] C. Crockford, R. M. Wittig, P. L. Whitten, R. M. Seyfarth, and D. L. Cheney, "Social stressors and coping mechanisms in wild female baboons (papio hamadryas ursinus)," *Hormones and Behavior*, vol. 53, no. 1, pp. 254–265, 2008.

[224] S. Cohen, R. Mermelstein, T. Kamarck, and H. M. Hoberman, "Measuring the functional components of social support," in *Social support: Theory, research and applications*, Springer, 1985, pp. 73–94.

[225] N. Bohnen, P. Houx, N. Nicolson, and J. Jolles, "Cortisol reactivity and cognitive performance in a continuous mental task paradigm," *Biological Psychology*, vol. 31, no. 2, pp. 107–116, 1990.

[226] C. S. Carter, "Neuroendocrine perspectives on social attachment and love," *Psychoneuroendocrinology*, vol. 23, no. 8, pp. 779–818, 1998.

[227] J. Panksepp, "Oxytocin effects on emotional processes: Separation distress, social bonding, and relationships to psychiatric disorders," *Annals of the New York Academy of Sciences*, vol. 652, no. 1, pp. 243–252, 1992.

[228] E. E. Nelson and J. Panksepp, "Brain substrates of infant–mother attachment: Contributions of opioids, oxytocin, and norepinephrine," *Neuroscience & Biobehavioral Reviews*, vol. 22, no. 3, pp. 437–452, 1998.

[229] T. R. Insel and L. J. Young, "The neurobiology of attachment," *Nature Reviews Neuroscience*, vol. 2, no. 2, pp. 129–136, 2001.

[230] B. S. Cushing and C. S. Carter, "Peripheral pulses of oxytocin increase partner preferences in female, but not male, prairie voles," *Hormones and behavior*, vol. 37, no. 1, pp. 49–56, 2000.

[231] J. Galgani and E. Ravussin, "Energy metabolism, fuel selection and body weight regulation," *International journal of obesity*, vol. 32, no. 7, S109–S119, 2008.

[232] M. R. Gunnar and C. E. Hostinar, "The social buffering of the hypothalamic–pituitary–adrenocortical axis in humans: Developmental and experiential determinants," *Social neuroscience*, vol. 10, no. 5, pp. 479–488, 2015.

[233] C. S. Carter, "Oxytocin pathways and the evolution of human behavior," *Annual review of psychology*, vol. 65, pp. 17–39, 2014.

[234] H. E. Ross and L. J. Young, "Oxytocin and the neural mechanisms regulating social cognition and affiliative behavior," *Frontiers in neuroendocrinology*, vol. 30, no. 4, pp. 534–547, 2009.

[235] R. M. Sullivan and R. E. Perry, "Mechanisms and functional implications of social buffering in infants: Lessons from animal models," *Social neuroscience*, vol. 10, no. 5, pp. 500–511, 2015.

[236] M. Nachmias, M. Gunnar, S. Mangelsdorf, R. H. Parritz, and K. Buss, "Behavioral inhibition and stress reactivity: The moderating role of attachment security," *Child development*, vol. 67, no. 2, pp. 508–522, 1996.

[237] M. M. Sanchez, "The impact of early adverse care on hpa axis development: Nonhuman primate models," *Hormones and behavior*, vol. 50, no. 4, pp. 623–631, 2006.

[238] K. D. Broad, M. L. Mimmack, E. B. Keverne, and K. M. Kendrick, "Increased bdnf and trk-b mrna expression in cortical and limbic regions following formation of a social recognition memory," *European Journal of Neuroscience*, vol. 16, no. 11, pp. 2166–2174, 2002.

[239] J. H. Wilson, "Prolactin in rats is attenuated by conspecific touch in a novel environment," *Cognitive, Affective, & Behavioral Neuroscience*, vol. 1, no. 2, pp. 199–205, 2001.

[240] M. Rukstalis and J. A. French, "Vocal buffering of the stress response: Exposure to conspecific vocalizations moderates urinary cortisol excretion in isolated marmosets," *Hormones and behavior*, vol. 47, no. 1, pp. 1–7, 2005.

[241] T. Kikusui, S. Takigami, Y. Takeuchi, and Y. Mori, "Alarm pheromone enhances stress-induced hyperthermia in rats," *Physiology & behavior*, vol. 72, no. 1-2, pp. 45–50, 2001.

[242] H. K. Laurent, "Clarifying the contours of emotion regulation: Insights from parent–child stress research," *Child Development Perspectives*, vol. 8, no. 1, pp. 30–35, 2014.

[243] C. R. Fox and D. M. Gelfand, "Maternal depressed mood and stress as related to vigilance, self-efficacy and mother-child interactions," *Early Development and Parenting*, vol. 3, no. 4, pp. 233–243, 1994.

[244] L. Meakins, L. Ray, K. Hegadoren, L. G. Rogers, and G. R. Rempel, "Parental vigilance in caring for their children with hypoplastic left heart syndrome.," *Pediatric nursing*, vol. 41, no. 1, 2015.

[245] G. Goodwillie, "Protective vigilance: A parental strategy in caring for a child diagnosed with adhd," *Journal of Family Therapy*, vol. 36, no. 3, pp. 255–267, 2014.

[246] W. Eriksen, “The role of social support in the pathogenesis of coronary heart disease. a literature review,” *Family Practice*, vol. 11, no. 2, pp. 201–209, 1994.

[247] C. E. Hostinar, “Recent developments in the study of social relationships, stress responses, and physical health,” *Current opinion in psychology*, vol. 5, pp. 90–95, 2015.

[248] Y. J. Zanstra and D. W. Johnston, “Cardiovascular reactivity in real life settings: Measurement, mechanisms and meaning,” *Biological psychology*, vol. 86, no. 2, pp. 98–105, 2011.

[249] G. P. Moberg, “Biological response to stress: Implications for animal welfare,” *The biology of animal stress: basic principles and implications for animal welfare*, vol. 1, p. 21, 2000.

[250] B. S. McEwen and J. C. Wingfield, “The concept of allostasis in biology and biomedicine,” *Hormones and behavior*, vol. 43, no. 1, pp. 2–15, 2003.

[251] C. Rabasa and S. L. Dickson, “Impact of stress on metabolism and energy balance,” *Current Opinion in Behavioral Sciences*, vol. 9, pp. 71–77, 2016.

[252] G. P. Chrousos, “Stress and disorders of the stress system,” *Nature reviews endocrinology*, vol. 5, no. 7, p. 374, 2009.

[253] T. Elsasser, K. Klasing, N. Filipov, F. Thompson, *et al.*, “The metabolic consequences of stress: Targets for stress and priorities of nutrient use,” in *The biology of animal stress: basic principles and implications for animal welfare*, CAB International Oxon, UK, 2000, pp. 77–110.

[254] M. Kanarik, “Inter-individual differences in vulnerability to depression: Regional brain energy metabolism, serotonergic function and behaviour in animal models,” PhD thesis, 2011.

[255] C. I. Martins, M. F. Castanheira, S. Engrola, B. Costas, and L. E. Conceição, “Individual differences in metabolism predict coping styles in fish,” *Applied Animal Behaviour Science*, vol. 130, no. 3-4, pp. 135–143, 2011.

[256] E. Charmandari, C. Tsigos, and G. Chrousos, “Endocrinology of the stress response,” *Annu. Rev. Physiol.*, vol. 67, pp. 259–284, 2005.

[257] S. Cohen, “Social relationships and health.,” *American psychologist*, vol. 59, no. 8, p. 676, 2004.

[258] T. W. Kamarck, S. B. Manuck, and J. R. Jennings, “Social support reduces cardiovascular reactivity to psychological challenge: A laboratory model.,” *Psychosomatic medicine*, 1990.

[259] A.-L. Unden, K. Orth-Gomer, and S. Elofsson, “Cardiovascular effects of social support in the work place: Twenty-four-hour ecg monitoring of men and women.,” *Psychosomatic Medicine*, 1991.

[260] B. N. Uchino and T. S. Garvey, “The availability of social support reduces cardiovascular reactivity to acute psychological stress,” *Journal of behavioral medicine*, vol. 20, no. 1, pp. 15–27, 1997.

[261] S. J. Lepore, K. Allen, and G. W. Evans, “Social support lowers cardiovascular reactivity to an acute stressor.,” *Psychosomatic medicine*, vol. 55, no. 6, pp. 518–524, 1993.

[262] B. N. Uchino, K. Bowen, M. Carlisle, and W. Birmingham, “Psychological pathways linking social support to health outcomes: A visit with the “ghosts” of research past, present, and future,” *Social science & medicine*, vol. 74, no. 7, pp. 949–957, 2012.

[263] M. D. Seery, “The biopsychosocial model of challenge and threat: Using the heart to measure the mind,” *Social and Personality Psychology Compass*, vol. 7, no. 9, pp. 637–653, 2013.

[264] J. Skelton and J. Pennebaker, *The verbal system. u: Jt cacioppo, lg tassinary (ur.), principles of psychophysiology: Physical, social, and inferential elements (str. 631-657)*, 1990.

[265] S. Ahadi and E. Diener, "Multiple determinants and effect size.," *Journal of Personality and social Psychology*, vol. 56, no. 3, p. 398, 1989.

[266] T. Onaka, Y. Takayanagi, and M. Yoshida, "Roles of oxytocin neurones in the control of stress, energy metabolism, and social behaviour," *Journal of neuroendocrinology*, vol. 24, no. 4, pp. 587–598, 2012.

[267] J. Gutkowska and M. Jankowski, "Oxytocin revisited: Its role in cardiovascular regulation," *Journal of neuroendocrinology*, vol. 24, no. 4, pp. 599–608, 2012.

[268] E. A. Lawson, D. A. Marengi, R. L. DeSanti, T. M. Holmes, D. A. Schoenfeld, and C. J. Tolley, "Oxytocin reduces caloric intake in men," *Obesity*, vol. 23, no. 5, pp. 950–956, 2015.

[269] E. A. Lawson, "The effects of oxytocin on eating behaviour and metabolism in humans," *Nature Reviews Endocrinology*, vol. 13, no. 12, p. 700, 2017.

[270] X. Barandiaran and A. Moreno, "Adaptivity: From metabolism to behavior," *Adaptive Behavior*, vol. 16, no. 5, pp. 325–344, 2008.

[271] C. Melhuish, I. Ieropoulos, J. Greenman, and I. Horsfield, "Energetically autonomous robots: Food for thought," *Autonomous Robots*, vol. 21, no. 3, pp. 187–198, 2006.

[272] M. D. Egbert, X. E. Barandiaran, and E. A. Di Paolo, "A minimal model of metabolism-based chemotaxis," *PLoS Comput Biol*, vol. 6, no. 12, e1001004, 2010.

[273] K. Kiryazov and R. Lowe, "The role of arousal in embodying the cuexdeficit model in multi-resource human-robot interaction," in *Artificial Life Conference Proceedings 13*, MIT Press, 2013, pp. 19–26.

[274] B. Wróbel, M. Joachimczak, A. Montebelli, and R. Lowe, "The search for beauty: Evolution of minimal cognition in an animat controlled by a gene regulatory network and powered by a metabolic system," in *International Conference on Simulation of Adaptive Behavior*, Springer, 2012, pp. 198–208.

[275] T. Romero, M. Nagasawa, K. Mogi, T. Hasegawa, and T. Kikusui, "Oxytocin promotes social bonding in dogs," *Proceedings of the National Academy of Sciences*, vol. 111, no. 25, pp. 9085–9090, 2014.

[276] D. Wittfoth-Schardt, J. Gründing, M. Wittfoth, H. Lanfermann, M. Heinrichs, G. Domes, A. Buchheim, H. Gündel, and C. Waller, "Oxytocin modulates neural reactivity to children's faces as a function of social salience," *Neuropsychopharmacology*, vol. 37, no. 8, pp. 1799–1807, 2012.

[277] H. IJzerman, J. A. Coan, F. Wagemans, M. A. Missler, I. v. Beest, S. Lindenberg, and M. Tops, "A theory of social thermoregulation in human primates," *Frontiers in psychology*, vol. 6, p. 464, 2015.

[278] L. A. Campbell, P. J. Tkaczynski, J. Lehmann, M. Mouna, and B. Majolo, "Social thermoregulation as a potential mechanism linking sociality and fitness: Barbary macaques with more social partners form larger huddles," *Scientific reports*, vol. 8, no. 1, pp. 1–8, 2018.

[279] L. A. Ebensperger, J. Ramırez-Estrada, C. León, R. A. Castro, L. O. Tolhuysen, R. Sobrero, V. Quirici, J. R. Burger, M. Soto-Gamboa, and L. D. Hayes, "Sociality, glucocorticoids and direct fitness in the communally rearing rodent, octodon degus," *Hormones and Behavior*, vol. 60, no. 4, pp. 346–352, 2011.

[280] J. A. Coan, "Toward a neuroscience of attachment.," 2008.

[281] M. Anderson, “Neural reuse: A fundamental organizational principle of the brain,” *Behavioral and brain sciences*, vol. 33, no. 4, p. 245, 2010.

[282] T. C. Adam and E. S. Epel, “Stress, eating and the reward system,” *Physiology & behavior*, vol. 91, no. 4, pp. 449–458, 2007.

[283] Y. Maejima, Y. Iwasaki, Y. Yamahara, M. Kodaira, U. Sedbazar, and T. Yada, “Peripheral oxytocin treatment ameliorates obesity by reducing food intake and visceral fat mass,” *Aging (Albany NY)*, vol. 3, no. 12, p. 1169, 2011.

[284] V. Ott, G. Finlayson, H. Lehnert, B. Heitmann, M. Heinrichs, J. Born, and M. Hallschmid, “Oxytocin reduces reward-driven food intake in humans,” *Diabetes*, vol. 62, no. 10, pp. 3418–3425, 2013.

[285] T. Onaka and Y. Takayanagi, “Role of oxytocin in the control of stress and food intake,” *Journal of neuroendocrinology*, vol. 31, no. 3, e12700, 2019.

[286] N. E. Rowland and S. M. Antelman, “Stress-induced hyperphagia and obesity in rats: A possible model for understanding human obesity,” *Science*, vol. 191, no. 4224, pp. 310–312, 1976.

[287] M. B. Wallach, M. Dawber, M. McMahon, and C. Rogers, “A new anorexigen assay: Stress-induced hyperphagia in rats,” *Pharmacology Biochemistry and Behavior*, vol. 6, no. 5, pp. 529–531, 1977.

[288] K. L. Tamashiro, M. M. Nguyen, and R. R. Sakai, “Social stress: From rodents to primates,” *Frontiers in neuroendocrinology*, vol. 26, no. 1, pp. 27–40, 2005.

[289] M. S. Spetter, G. B. Feld, M. Thienel, H. Preissl, M. A. Hege, and M. Hallschmid, “Oxytocin curbs calorie intake via food-specific increases in the activity of brain areas that process reward and establish cognitive control,” *Scientific reports*, vol. 8, no. 1, pp. 1–11, 2018.

[290] J. F. Mattanah, F. G. Lopez, and J. M. Govern, “The contributions of parental attachment bonds to college student development and adjustment: A meta-analytic review.,” *Journal of Counseling Psychology*, vol. 58, no. 4, p. 565, 2011.

[291] J. Kandiah, M. Yake, and H. Willett, “Effects of stress on eating practices among adults,” *Family and Consumer Sciences Research Journal*, vol. 37, no. 1, pp. 27–38, 2008.

[292] R. E. Henriksen, T. Torsheim, and F. Thuen, “Loneliness, social integration and consumption of sugar-containing beverages: Testing the social baseline theory,” *PloS one*, vol. 9, no. 8, 2014.

[293] A. V. Jaeggi, J. M. Stevens, and C. P. Van Schaik, “Tolerant food sharing and reciprocity is precluded by despotism among bonobos but not chimpanzees,” *American Journal of Physical Anthropology*, vol. 143, no. 1, pp. 41–51, 2010.

[294] I. C. Gilby, “Meat sharing among the gombe chimpanzees: Harassment and reciprocal exchange,” *Animal Behaviour*, vol. 71, no. 4, pp. 953–963, 2006.

[295] N. G. Blurton Jones, “Tolerated theft, suggestions about the ecology and evolution of sharing, hoarding and scrounging,” *Information (International Social Science Council)*, vol. 26, no. 1, pp. 31–54, 1987.

APPENDIX

Appendix A

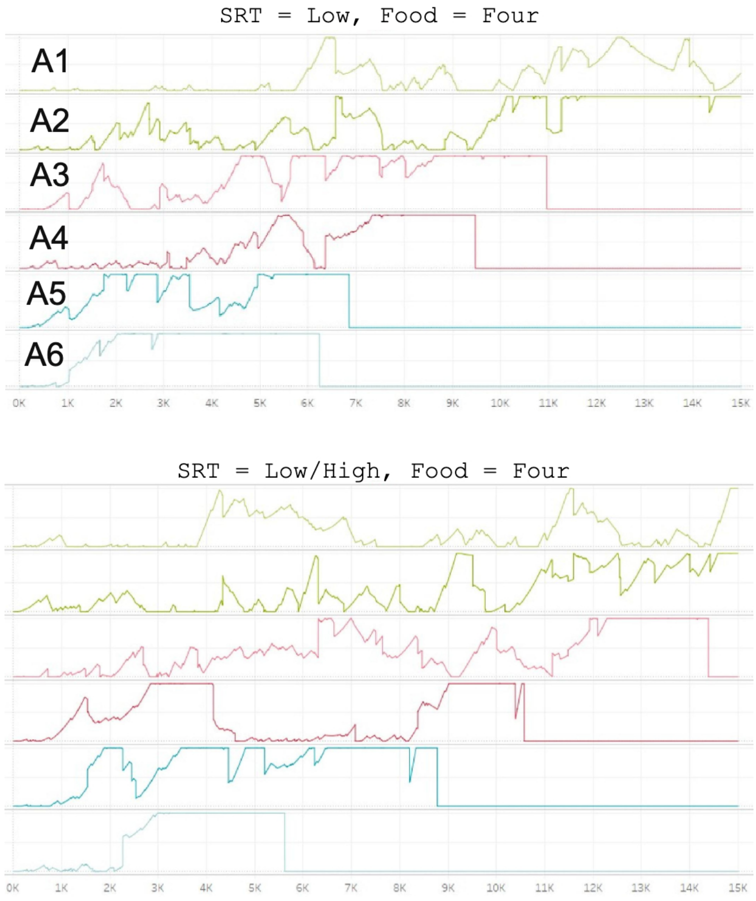
SRT = Low, Food = Four

SRT = Low/High, Food = Four

Figure A.1 Evolution of cortisol level trends for all agents in one simulation run in two *StressThreshold* conditions: Low (top) and Low→High (bottom). In both instances, cortisol levels of subordinate agents (A5 and A6) become elevated before dominant agents (A1, A2).

Appendix B

List of journal articles and conference proceedings contributions from the Ph. D. :

- Khan, I., Lewis, M. and Cañamero, L., 2018, April. Adaptation and the social salience hypothesis of oxytocin: Early experiments in a simulated agent environment. *In Proceedings of the 2nd Symposium on Social Interactions in Complex Intelligent Systems (SICIS)*, Liverpool, UK (pp. 4-6). (Chapter 3)
- Khan, I. and Cañamero, L., 2018. Modelling adaptation through social allostasis: Modulating the effects of social touch with oxytocin in embodied agents. *Emotions in Robots: Embodied Interaction in Social and Non-Social Environments, MTI Special Issue*, 2(4), p.67. (Chapter 4)
- Khan, I., Lewis, M. and Cañamero, L., 2019, September. The Effects of Affective Social Bonds on the Interactions and Survival of Simulated Agents. In *2019 8th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)* (pp. 374-380). IEEE. (Chapter 5)
- Khan, I., Lewis, M. and Cañamero, L., 2020. The Effects of Rank-Related Differences in Stress Tolerance on the Social Dynamics and Group Viability of Artificial Affective Agents. *IEEE Transactions on Affective Computing* (Chapter 6, submitted).
- Khan, I., Lewis, M. and Cañamero, L., 2020, July. Modelling the Social Buffering Hypothesis in an Artificial Life Environment. *In Artificial Life Conference Proceedings (pp. 393-401).* MIT Press. (Chapter 7)

Appendix C

Justification for Artificial Life Approach

Though this list is not exhaustive, an artificial life approach proposes a number of advantages for studying autonomous agents over physical agents (i.e. robots), particularly when studying multi-agent social systems:

- **Scalability and Cost:** Societies of artificial agents can be scaled easily and at no cost. A similar approach using physical agents would have additional financial implications (for instance, the purchase of additional agents).
- **Resource Requirements:** Once a model has been developed, simulations of artificial life models can be performed quickly and concurrently, with minimal computational resource. As computational resource increases, the number of simulations running concurrently can be increased.
- **Noise Reduction:** One of the difficulties in working with physical agents is the effect of extraneous variables.
- **Open-Access & Accessibility:** Artificial life models can be developed with an accessible front-end graphical user interface, allowing non-domain experts to use such models with a minimal learning curve.

Appendix D

Trended results for Life Length and Comfort across all values of *SocialSalienceDirection* for Experiment 1 and 2 in Chapter 4

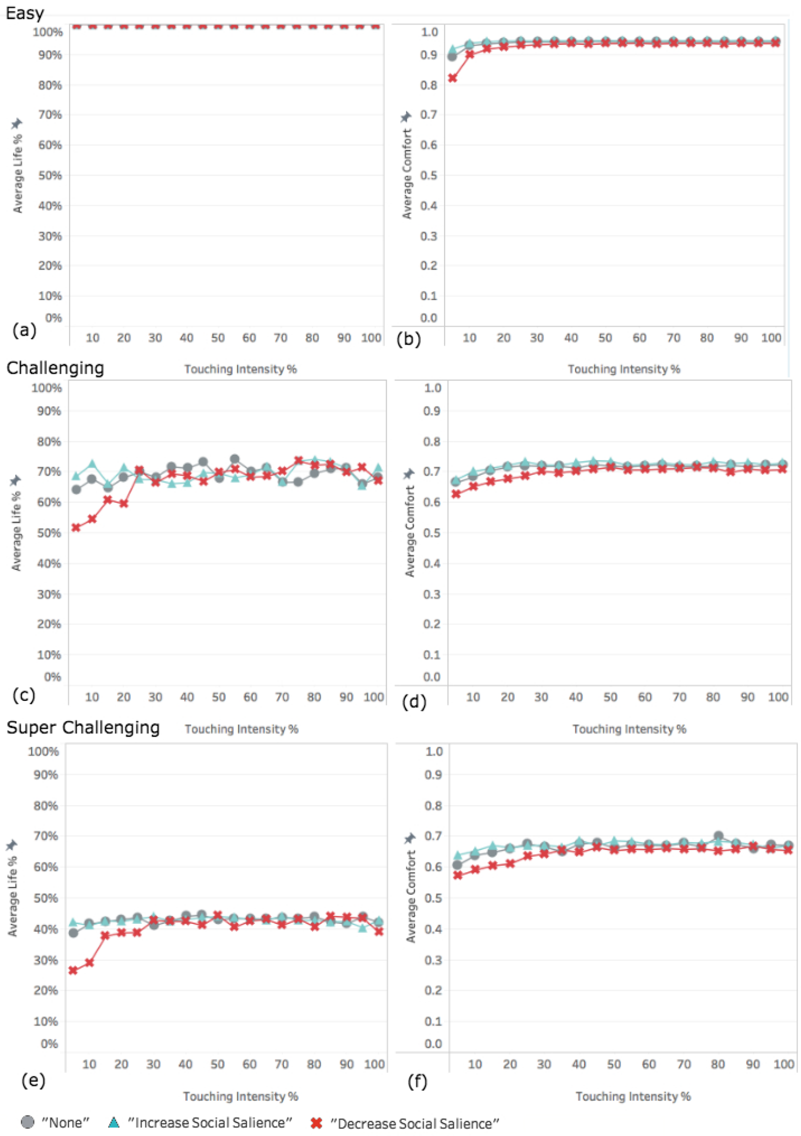

Figure D.1 Trended Life Length and Comfort results for Experiment 1, Chapter 4. Results are aggregated across all experimental runs.

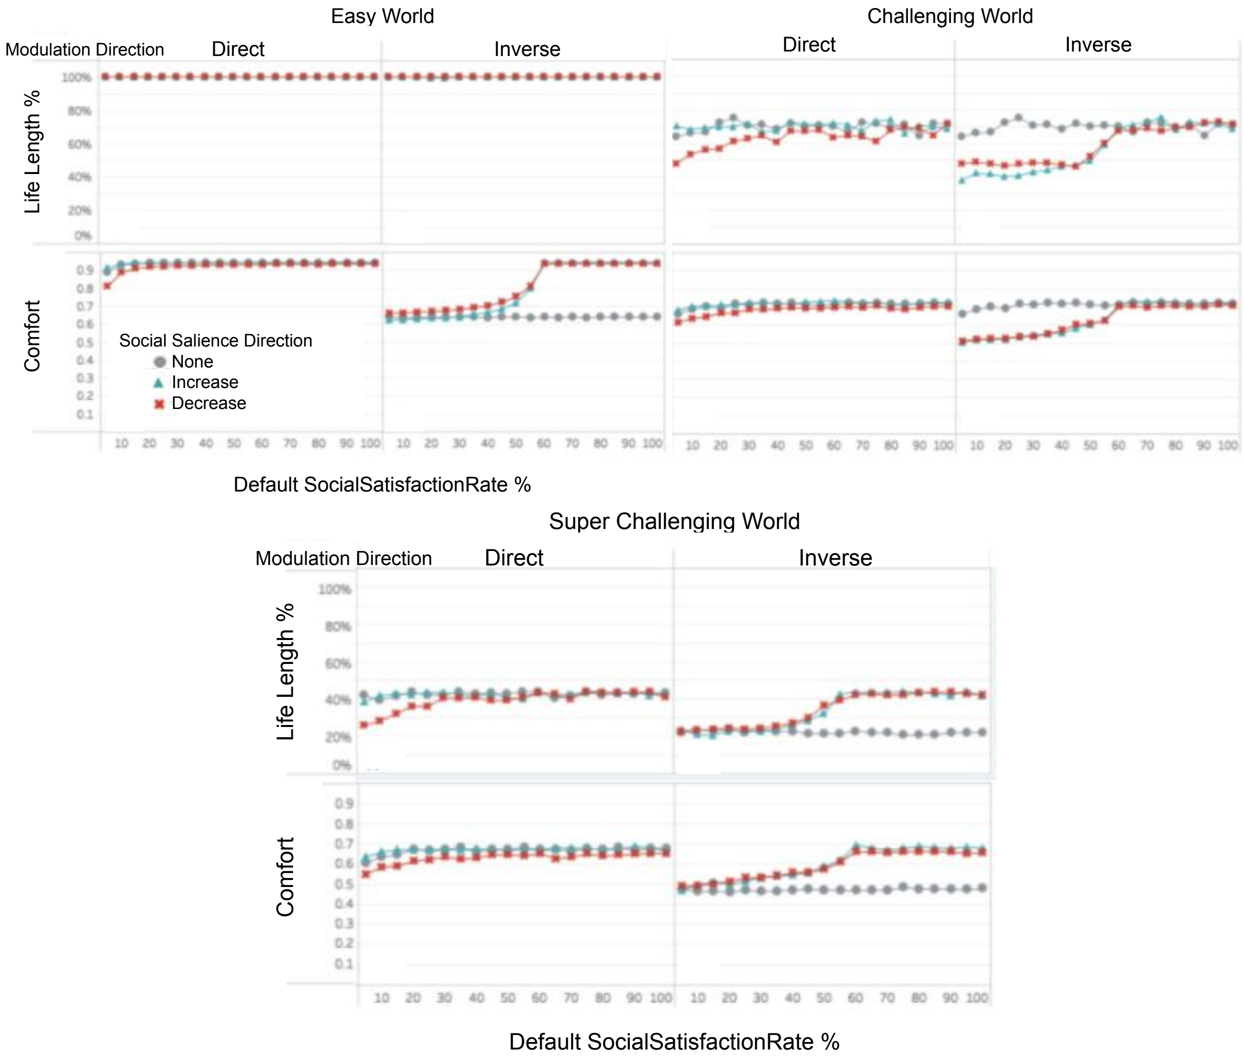

Figure D.2 Trended Life Length and Comfort results for Experiment 2, Chapter 4. Results are aggregated across all experimental runs.

Appendix E

Short Review of Chimpanzee Mechanisms that Inspired Mechanism Decisions in Chapter 5

Partner Selection is Value–Based

Mielke et al. [142] found that primates were flexible in their partner selection, and that a number of factors, including relative ranking and bond strength of potential partners, would influence the partner selection process. The relative "value" of a potential partner is therefore dynamic, based on a number of social conditions.

Social Bonds and OT Facilitate Food Sharing

Samuni et al. [180] discusses numerous hypotheses of food sharing within chimpanzee populations. They found that previously–formed social bonds were strongly correlated with food sharing (whether passive or active) in a population of wild chimpanzees. This is also noted by [293]. The authors also highlight the possibility of emotions playing an important role in animal reciprocal behaviour — including food sharing — and that OT systems may underpin the reinforcement of these emotional mechanisms. This hypothesis is also suggested by [109], [174].

Passive Food Sharing Through "Tolerated Theft"

Where [180] observed active transferring of food resources, [293] found that food transfers in captive chimpanzees was largely a result of tolerated transfers, with non–owners (A) approaching a food owner (O) and taking the food with (O) passively allowing it. It has been suggested that tolerant approaches to food sharing within primates may be due to them avoiding the cost of defending their food against beggars[294]. This "cost analysis" approach towards food sharing has been suggested to be a form of "tolerated theft" [295], which may play an important role in survival in highly competitive environments. As a first step in testing our new architecture, we implement a basic mechanism of this "tolerated theft" within our embodied agents.

Appendix F

Results of correlation analysis between agent rank and cortisol levels in Chapter 6.

	World Condition			
Stress Response Threshold	**A**	**B**	**C**	**D**
Low	-0.486*	-0.581*	-0.678**	-0.827**
Neutral	-0.299*	-0.476**	-0.628**	-0.681**
High	-0.252**	-0.557*	-0.587**	-0.735**
Low-High	-0.226*	-0.525*	-0.599*	-0.741**
High-Low	-0.337**	-0.472*	-0.701*	-0.700*

Table F.1 r values of relationship between rank and CT levels across all conditions. $* = (p \leq 0.05), ** = (p \leq 0.01)$

Appendix G

Results of correlation analysis between Stress Threshold condition and cortisol levels (top) and world condition and cortisol levels (bottom) in Chapter 6.

Stress Response Threshold	**r**	**sig**
Low	-0.647	**
Neutral	-0.792	**
High	-0.706	**
Low→High	-0.746	**
High→Low	-0.738	**
World Condition	**r**	**sig**
A	-0.390	**
B	-0.17	*
C	-0.251	*
D	0.538	

Table G.1 r values of relationship between world conditions and group cortisol levels, broken down by *StressResponseThreshold* conditions. $** = (p \leq 0.01)$.

Appendix H

Bond Condition	**World Condition**	r^2	**p-value**
A	Static	0.53	0.02
B	Static	0.49	0.04
C	Static	0.56	0.01
A	Seasonal	0.78	<0.01
B	Seasonal	0.70	<0.01
C	Seasonal	0.84	<0.01
A	Extreme	0.87	<0.01
B	Extreme	0.87	<0.01
C	Extreme	0.89	<0.01

Table H.1 Results of statistical analysis of correlation between mean oxytocin levels and mean Life Length for bonded agents in Chapter 9, across all conditions. Analysis was conducted using Pearson's Correlation Coefficient. Significance declared at 0.05.