Convention Emergence and Destabilisation in Multi-Agent Systems

by

James Michael Marchant

A thesis submitted to The University of Warwick

in partial fulfilment of the requirements

for admission to the degree of

Doctor of Philosophy

Department of Computer Science

The University of Warwick

September 2017

Abstract

Ensuring coordination amongst individual agents in multi-agent systems (MAS) helps to reduce clashes between them that waste resources and time and facilitates the capability of the agent population to solve mutually beneficial problems. Determining this coordinated behaviour is not always possible *a priori* due to technical issues such as lack of access to individual agents or computational issues due to the large number of possible clashing actions. Additionally, in systems lacking centralised authorities, dictating rules in a top-down perspective is difficult or impossible.

Conventions represent a light-weight, decentralised and emergent solution to this problem. Acting as a socially-accepted rule on expected behaviour they help to focus and constrain agent interactions to facilitate coordination. Understanding how these conventions emerge and how they might be encouraged allows scalable coordination of behaviour within MAS with little computational or logistical overhead.

In this thesis we consider how fixed strategy Intervention Agents (IAs) may be used to encourage and direct convention emergence in MAS. We explore their efficacy in doing so in various topologies, both static and time-varying *dynamic* networks, and propose a number of methods and techniques to increase this efficacy further. We consider how these IAs might be used to *destabilise* an existing convention, replacing it with a more desirable one and highlight the different methods required to do this. We also explore how various limitations such as time or observability of topological structure can impact the emergence of conventions and provide mechanisms to counteract these issues.

To my parents, Christopher and Liane, for always believing in me and supporting me.

To my younger brother, Joseph, for all the laughs and adventures throughout the years.

And to Aimée. My Evenstar. My most ardent fan. My love.

Acknowledgements

Over the course of this PhD many people have given me no end of support, advice, friendship and laughter. Without their myriad contributions I would not have completed this thesis nor look back at these last few years with nearly as much fondness. This list is by no means exhaustive and there are too many people to thank with not enough space.

I would like to acknowledge the friendship and support given to me by my housemates, James Archbold, James Dickson and Danielle Lloyd. The number of stress-free evenings and good humour they have provided was invaluable during the more trying times of this PhD and they always provided a listening ear and sound advice.

The friendship of Faiz Sayyid and Leanne Barwick has also been invaluable, who, along with James Archbold, Danielle and Aimée Crickmore, have provided hundreds of hours of stress-free gaming time, forming some of my fondest memories and acting as a wonderful distraction.

The advice and guidance provided by Peter Coetzee could not have been replaced. Always willing to offer a thought and helping me become a substantially better programmer than I was at the start, he has always been willing to listen. I would also like to thank Steven Wright and Phillip Taylor whose insights and technical help have always been appreciated from my time as an undergraduate through to this PhD.

The Department of Computer Science could not continue to function without the help of the numerous support staff and I would like to thank them for their efforts during my time here. In particular, I would like to thank Roger Packwood who has helped me fix more problems than I can count and ensured the smooth running of the technical aspects of the department for many years.

Some of the most important thanks must go to my supervisor, Nathan Griffiths, whose support throughout this PhD has been beyond invaluable. His calm, understanding demeanour and technical and personal guidance have been formative and I would not have completed this PhD without his help. I am truly grateful for his support even when deadlines have been missed and work has not gone as expected. His contribution to this thesis is certainly the most significant.

My deepest thanks go out to my family as well, who have been supportive and understanding throughout my academic career and without whom I would not have reached anywhere near where I am today.

I also wish to offer special thanks to Aimée who has selflessly supported me throughout this PhD but particularly over the last few months whilst writing up. She has been kind, caring and understanding and this PhD would not have been completed without her.

Sponsor

The research presented in this thesis was made possible by the generous support of The University of Warwick Chancellor's Scholarship.

Publications

Parts of this thesis have been published in the following:

- James Marchant, Nathan Griffiths, Matthew Leeke, & Henry Franks. "Destabilising Conventions Using Temporary Interventions". In: *Coordination, Organizations, Institutions, and Norms in Agent Systems X*. vol. 9372. LNCS. 2015, pp. 148–163
  - Introduces the notion of *destabilising* established conventions using fixed strategy agents in static networks and shows the efficacy of this. This work is discussed in Chapter 3. Additionally introduces the idea of temporary interventions, which is discussed in Chapter 5
- James Marchant, Nathan Griffiths, & Matthew Leeke. "Destabilising Conventions: Characterising the Cost". In: *Proc. of the 8th IEEE International Conference on Self-Adaptive and Self-Organizing Systems*. 2014
  - Builds on the work of the above and explores the notion of *minimum interventions* and the cost of these, as discussed in Chapter 5.
- James Marchant, Nathan Griffiths, & Matthew Leeke. "Convention Emergence and Influence in Dynamic Topologies". In: *Proc. of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. Istanbul, Turkey, 2015, pp. 1785–1786
  - Explores convention emergence and destabilisation in dynamic topologies and introduces Life-Degree as a metric for exploring aspects unique to dynamic topologies. Work from this is included in Chapter 3.
- James Marchant & Nathan Griffiths. "Manipulating Conventions in a Particle-Based Topology". In: *Coordination, Organizations, Institutions,*

and Norms in Agent Systems XI. ed. by Virginia Dignum, Pablo Noriega, Murat Sensoy, & Jaime Simão Sichman. Vol. 9628. LNCS XI. 2016, pp. 242–261

  - An extension of the above, this work includes more details on the nature of convention emergence in dynamic topologies and explores the effect of differing payoff matrices on said nature. This work is included in Chapter 3.

- James Marchant & Nathan Griffiths. "Limited observations and local information in convention emergence". In: *Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems.* International Foundation for Autonomous Agents and Multiagent Systems. 2017, pp. 1628–1630

  - Introduces the idea of partial observability as a limiting factor for convention emergence and shows the efficacy of our approach to dealing with this. This work forms the basis of Chapter 4.

- James Marchant & Nathan Griffiths. "Convention emergence in partially observable topologies". In: *Proceedings of the International Workshop on Coordination, Organizations, Institutions and Norms in Agent Systems (COIN@AAMAS 2017).* 2017

  - An extended version of the above, this work includes a more in depth analysis of the algorithms presented and the various aspects of partial observability that may impact convention emergence. This work is also used in Chapter 4.

Abbreviations

BC	betweenness centrality
CC	closeness centrality
Degree	degree centrality
EC	eigenvector centrality
EM	external majority
FS	fixed strategy
GSM	generalised simple majority
HCR	highest current reward
HEE	highest edge embeddedness
HITS	hyperlink-induced topic search
IA	Intervention Agent
MAS	multi-agent system
RFR	recruitment based on force with reinforcement
SF	scale-free
SM	simple majority
SW	small-world
u.a.r.	uniformly at random
WCC	weakly-connected component
WoLF-PHC	"Win or Learn Fast" policy hill-climbing

Contents

List of Figures

List of Tables

List of Algorithms

CHAPTER 1
Introduction

Coordination is a fundamental problem in multi-agent systems, where multiple, often independently controlled and owned, agents must interact with one another to facilitate their goals and improve their own, and the group's, welfare. Coordination amongst these agents acts to reduce clashes and wasted resources from agents choosing incompatible actions with one another. Finding coordinated actions is often difficult however, with large action spaces available to individuals and not necessarily any *a priori* knowledge about which actions will clash and which will not. Finding ways to constrain agent action choices so as to minimise clashes helps to maximise the potential of the system.

Designing off-line solutions to these problems is often computationally infeasible [Tinnemeier et al., 2010] or inapplicable in many domains where there is no centralised control of agent populations. Instead, recent research efforts have focused on the use of *emergent*, on-line behaviour where solutions are developed by the interactions between agents themselves. *Conventions* represent one of these approaches, acting as self-imposed, socially-adopted constraints on expected agent behaviour which allows coordination by restricting the likely action choices of agents. Conventions are light-weight, requiring few assumptions about agent architecture and have been shown to emerge unaided assuming only agent rationality and the ability to learn from interactions [Delgado, 2002; Walker & Wooldridge, 1995].

Directing and manipulating convention emergence in order to facilitate rapid and robust convergence and minimise the amount of time that agents are in conflict is a recent development. It has been shown that this is possible through the use of small numbers of fixed strategy agents who are able to direct the emergent conventions of populations much larger than themselves [Griffiths & Anand, 2012; Sen & Airiau, 2007].

In this thesis we concern ourselves with this notion of manipulating and directing convention emergence through the use of fixed strategy Intervention Agents (IAs) to facilitate fast and stable convergence. We explore the effect of these IAs in manipulating conventions, both emerging and already established, and their efficacy at causing change within the system to reach a desired outcome.

1.1 Multi-Agent Systems

The use of multi-agent systems (MAS) to model abstract populations of individuals has seen a dramatic increase in use over recent decades as computing power has reached levels applicable to the problems in this domain. As the use of MAS becomes more pervasive, the need for rapid, scalable, decentralised methods of enforcing beneficial and cooperative behaviour has become more important [Durfee, 2004; Durfee, 1999; Jennings, 1993]. As we begin to enter the age of the Internet of Things (IoT) and MAS begin to enter more safety critical arenas such as self-driving cars, ensuring coordinated behaviour amongst disparate and often separately controlled agents becomes paramount.

These domains present a range of different issues to contend with: decentralised communication, dynamic populations and networking, heterogeneous agent architectures and conflicting intentions and goals are all common aspects in scenarios including agents from multiple parties. The inability to dictate, without legal fiat, internal aspects of the agent architectures means that any mechanism that increases coordination or cooperation must work at a high level in order to be generalisable. Finding such a general solution remains an ongoing area of research but finding a solution that is powerful, scalable and domain-agnostic is difficult.

The constraints identified above mean that dictating behaviour is unlikely to be a viable solution to the problem of coordination in open MAS. Ensuring that rules are properly formalised, communicated and enforced requires a number of system-level design decisions that limit its applicability to tightly regulated or controlled domains. Instead, allowing agents to facilitate the emergence of mutually beneficial behaviour amongst themselves as a means to solve problems offers a solution that fits within all of the constraints.

1.2 Conventions

Conventions fulfil the criteria outlined above, representing convergence of agent choices to mutually beneficial actions. They do not require explicit punishment mechanisms as the punishment for going against the established convention is the penalty of clashing actions [Kandori, 1992; Savarimuthu et al., 2011]. They are able to form in an entirely decentralised manner and as such are applicable to the types of MAS likely to be encountered above. Conventions represent "an equilibrium everyone expects in interactions that have more than one equilibrium" [Young, 1993] and have been described as a regularity in behaviour amongst a collection of agents.

Previous work has shown that global convention emergence amongst a population is possible, even with minimal assumptions [Walker & Wooldridge, 1995] and indeed many do not consider a convention to have emerged until nearly universal adoption has taken place [Kittock, 1995]. Conventions are also self-reinforcing, with a positive feedback loop of increased utility the more the convention is being used [Boyer & Orléan, 1992]. With such high levels of usage and their self-reinforcing nature, understanding how to *remove* conventions and overcome these issues is necessary in situations where there are time-varying notions of optimality.

They have been used to describe and quantify a number of coordinated behaviours in fields as diverse as economics [Akerlof, 1980; Jones, 1984], marketing [Delre et al., 2010] and politics [Snidal, 1985] and are widely accepted as good models for interactions that require a common behaviour to emerge. Understanding how these conventions emerge, how to direct them and how to remove undesirable conventions is an open research problem and identifying and investigating these is the focus of this thesis.

1.3 Constraining Environments

MAS are often situated within a topology that constrains the interactions between agents to those who are neighbours within the topology. This models a number of real-world situations including computing, communication and social networks. They have been shown to have a dramatic effect on convention emergence [Delgado, 2002; Delgado et al., 2003; Kittock, 1995] and as such modelling their inclusion is an intrinsic part of studying convention emergence that is applicable to wider domains. These underlying topologies often exhibit complex structural properties within themselves that modify the nature of a society of agents placed within it [Barabási & Albert, 1999; Kleinberg, 2000a].

In particular, the notion of time-varying, dynamic networks better models the nature of many open MAS where the population is able to change over time as agents join and leave the network. Additionally these dynamic networks are able to model scenarios where agents are moving in 3D space and the interactions between them change because of this. These are known to have different dynamics than similar static networks [Brandt & Sigmund, 2005] and as such are likely to induce different dynamic within convention emergence.

Because of these features, the underlying network topology and how information on agent influence can be tied to it forms an integral part of the study of convention emergence in this thesis and some agents will be notably more influential and desirable because of the local network structure. As such, in this thesis we always assume that the agent society is situated in a connective topology and the investigation of the effects this has, particularly the effects of dynamic topologies, forms part of the research contained within.

1.4 Objectives of the thesis

This thesis aims to expand the knowledge on the emergence of conventions in MAS and to understand how they might be manipulated in order to facilitate higher levels of coordination amongst agents with minimal effort.

More precisely, in this thesis we aim to do the following:

1. Explore how convention emergence in MAS can be influenced and manipulated to ensure a rapid and robust convergence to coordinated behaviour in scenarios with no centralised control and complex interconnecting topologies.

2. Investigate how already established conventions can be destabilised and replaced with more desirable or optimal ones and how the force of precedence that grants them stability may be counteracted.

3. Garner an understanding of the differences in convention emergence in time-varying dynamic topologies and how their emergence may be directed or destabilised using unique facets of the dynamic topological structure.

4. Develop techniques for fostering convention emergence and destabilisation in scenarios where topological information is restricted, as is often the case in unknown, real-world networks.

5. Identify the minimum level of intervention that is necessary to foster permanent change in a population and how this limit may be exploited to increase the performance of mechanisms that support coordination in MAS.

1.5 Contributions of the thesis

In developing techniques and understanding of the nature of conventions in MAS we identify the following main contributions that this thesis makes:

1. **Introduction and analysis of the concept of destabilisation of an existing convention and techniques to facilitate this.**

Whilst previous work has been done that explores the contribution of fixed strategy agents to convention emergence when inserted at the start of a simulation [Griffiths & Anand, 2012; Sen & Airiau, 2007], we consider the case where a convention has already become established and we wish to remove it and replace it with another.

We introduce the concept of *destabilisation* of an existing convention and present techniques to manipulate and remove these convention by inserting a number of fixed strategy Intervention Agents (IAs) at topologically influential locations. We show that (i) a small proportion of the convention population used as these IAs is enough to guarantee destabilisation and replacement and (ii) that topological effects contribute to the effectiveness of the destabilisation efforts and influence the number of IAs needed.

2. **An exploration of convention emergence in dynamic topologies and the creation of placement metrics to encourage convention emergence and destabilisation in these topologies.**

   Dynamic networks, those that vary over time, are known to exhibit substantially different dynamics than static networks [Brandt & Sigmund, 2005; Savarimuthu et al., 2007]. Despite this, no research has focused on how these differences manifest in convention emergence. We investigate convention emergence and destabilisation in a number of dynamic topologies and clarify the differences that exist. We introduce new placement mechanisms for IAs that are specific to the time-varying nature of dynamic networks and show their efficacy as well as an analysis of topological features unique to such networks and their impact on influential agents.

3. **Development of placement algorithms that find influential locations in topologies with restricted observability.**

   In the age of billion member social networks, it is no longer the case that an entire topological network can be analysed and searched completely for influential locations [Avrachenkov et al., 2014; Brautbar & Kearns, 2010]. Given the importance of these locations in increasing the efficacy of IAs we contribute algorithms that can find influential locations in both static and dynamic topologies given finite, limited observation of the local area of the graph. We show that these are able to closely approximate the performance of IAs with full graph knowledge when manipulating conventions.

4. **Mechanisms for assessing minimal interventions, their costs and an investigation into their effectiveness at manipulating conventions.**

   The self-reinforcing nature of conventions [Boyer & Orléan, 1992; Lewis, 1969] means that the permanent inclusion of IAs to manipulate conventions is unnecessary but is the only approach used so far [Airiau et al., 2014; Franks et al., 2013]. We explore the nature of temporary interventions and identify the minimum timeframes that they must be located in a system to have the same effectiveness as permanent inclusion. We introduce a number of mechanisms for assessing these minimum interventions and explore the relationship between minimum interventions and number of IAs. Given that the creation and maintenance of IAs is likely to have an associated cost in real-world scenarios, we analyse the minimum interventions within this framework to identify ways to minimise cost. We then use these findings to inform design of techniques for budgeted placement of IAs.

1.6 Structure of the thesis

The remainder of this thesis is structured as follows. Chapter 2 presents an overview of the literature regarding coordination in MAS, the history and necessity of conventions and how fixed strategy agents may be used to direct and encourage convention emergence. We also discuss related areas of work such as normative behaviour and the problem of cooperation between agents as areas that provide solutions to similar problems. A brief overview of the game theoretic underpinnings that are used to model convention emergence are presented and we analyse what features in these games make the problems amenable to solution by convention. A summary of the network topologies in use throughout this thesis and the network metrics that will be utilised for the identification of influential locations concludes this chapter.

In Chapter 3 we investigate the concept of *interventions* in MAS, using IAs who adhere to a singular fixed strategy in favour of all others to influence and direct the rest of the population in convention emergence. We explore the nature of convention emergence in both static and dynamic topologies and examine how the different features of these affects the manner in which convention emergence can be manipulated. We introduce the concept of *destabilisation* and show how IAs can be used to remove already established conventions in multiple paradigms.

Chapter 4 investigates the nature of convention emergence under partial observability where access to information about the underlying topology is restricted. We provide a number of tools and insights to allow the effective placement of IAs under constraints to maximise their effectiveness at directing convention emergence. We investigate partial observability in both static and dynamic networks and highlight the differences between them, introducing additional tools for the latter.

In Chapter 5 we introduce the concept of a *temporary intervention*, including IAs within the population for a finite time rather than permanently. We investigate how these temporary interventions may be used to elicit the same level of change we have seen previously and consider what how short temporary interventions can be to still cause this. We explore the notion of the cost of an intervention and show how this can be used to inform budgeted placement of IAs.

Finally, in Chapter 6 we present our conclusions and highlight the research contributions presented as well as identifying avenues for additional future work built on the findings of this thesis.

CHAPTER 2

Background and Related Work

In this chapter we present a summary of the related work in the field of convention emergence, both as an underlying concept of societal interaction and as a method for facilitating coordination in multi-agent system (MAS). Section 2.1 examines the history of conventions in the wider literature as well as its use in recent years as a tool to foster cooperation and coordination. Section 2.3 then highlights the contributions of the use of fixed strategy (FS) agents in effecting robust convention emergence in agent-based systems whilst Section 2.4 describes the game theoretic approaches that are often used to describe agent interaction. Section 2.5 discusses the work so far on the notion of *stability* in emerged conventions. Finally, Section 2.6 describes the network topologies that will be used throughout the rest of this thesis and Section 2.7 explores some of the common network metrics that are used in deciding influential locations within these topologies.

2.1 Conventions, Coordination and Normative Behaviour

We begin with an overview of the nature of conventions: what they are, why they are beneficial and how they can be used to solve the coordination problem in MAS. We will distinguish them from norms and cooperation before examining some of the ways in which they may be manipulated in order to ensure robust and rapid convention emergence even in systems that seem resilient to it.

Conventions amongst members of a society represent unwritten rules that govern interactions between individuals. They take the form of constraints on the expected behaviour that will be encountered in specific situations. In human society they can be as simple as an expectation of what hand to use to shake another's or as complex as the many unwritten rules that govern the actions of the UK Government (so called "Constitutional Conventions"). When adhered to, conventions facilitate easy, unspoken interactions to proceed smoothly without the need of formalised rules. When conventions are broken though, with individuals going against the expected behaviour, the interactions become more complex as the situation must now be navigated without knowledge of likely behaviour.

Conventions can thus be thought of as *socially-accepted rules in the form of expected behaviour*. They represent self-imposed restrictions over a wide, perhaps infinite, range of possible actions down to a smaller (preferably unitary) number such that any individual can, with some degree of certainty, know how an interaction is likely to proceed before it even happens.

These features lend themselves well to addressing the issue of uncoordinated action choices amongst agents in MAS [Durfee, 1999; Huhns & Stephens, 1999; Jennings, 1993; Vylder, 2007]. Agents that lack coordination often conflict with one another which causes wasted resources [Durfee, 2004] either in the form of time, energy or some other cost to the agent due to incompatible action choices being made. Shoham & Tennenholtz [1992b] argue that in MAS agents *have to* agree on common rules to promote cooperative behaviour and decrease clashes implying it is a fundamental requirement. Additionally, if the desired goal that the system is trying to reach requires the combined efforts of multiple agents, this lack of coordination may make it impossible for the goal to be achieved. As Durfee notes, it is unlikely that there is a universal, always applicable coordination strategy, otherwise human society could adopt it for all problems that require a coordinated front. As such, conventions and their manipulation and convergence are a useful abstraction that can be used to facilitate this aim.

A number of views on what constitutes a convention can be found in the literature. The seminal work in the formalisation of convention emergence in the scientific sphere comes from Lewis [1969]. Lewis discusses the notion of conventions as a way to describe the emergence of language amongst early humans. Whilst there was no formal establishment of rules (which in modern parlance would be called a *prescriptivist* view of language) humans obviously established consensuses regarding meaning and method of communication in the form of conventions on how linguistic interactions should be governed. He shows that a number of day-to-day interactions that we take for granted are in many ways governed by conventional behaviour. In the words of Young [1993], Lewis formally defined a convention as a *regularity* in behaviour to which "everyone conforms, everyone expects others to conform, and everyone wants to conform given that everyone else conforms". Schotter [1981] prominently agrees with Lewis calling conventions a "regularity in behaviour which is agreed to by all members of a society".

In his formalisation, Lewis is one of the first to take a game theoretic approach to describe the reasons why conventions emerge, the benefits of them and how that might be explored mathematically. He explicitly highlights the notion of *coordination games* in his theory of convention, contrasting it to the previous explorations of pure conflict games where the aims and outcomes for individuals are diametrically opposed.

Goyal & Janssen [1997] go even further, noting that conventions are, in actuality, arbitrary in their solution to the underlying social problem. That is, several of the choices available could solve the problem equally well but individuals conform to one because they expect others to do so as well, elevating it to a higher level of preference than the other options partly because it already is more preferable. This positive feedback mechanism of individuals conforming to a convention because others conform to it, underpins the benefits that conventions introduce to a system. As Jennings [1993] argues: "All coordination mechanisms can ultimately be reduced to (joint) commitments and their associated (social) conventions."

Boyer & Orléan [1992] note that one of the essential characteristics of conventions over other forms of coordination and cooperation is their self-reinforcing quality. The more individuals subscribing to the convention, the greater the utility of choosing that convention as you are more likely to meet other users of it than any other. Indeed, Schelling [1980] argues that there are no distinguishing characteristics between the multiple Nash equilibria of the coordination games often used to explore coordinated behaviour but that agents still exhibit an "intrinsic magnetism" towards certain solutions based on common experience, prominence or precedent.

Within the MAS community, Shoham & Tennenholtz [1997] are some of the first to utilise the notion of conventions to solve coordination between agents. They utilise a game-theoretic approach to describe the positive effect that agents receive when they successfully coordinate or cooperate. They define a *social convention* as a restriction on the set of actions available to the agents down to a singular action that maximises some variable within the game. Importantly, they note that this variable may not necessarily be optimal for an individual agent but maximises some other utility such as the group reward. In this way they show that conventions can be used to foster not just coordination but *cooperation*; enforcing behaviour that is not necessarily rational for an individual but aims to be "best" based on some other metric.

Kittock [1995] similarly applies the notion of games as a way to allow conventions to emerge through iterated play of the coordination game and Prisoner's Dilemma (see Section 2.4 for more details) where agents learn through monitoring the highest payoff received for each action choice and changing choice once a better paying action has been observed. Kittock defines a convention as being established when some threshold, usually 90%, of all agents in the population are choosing the same action. We refer to this as the *Kittock Criteria*. The distinguishing feature of the approach of Kittock compared to the others is that Kittock does not seek to address restrictions on the actions on agents (either explicit or implicit) but instead merely notes the emergent behaviour is one that is being used as it represents the greatest benefit to all agents.

In each of these definitions the fundamental property is that the value of the convention is partially due to its self-reinforcing nature. The expectation that others will use it makes others value it beyond what the payoff they receive from it might be. Indeed, there is not necessarily any explicit benefit in the action represented by the convention over other possible actions and many games often have multiple equally valid equilibria [Myerson, 1991] Thus a convention can be thought of as "an equilibrium everyone expects in interactions that have more than one equilibrium" [Young, 1996]. The convention is distinguished by what Lewis [1969] and Young [1993] identify as two major characteristics: *salience* and *precedence*. *Salience* is a feature of an action choice that marks it as different in the eyes of the agents. It is not necessarily explicit, representing a fundamental difference in the action compared to others, but simply something that marks the action as unique amongst equal choices. *Precedence* can be viewed as a special case of a salient property and makes a certain action choice more desirable simply because it has been observed to be chosen before. Indeed, Young [1996] describes the "gradual accretion of precedent" as the main mechanism through which otherwise identical equilibria can become conventions.

These special equilibria are also known as "focal points" [Schelling, 1980] and are "equilibrium more likely to be chosen by the players because it seems special, natural or relevant to them, although other equilibria are equally good". Young [1996] notes that conventions are effectively focal points evolved through learning. Vylder [2007] formalises the notion of this different yet equal nature of possible conventions by way of describing the *convention space* of a given scenario. He delineates between *flat convention spaces* where there is no preference at all between possible conventions and *structured convention spaces* where there is some salient preference (for instance a convention space which consists of possible days of the week to go to the shops is likely to exhibit an intrinsic preference for the days of the weekend). In the work presented in this thesis we concern ourselves only with flat convention spaces and choose games and settings to ensure this is the case.

Whilst the *emergence* of conventions, the manner in which they arise naturally amongst agents in a population, is the primary target of exploration, they have also been considered for *a priori*, offline generation. Shoham & Tennenholtz [1995] and Shoham & Tennenholtz [1992b] were the first to address the notion of formalizing conventions and creating them during the design of the system. They design *social laws* for use by systems of robots to dictate desirable behaviour from the very beginning of the system and include the capability of the agents to be aware of these social laws. They show that whilst designing explicit laws is simple the general problem of predetermining rules that will allow no need for online conflict resolution is intractable. Additionally, offline generation limits the system as it does not allow for change over time and would require the expensive reprogramming of the conventions within the system [Tinnemeier et al., 2010].

2.1.1 Cooperation

Having described the prominence of conventions in solving the problem of coordination we must briefly address a related but distinct problem, that of encouraging *cooperation* between agents rather than coordination. Cooperation in many scenarios can be viewed as a subtype of coordination [Axelrod & Hamilton, 1981]. However, whilst in the coordination problem the underlying assumption is that equilibria are choices that are maximally beneficial to each individual agent the cooperation problem is defined by the fact that agents are incentivised to act selfishly in order to try to maximise their own personal rewards. They must hence be encouraged to act in a way that is mutually beneficial even if this comes at a personal cost. This notion of cooperation introduces a number of different dynamics that convention emergence is not intrinsically suited for as "agent benevolence" can no longer be assumed [Genesereth et al., 1986; Ullmann-Margalit, 1977] and is most frequently described using the Prisoner's Dilemma game which is explored in more detail in Section 2.4. Fundamentally, the mechanisms available to address these issues are applicable to the coordination problem, although in many cases they may represent solutions that are unnecessarily strict given the presence of agent benevolence.

The work of Nowak [2006] represents one of the foundational approaches to encourage the evolution of cooperation. Nowak identifies 5 rules from amongst the wider literature and drawing inspiration from nature that can be credited with encouraging cooperation: kin selection, direct reciprocity, indirect reciprocity, network reciprocity and group selection. Several of these have direct applicability in encouraging coordination as well: direct reciprocity and indirect reciprocity focus on agents learning from their interactions with others or observing second-hand interactions. They allow agents to learn the expected behaviours of others and hence can inform agent decisions in their own interactions to maximise their payoff. Network reciprocity and group selection offer justification for alternative mechanisms that make use of the nature of those agents connected to each other in a topology or allowed to self-organise into distinct groups. Nowak shows that each of these mechanisms can facilitate cooperation amongst agents. Franks [2013] uses these as the basis for informing designs that encourage convention emergence and hence show the general applicability of cooperation mechanisms in the coordination game.

Purvis et al. [2006] introduce the concept of *monitor agents* to facilitate robust cooperation. They allow agents to self-organise into different subgroups and to include or exclude players based on their history of cooperation or defection. They show that these mechanisms cause agents to foster more cooperation than if allowed to interact freely but that the system is still open to new membership from unproven agents. Kułakowski & Gawroński [2009] look at the related mechanism of allowing agents to assign *reputation* to each other to determine how they will interact with one another. They show that this results in substantially higher levels of cooperation than would otherwise be expected as agents are able to more accurately predict what behaviour to expect from those they interact with.

2.1.2 Norms

Norms represent a stricter set of restrictions placed upon agents than those due to the expected behaviour that is encountered within conventions. Despite this, the terms are often used interchangeably in the literature [Mukherjee et al., 2007; Sen & Airiau, 2007]. In this thesis we thus differentiate between conventions and norms.

Norms typically imply an *obligation* (things an agent *must* do [Axelrod, 1986; Tuomela, 1992]) or *prohibitions* (things an agent must *not* do) on agents with regards to a specific action often specified as explicit logical rules on what behaviours are allowed or forbidden. Failure to adhere to norms and exhibit the expected behaviour is often associated with punishments or sanctions [Axelrod, 1986; Bicchieri et al., 1997; Kandori, 1992; Savarimuthu et al., 2011]. Alternatively, agents may be explicitly rewarded for adherence to norms. Thus, norms generally require additional system or agent capabilities as well as incurring a system-level overhead for punishment/reward. Conventions can be considered a light-weight alternative to norms, requiring little additional overhead and emerging solely from agent-agent interactions without the requirement of formalisation. Normative systems can also emerge amongst a population but require additional capabilities of agents to internally represent the norms [Mahmoud, 2013].

Tuomela [Tuomela, 1992; Tuomela, 1995] distinguishes between *r-norms* and *s-norms* in the context of normative behaviour amongst agents. They define *r-norms* (or *rules*) as norms created by a centralised authority which represents the agent population and can be based on either explicit or implicit agreement-making. By contrast, *s-norms* (or proper social norms) are based on mutual belief about beneficial actions and correspond to what we and other literature call conventions. Tuomela further divides *r-norms* into formal and informal types: formal rules are specifically articulated norms with specified sanctions, whereas informal rules, whilst still articulated, use informal, social sanctions. Similarly, Tuomela subdivides *s-norms* into *conventions* which are relevant for the whole society and group-specific *s-norms* which only concern a group of agents within the society.

Boella & Torre [2008] offer an alternative classification of norms between those that are *constitutive* (facts about the system that can support regulative norms), *regulative* (the rules that govern agent behaviour) and *procedural* (norms which seek to guarantee the others, often achieved by providing agents with the mechanisms necessary to detect and enforce violations). Mahmoud et al. [2017] and others also consider the notion of *meta-norms*, "rules about rules", which deal with situations where norm enforcement is not guaranteed and allow hierarchies of rules and agents who can punish those who don't punish others.

Whilst there are differences between conventions and norms, many believe that conventions can be subsumed into the hierarchy of norms due to the similarities between them. In particular, when considering social norms, results from one give insight into the other.

2.2 Conventions in Multi-Agent Systems

Having examined the literature for both coordination and cooperation and highlighted the differences between conventions and norms we now focus on the exploration of convention emergence in MAS.

Many different approaches have been used to study convention emergence in MAS, both to understand the underlying nature of conventions as well as to provide mechanisms to encourage their emergence in agent populations. Many different learning algorithms have been presented and analysed for their efficacy in allowing norms or conventions to emerge [Airiau et al., 2007; Hasan, 2014; Nudelman et al., 2004; Vu et al., 2006]. Ensuring that agents have the capability to learn from their interactions and designing an approach that allows them to do so well is paramount to ensuring that conventions can emerge. If agents learn ineffectively, they may never converge to a convention. If they are too quick to learn when presented with new information they are likely to switch strategies frequently and cause clashes. As Walker & Wooldridge [1995] note: "The key problem is to design a strategy update function, representing an agent's decision making process, that when used by every agent in the society, will bring the society to a global agreement as efficiently as possible".

In this thesis we make minimal assumptions about agent capabilities, assuming only that agents are rational (and hence wish to maximise their payoffs in any given interactions) and have access to, at most, a (limited) *memory* of past interactions. This setting has been widely studied [Delgado et al., 2003; Griffiths & Anand, 2012; Sen & Airiau, 2007; Walker & Wooldridge, 1995] and is able to support effective convention emergence without requiring architectural changes to agents in order to facilitate it.

Walker & Wooldridge [1995] were amongst the first to produce a formal model of convention emergence with these minimal assumptions. They present a model in which a global convention emerges where agents choose their action based solely on observations of others. A population of agents is located on a grid searching for food, with the aim to maximise their intake with conflict occurring when multiple agents try to eat the same food. Each agent updates their strategy (to give precedence to others based on their relative location on the grid) using the simple majority (SM) mechanism where they will change to an alternative strategy when they have observed more instances of that strategy than their current one. Whilst simplistic in nature, and not generally applicable in many regards, the work of Walker & Wooldridge shows that global convention emergence is indeed possible with the given assumptions and without use of memory.

Shoham & Tennenholtz [1992a] examine a simple coordination game where agents are rewarded if they match on an internal value that is drawn from the set $\{0, 1\}$. They propose a new strategy update rule, external majority (EM), which makes use of an internal memory to the agents to monitor the conventions encountered and to update if the other was encountered more than your own. They show that the size of the memory available affects the speed of convention emergence with limited memory reducing efficiency. They also explore the notion of larger convention spaces and show that, whilst increasing the number of conventions reduces the likelihood of emergence it does so in a less than logarithmic fashion. However the work does not consider the effect of topologies with all agents able to interact with all others.

Shoham & Tennenholtz [1997] also introduce the concept of highest current reward (HCR) as an update rule where agents monitor the payoffs they have received for using each action and switch to that action providing the highest current payoff. They show that for certain classes of games using HCR is guaranteed to converge to a convention and that a limited memory of 2 to 3 times the number of agents provides optimal performance in increasing the speed of convention. Kittock [1995] expands on their work by exploring the effect of network topology on convention emergence, showing that it introduces different dynamics. In particular he shows that increasing graph diameter has a detrimental effect on convention emergence in the coordination game. However, Kittock only considers a ring topology with limited interaction radius; a poor approximation of many real-world topologies.

Delgado [2002] expands on the methods of Walker and Wooldridge and introduces a new variant of their update strategy, generalised simple majority (GSM), that allows for its use with agents located in a topology. Delgado uses this and HCR with both making use of a memory of finite size to monitor how the agent's choices have been rewarded previously. He shows that utilising such information can increase the rate at which conventions emerge. However, the use of memory places additional requirements on agent capabilities. They also explore the effect of network topology on the speed of convention emergence finding that complex topologies such as small-world and scale-free allow much more rapid convention emergence than equivalent regular graphs (where all agents are connected to all others).

Urbano et al. [2009] expand upon this further, exploring the effectiveness of both EM and HCR in numerous graph topologies: regular graphs, random graphs with uniform degree distribution, scale-free and small-world. They show that the different topologies exhibit drastically different behaviours with scale-free and small-world networks producing more rapid convention emergence than regular graphs in keeping with [Delgado, 2002]. They also evaluate the performance of recruitment based on force with reinforcement (RFR) [Urbano & Coelho, 2005] which makes use of “force” of a convention, a measure that increases in strength every time the convention successfully causes an agent to switch to it. They show that RFR outperforms all other update strategies in encouraging rapid convention emergence.

Many of the previous mechanisms for convention emergence have relied upon agents having access to a memory of interactions or having the ability to observe the state or interactions of other agents to determine the dominant convention. In open MAS where the agents are often controlled by multiple independent parties, these assumptions about capabilities may not be reasonable and changing agent architecture would be difficult or impossible. Sen & Airiau [2007] investigated the use of “social learning” for convention emergence, where agents receive a payoff corresponding to a game from their interactions which informs their learning (via Q-Learning [Watkins, 1989] or “Win or Learn Fast” policy hill-climbing (WoLF-PHC) [Bowling & Veloso, 2001]). They showed that convention emergence can occur even when agents have no memory of interactions and only observe their own rewards. The payoffs directly quantify the notion of an action clash costing resources. This approach does not rely on knowledge of other agents and indeed does not require identification of whom the agent is interacting with (the *obliviousness* property). It also allows for interactions to be *private* and is particularly applicable as an approach in open MAS because of this.

However, their model is limited in that agents are able to interact with any other member of the population rather than being situated in a network topology. Additionally, the convention space considered is restricted to only two possible actions. In more realistic settings larger convention spaces and more restrictive connecting network topologies are likely. Airiau et al. [2014] build upon this work, showing that a population of Q-Learning agents converges to a convention faster than a population of WoLF-PHC agents.

Other memoryless approaches to convention emergence exist. Recently, Mihaylov et al. [2014] introduced another strategy update rule based on developments from game theory: Win-Stay Lose-probabilistic-Shift (WSLpS) where agents continue to use their chosen strategy if they are being rewarded but will change strategy with some probability, p, when they receive a negative reward. They show that this approach outperforms other state-of-the-art coordination mechanisms and guarantees full convergence in any topology. This highlights the fact that convention emergence is possible simply from repeated play of the coordination game without the need for monitoring other agents' behaviours.

Another memoryless mechanism that has been shown to allow convention emergence is that of imitation [Borenstein & Ruppin, 2003]. Agents mimic the choices of others, either universally or only of those agents that they have reason to believe are informed better than themselves. This "tit-for-tat" approach has been shown to be highly effective at fostering cooperation [Axelrod & Hamilton, 1981] and has similar applicability in the coordination domain. Savarimuthu et al. [2008] introduce a similar mechanism for the purposes of norm emergence. They allow agents to select those they view as influential (so-called "role model" agents) and to mimic their behaviours after interacting with them. They show that the resultant set of "role model" nodes follows a power-law distribution (even in random graphs that have no such distribution among the general populace) and that complete norm emergence is possible under this mechanism.

As discussed above, the underlying topology has been shown to have a significant effect on convention emergence [Delgado, 2002; Delgado et al., 2003; Kittock, 1995; Pujol et al., 2005; Santos et al., 2006; Villatoro et al., 2009]. Much of the work investigating topology has been restricted to a small convention space (typically with just two actions). More recent work has explored the effect of increasing the number of available actions and has shown that doing so typically increases the time taken for convergence [Franks et al., 2014; Griffiths & Anand, 2012; Salazar et al., 2010].

Because of this effect of topology, [Hasan et al., 2014] introduced Topology-Aware Convention Selection, an intelligent strategy update rule that makes use of local information to decide the best strategy update rule to use in a given situation. They show that this approach outperforms the individual strategy update rules in both random and scale-free networks and this lends credence to the idea that local information can be sufficient to effect convention emergence, something we explore in Chapter 4.

2.3 Manipulation of Conventions

The previous section examined the work in the literature regarding the various mechanisms that allow conventions to emerge and how the speed and quality of convention emergence can be affected by the parameters of these mechanisms and by the underlying topologies. Each of these approaches does not provide an answer to one major question though: what if you want to ensure a specific convention is the one that emerges?

The concept of *manipulating* convention emergence to ensure a desirable outcome between the multiple possible equilibria is one that has seen relatively little work dedicated to it. Kittock [1994] utilised the notion of *authority*, giving some agents the ability to ignore feedback from those "beneath" them whilst still being able to propagate feedback to them. Kittock shows that this approach dramatically changes the convention emergence within the system with the strategy of the authoritative nodes being "pushed" to their descendent.

The use of fixed strategy agents, who always choose the same action regardless of others' choices, to influence convention emergence has also been explored. Sen & Airiau [2007] show that a small number of such agents can cause a population to adopt the fixed strategy as a convention over other equally valid choices with high probability. They consider the case where agents must choose which side of the road to drive on, modelled as a coordination game. They show that, with a population of 3000 agents, having only 4 fixed strategy agents using the strategy “drive on the right” causes the entire population to nearly always converge to this strategy. This indicates that small numbers of such agents can affect much larger populations.

In Sen and Airiau’s model, due to the lack of connecting topology, all agents are identical in terms of their ability to interact with others. However, in many domains, agent interactions may be limited to neighbours in the network. As such, some agents will have larger sets of potential interactions than others. In the context of static topologies, Griffiths & Anand [2012] establish that which agents are selected and *where* they are in the topology is a key factor in their effectiveness as fixed strategy agents. They show that placement by simple metrics such as degree offers better performance than random placement and increases the ability of the fixed strategy agents to influence the population.

Franks *et al.* [Franks et al., 2014; Franks et al., 2013] investigated fixed strategy agents where interactions are constrained by a static network topology and agents are exposed to a large convention space. They found that topology affects the number of fixed strategy agents required to increase convergence speed. This also expanded on the work of Griffiths & Anand [2012] by investigating the effectiveness of placing by more advanced metrics such as eigenvector centrality. They showed that, even in much larger populations and with a much larger convention space, these *Influencer Agents* were able to elicit convention emergence towards the desired convention.

The ability of these fixed strategy agents to direct and encourage convention emergence by being placed at influential locations is the primary focus of this thesis. The work of Sen & Airiau and Franks et al. is foundational in this regard and the work presented here builds on much of their work, expanding and extending it by applying fixed strategy agents to new domains and problems, and we will revisit the importance of these works in Chapter 6. We examine the effect of these fixed strategy agents in causing convention emergence by allowing them to *intervene* in the process. As such, we call them Intervention Agents (IAs) from this point onwards.

2.3.1 Convention Emergence Threshold and Intervention Agents

When including IAs in a population to elicit convention emergence, we must consider how they affect the calculation of said emergence. In the case with no IAs we can use the Kittock Criteria [Kittock, 1995] to establish a percentage threshold of the population that has converged to the same action choice. When this threshold is surpassed then we consider a convention to have emerged. Previous work subscribes to this definition with various percentages being used as a threshold: 90% [Delgado, 2002; Griffiths & Anand, 2012; Kittock, 1995] and 95% [Airiau et al., 2014] are common.

However, there is seemingly no consensus on how the inclusion of IAs should affect this calculation. As agents that will always choose the same action, regardless of other agents' choices or detrimental effects, these fixed strategy agents can either boost the proportion measured for convention emergence (if they are using the action that matches that of the rest of the population) or reduce it (if they are using a differing action). However, these agents are unable to change their strategy and as such are distinct from those members of the convention who have learnt or converted to it. On the other hand, as part of the population, the actions used by the IAs are still choices within the system and, as far as the traditional Kittock Criteria is concerned, represent selections.

Previous work is divided on how to consolidate these two seemingly disparate ideas. Airiau et al. [2014] and Sen & Airiau [2007] continue to count the "choices" of the fixed strategy agents as though they were actively choosing when comparing to the Kittock Criteria whereas Griffiths & Anand [2012] choose to only count the proportion of non–fixed-strategy agents adhering to the convention. Both of these have the effect of changing the number of non–fixed-strategy agents that must join the convention before it is considered emerged compared to the case with no IAs, sometimes to a significant level.

To illustrate this, we can consider a graph of 1000 nodes ($n = 1000$) with 100 IAs ($f = 100$) and the Kittock Criteria set to 90% (or $K = 0.9$). We consider the three cases:

No IAs 900 agents of the population must become members of the convention.

Count the IAs as choosing their fixed strategy With 100 IAs already "in" the convention 800 additional agents must become members.

Exclude the IAs from the proportion calculation With 900 non–fixed-strategy agents, in order to reach 90% convention emergence, 810 of those 900 must join the convention.

These differences, 100 and 90 fewer agents (or 10% and 9% of the population respectively) highlight the difficulties with both approaches: they reduce the additional number of agents that must be persuaded to join the convention. As the research on convention emergence is concerned with how conventions behave in populations of specific size, this artificial depression of the number of agents that must be converted will affect the patterns and behaviours found. In previous work this has not been a concern due to the low numbers of IAs used but, as we seek to explore destabilisation as well as emergence, the work presented in this thesis is likely to use numbers of IAs where the effect becomes noticeable.

One solution is to consider IAs as additional – agents who are not part of the base population but are instead inserted as extra agents. However, this does not translate well to scenarios with underlying topologies (discussed below) which are the main focus of the following chapters. Inserting additional agents into the topologies would require making additional artificial edges and the creation and number of these will affect the simulation.

Instead, unless noted otherwise, we address this problem by allowing those agents who are chosen as IAs to continue to learn from their interactions with others. When interacting they will still use the fixed strategy assigned to them (as the IAs in Griffiths & Anand [2012] and Sen & Airiau [2007] do) but for the purposes of counting members of a convention we query them as we would any other agent, finding what strategies they would choose when not exploring. This addresses both concerns raised above: the IAs are counted as part of the population as in the work of Kittock [1995] but will not artificially inflate the number of agents in the convention if they don't believe it to be the best choice.

Additionally, this in effect changes the measurement of convention emergence from *a posteriori* (what *did* agents choose) to *a priori* (what *would* they choose, all other constraints removed). This small distinction allows us to measure convention emergence even in settings with high levels of exploration (selecting what the agent considers to be a non-optimal action in order to ensure that this consideration is true), an aspect not considered by Kittock [1995]. Additionally, it allows the learnt behaviour from agents' last interactions to be included in the convention emergence evaluation, unlike the approach taken by Airiau et al. [2014] and Sen & Airiau [2007] who measure what *was* chosen and do not consider that agents' choices may have changed because of this interaction. In this manner we are effectively looking at what agents will do during the next timestep rather than what they have just done. In the context of real-world applications of convention emergence, we believe this approach is more sensible; if a customer has been regularly using your services but is unlikely to do so next time this is more valuable information than knowing once they've stopped using it. Overall though, this change is minor due to the timescales involved but offers much greater flexibility in approach. In particular, when we consider the non-permanent inclusion of IAs in Chapter 5 measuring what agents wish to do next is more important for conventions than measuring what they have already done.

In initial intervention (using IAs to encourage convention emergence at the start of a simulation), depending on the learning mechanism used, the behaviour of our approach is unlikely to differ much compared to the approach of Sen & Airiau [2007] where they explicitly count the IAs as members of the convention.

This is because the IAs will not necessarily have had time to form opinions about the benefits of other potential conventions and thus would only wish to choose differently from the fixed strategy if the fixed strategy is rejected by other agents. Given the performance attributed to IAs in the literature [Airiau et al., 2007; Franks et al., 2013; Griffiths & Anand, 2012] this is unlikely but our approach offers the flexibility to address it if necessary.

When we are considering destabilisation of existing conventions however, the difference is likely to become more marked. Given the self-reinforcing nature of conventions [Boyer & Orléan, 1992; Goyal & Janssen, 1997] we expect larger numbers of IAs to be required and the considerations noted above become more prominent. We hypothesise that allowing the IAs to continue to learn from the time of creation will enable them to learn the chosen intervention strategy as "good" due to their interactions – assuming destabilisation occurs. Additionally, in the most extreme case, many of them may still view the original dominant strategy as "good" due to not using it in interactions since becoming IAs. This is a true representation of the agents' beliefs based on their interactions and is important to accurately reflect in the proportion calculation. If the agents were to cease being IAs they would act on these beliefs and this is something that may occur in dynamic networks (as we explore in Chapter 3) or if the presence of IAs is only temporary (as we explore in Chapter 5).

This means that we would expect some of the IAs to continue to "choose" between the intervention strategy and the previously dominant one when queried. In the case where the agent was a prominent adherent of the dominant strategy (which is likely given the nature of dominance), we might expect it to choose between the two 50:50. As such, the number of non–fixed-strategy agents required to convert to the intervention strategy will still be reduced, but not artificially. The accurate measuring of agent intentions and beliefs means that this reduction will be smaller than that described above. Returning to the previous example of a 1000 agent population, we would expect, at any given moment, 50 of the 100 IAs to be choosing the intervention strategy. As such, 850 other agents must

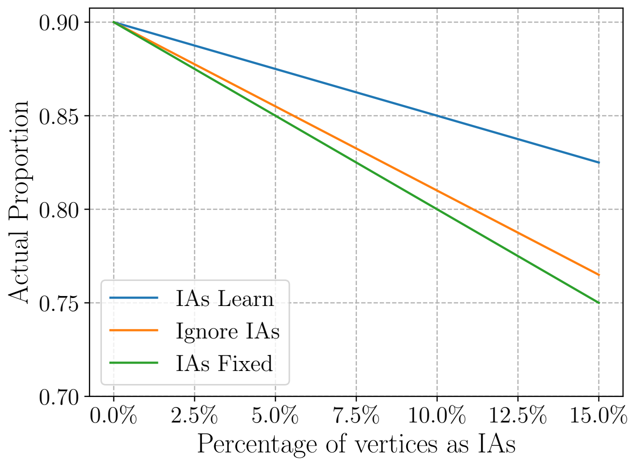

Figure 2.1: The actual proportion of agents in the population that must be converted in order to reach 90% convention emergence when using each of the three approaches for including IAs in the Kittock calculation.

also be members of the convention to reach the 90% threshold. This is much closer to the actual 900 agents we would hope to treat as the threshold and, due to the ability of the IAs to learn, actually represents 900 agents or 90% of the population who believe in the convention.

With each of these approaches, the difference between the desired proportion of the population that are members of the convention and the actual proportion that must be reached beyond the IAs varies with the number of IAs placed into the system. Figure 2.1 shows the actual proportion that must be reached against the percentage of the population being used as IAs for the three different approaches. As can be seen, the approach we have described above has a marked benefit over the other two. It represents a more accurate proportion of the population and this difference only increases at higher numbers of IAs. Whilst IAs will still affect the calculation, our approach aims to minimise this, particularly when considering destabilisation.

2.4 Games

The application and use of game theoretical constructs to model social interactions is well established in numerous fields including economics [Akerlof, 1980; Börgers & Sarin, 1997; Boyer & Orléan, 1992], linguistics [Meara, 2006], law [McAdams, 2008] and sociology and politics [Brams, 2011; Snidal, 1985]. Modelling potentially complex human interactions as games allows higher levels of mathematical analysis and the abstraction from the actual interactions to focus instead on the outcome allows these games to be generic and applicable to numerous situations.

As Schelling [1980] argues, nearly all games can be placed somewhere on the spectrum of *pure conflict games*, where one agent's gain is another's loss, and *pure coordination games*, where agents win or lose together and equally. Within this spectrum there are distinct notions of other types of games including those that require *cooperation*, where agents must mutually make choices in order to provide the best outcome for all, and games where coordination does not result in equal benefit to each of the agents. In this section we provide an overview of the different categories and types of games frequently encountered in the MAS community and discuss the properties and distinctive features of each.

An n-action-k-person game is one played by k agents at the same time, each of whom has n action choices available to them. Each agent makes a choice from the n available actions and the combination of their choices (known as the *joint strategy* or *joint action*) determines the payoffs that each agent receives. These are generally described in a *payoff matrix*, an example of which is shown in Table 2.1. In this table we are looking at a 2-person-2-action game with both agents choosing from the set of actions $\{A, B\}$. The payoffs received are specified in the cell that aligns with the joint action. The left hand, uppercase letter represents the payoff of the row-player with the right hand, lowercase letter representing that of the column-player.

As in the bulk of the literature [Delgado et al., 2003; Kittock, 1995; Shoham

	A	**B**
A	X,x	U,u
B	V,v	Y,y

Table 2.1: General form of a 2-person-2-choice payoff matrix

	A	**B**
A	x,x	u,v
B	v,u	y,y

Table 2.2: General form of a 2-person-2-choice symmetric game payoff matrix

& Tennenholtz, 1997], in this thesis we focus on 2-person *symmetric* games for modelling agent interactions. A symmetric game is one taking the form shown in Table 2.2 and is thus called due to the fact that the payoff matrix is symmetrical for the row- and column-players; the reward a player receives from a specific joint action is not reliant on whether they are the row-player or column-player. This aspect is important because it means there are no different roles or identities that affect agent payoff, only their choices. This allows us to treat agents as heterogeneous without special consideration for their role in a given interaction and allows convention emergence in the population to be described as convergence to a simple action choice. Whilst this may seem limiting, requiring agents to have access to exactly the same actions, Lewis [1969] notes that problems can often be suitably reframed such that *corresponding* but not identical actions are described in such a way that these actions become identical. This assumption is not general however with Sen & Airiau [2007] allowing agents to learn and emerge conventions separately as both row- and column-players. The restriction to 2-player games is borne from the expectation that many agent populations will be situated in a network topology (this is discussed in Section 2.6). As each edge only links two agents, interactions are limited to this as well.

An important notion in all games that are expected to emerge consistent behaviour, and one that underlies the problems of coordination and cooperation, is that of *Nash equilibria* [Nash, 1950]. A Nash equilibrium is a joint action that is stable in the sense that no single agent can increase their payoff by changing actions if no other does. Thus, no agent has any reason to unilaterally change their action without assurances of some form that others will do so too. In

	C	**D**
C	-1,-1	-3,0
D	0,-3	-2,-2

Table 2.3: Prisoner's Dilemma payoff matrix

	S	**H**
S	10,10	0,8
H	8,0	7,7

Table 2.4: Stag Hunt payoff matrix

scenarios where games are repeated (what Shoham & Tennenholtz [1997] call "stochastic social games") the Nash equilibria are particularly important as they represent joint actions that are local optima, in terms of payoff, in the joint–action-space. Encouraging agents to move away from these equilibria is one of the fundamental questions of convention emergence and cooperation.

Lewis [1969] introduces an extension of Nash equilibria, what he defines as a *coordination equilibria*. Whilst a Nash equilibria is a combination of actions in which no agent would be better off if they unilaterally changed their action, a coordination equilibria is a combination of actions in which *no one* would be better off if *anyone* unilaterally changed their action. All coordination equilibria are thus Nash equilibria but Nash equilibria are only coordination equilibria if no agent could increase the payoff of themselves or another with a different choice. This distinction, between coordination and *non-coordination equilibria*, as we shall see, underlies the difference between situations and games that can be solved purely by coordination and those which require cooperation to ensure agents receive the best possible payoff.

A dominant strategy is one which an agent always prefers as it always gives the best payoff possible regardless of the choices of others. Because of this it is necessarily the case that combinations of dominant choices must be an equilibrium by definition. However, it is not guaranteed that this will also be a coordination equilibrium. This is best highlighted in the classic Prisoner's Dilemma game, the payoff matrix for which is shown in Table 2.3. In the Prisoner's Dilemma the players are given two options, cooperate (with each other) or defect. If both cooperate they will receive a small punishment. If one defects and the other cooperates the one who defected will receive no punishment, the one who cooperated a larger punishment. If both defect, they will both receive punishment greater than that they would have received if both cooperated. The Nash equilibrium in the Prisoner's Dilemma game is for both agents to defect (combination $< D, D >$). However it is not a coordination equilibrium as a change in strategy for either player will result in the other player's payoff improving. Intuitively, this means that each agent wants their opponent to change strategy to increase their own payoff (to the detriment of the opponent's payoff). Indeed, it is this property, the lack of a coordination equilibria, that distinguishes the Prisoner's Dilemma as a game *requiring* cooperation to increase agent payoff; rational agents (those which seek to maximise their own payoff) have no joint action available to them that it is logical to choose and that no-one would wish to change from. In order to make rational agents in the Prisoner's Dilemma choose $< C, C >$ (which increases both player's payoffs) there needs to be cooperation amongst them, potentially in the form of externalities such as norms [Shoham & Tennenholtz, 1997], control, or incentives and sanctioning [Villatoro et al., 2011b] of defecting players), or allowing them to learn over time that continued use of Defect when all others are doing it will result in worse payoffs.

The Iterated Prisoner's Dilemma as described by Axelrod & Hamilton [1981] highlights these issues. Axelrod shows that the best strategy is that of "tit-for-tat with forgiveness". Agents begin by choosing Cooperate and then simply repeat the action taken by their opponent, treating defection with defection and cooperation with cooperation. The aspect of "forgiveness" is necessary however to avoid a cycle of defection and shows the difficulty of naturally converging to pure cooperation. The work of Kittock [1995] and Shoham & Tennenholtz [1997] also address these issues. Both make use of HCR to allow agents to learn behaviour of multiple interactions but the conclusions of both works is that stable convergence to the cooperate strategy is difficult and the time to do so grows exponentially with the number of agents in the population. Ullmann-Margalit [1977] go even further, noting that, whilst cooperation is an equilibrium in the Iterated Prisoner's Dilemma it "enjoys what might be termed *precarious stability*" waiting only for some agent to exploit defection. Thus the norms and other mechanisms to enable cooperation must continue to be used within the system.

Cooperation can also be required in games that do not exhibit the features found in the Prisoner's Dilemma. The game known as Stag Hunt is depicted in Table 2.4. The Stag Hunt is often considered as similar to the Prisoner's Dilemma [Fang et al., 2002; McAdams, 2008; Ullmann-Margalit, 1977] as both represent the necessity of trust and cooperation in order to maximise the payoff for the agents. The game represents a situation of two hunters who can hunt either hares (H) or a stag (S). A single hunter is unable to successfully hunt the stag alone and the other can exploit his absence to hunt more hares. But if both hunt the stag they will receive a higher payoff than they can get by hunting hares. The primary difference between this and the Prisoner's Dilemma is that the cooperative strategy $< S, S >$ is a coordination equilibrium but the defecting strategy $< H, H >$ is still a Nash equilibrium. The Stag Hunt game thus highlights the problem of *trust* as, assuming your opponent is rational, there is no reason that $< S, S >$ would not be chosen. If your opponent is antagonistic however, or untrusting, then they may not expect you to choose S and doing so themselves would result in a worse payoff than if they choose H [Fang et al., 2002]. In iterated interactions against differing opponents, such as we concern ourselves with, the likelihood is increased as the agent is unaware how the other's previous games have gone. In this scenario, Rankin et al. [2000] found that many players may instead tend to the "risk-dominant" equilibrium of $< H, H >$ which has an expected return of 7.5 if you are unsure of your opponent's choice compared to the "payoff-dominant" equilibrium of $< S, S >$ which only has an expected payoff of 5. As noted by Ullmann-Margalit [1977] the Stag Hunt game is thus likely to still require norms to ensure that the cooperative behaviour prevails, although to a lesser extent than the norms required to ensure cooperative behaviour in the Prisoner's Dilemma case.

	0	**1**	...	**n**
0	4,4	-1,-1	...	-1,-1
1	-1,-1	4,4	...	-1,-1
⋮	⋮	⋮		⋮
n	-1,-1	-1,-1	...	4,4

Table 2.5: n-action coordination game payoff matrix

The previous discussions have highlighted the requirement of cooperative behaviour to produce the desired mutually beneficial outcomes for all agents of a population in these games. However, for the work in this thesis we seek to highlight how conventions can be directed by the inclusion of IAs. The nature of both the Prisoner's Dilemma and Stag Hunt games means that there is pressure on the agents, that can be curbed through enforcing cooperation, to converge to one particular solution either due to a lack of coordination equilibria or the presence of a risk-dominant one. In order to explore scenarios where this effect is not influencing the convention emergence we focus on pure coordination games which inherently address a number of these issues.

A *pure coordination game* is one in which agents' interests are perfectly aligned. In game theoretic terms this means that the two payoffs are equal in every square of the payoff matrix; each agent receives the same payoff for the action combination, there is no situation in which one agent benefits more than the other. An example of the 2-player-n-choice coordination game is shown in Table 2.5.

The pure coordination game thus has a number of properties that make it amenable to the study of convention emergence: all of its Nash equilibria are also coordination equilibria and all of them are Pareto optimal[1]. This means that there is no preference amongst the agents for which one becomes the convention, only that one of them does. No equilibrium is risk-dominant or payoff-dominant, all provide the same payoff to the agents. The only aspect that differentiates which of the equilibria is of more benefit to the agents is the precedent of which

[1] A joint action is *Pareto optimal* if there is no other action combination that increases the payoff of one agent without decreasing the payoff of another.

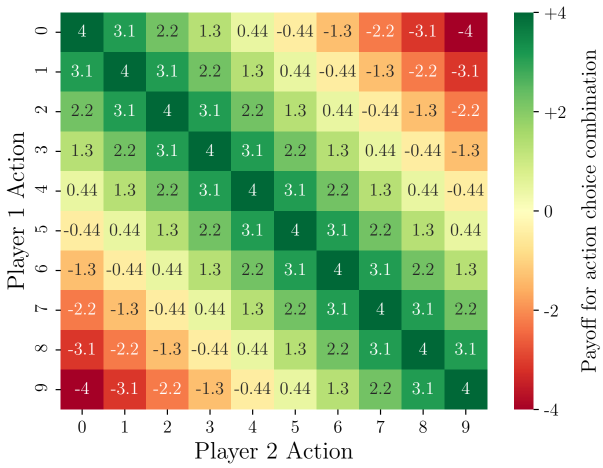

Figure 2.2: Gradient Payoff matrix for a 10-action game with $pay_{max} = 4$ and $pay_{min} = -4$.

is being used by other agents they interact with; there is no salient difference between them. Thus the pure coordination game allows us to focus on the nature of convention emergence itself rather than on underlying features of the game. These features have made the pure coordination game prevalent amongst the convention emergence literature [Franks, 2013; Griffiths & Anand, 2012; Kittock, 1995; Sen & Airiau, 2007] and, unless mentioned otherwise it is what we utilise throughout this thesis.

2.4.1 Gradient Coordination

In many real-world interactions there is unlikely to be a hard divide between entirely compatible and entirely incompatible action choices as is the case in the pure coordination game. Many interactions can be solved by compromise or making use of the aspects of the action choice that are compatible and negating or ignoring those which are not. Consider the overlap in agents choosing to speak British English and American English language. Whilst not entirely compatible (and hence not truly the same action) there is enough similarity that interaction is possible. Whilst the language game [Lakkaraju & Gasser, 2008; Steels, 1998] can be used to address this problem, we wish to examine this notion of "similar enough" in the context of the coordination game. To this end we extend the notion of the pure coordination game to investigate the effect that a *gradient* of payoffs has on the ability for conventions to emerge.

We define the payoff that agents receive from their joint action in the gradient payoff as:

$$\mathit{Gradient\text{-}Payoff}(i, j, n, pay_{max}, pay_{min}) = pay_{max} - \frac{dist(i,j)}{n-1}(pay_{max} - pay_{min}) \tag{2.1}$$

where i is the action of player one, j is the action of player two, n is the number of actions, pay_{max} and pay_{min} are the maximum payoff (same action chosen) and minimum payoff (most different actions chosen) respectively, and $dist(i, j)$ is the distance between the two actions when they are ordered by similarity. Figure 2.2 shows a representative payoff matrix created in this manner for the 10-action game with maximum and minimum payoffs of $+4$ and -4 respectively.

2.4.2 Other Games

A number of other games have been explored within the notion of convention emergence that do not conform to the traditional definition of a coordination game as proposed by Lewis [1969] and Schelling [1980].

The El Farol Bar problem is an example of a *congestion problem* where the payoff that the agent receives for a particular choice is related to the number of other agents choosing the same. The problem as originally formulated by Arthur [1994] is that of attending a popular bar on a given night or staying home. If the bar is not too crowded (60% of its capacity) the player enjoys themselves more than if they had stayed at home. If the bar *is* too crowded the player does not enjoy themselves and would have been better off remaining at home. It is impossible to definitively say ahead of time how many others will be attending the bar on the same night, all players decide whether to attend or not at the same time. As such players must rely on their expectations (and higher order meta-expectations [Lewis, 1969]) on the behaviour of others in order to deduce whether to attend or not. The optimal outcome to the El Farol Bar problem is one of several conventions existing simultaneously with subsets of agents (of size below the crowdedness threshold) each electing to attend on different days. Systems naturally emerging this outcome have been shown to be unlikely however. Cara et al. [1999] show that most initial conditions tend to the case where agents have failed to coordinate the optimal behaviour amongst themselves and instead perform no better than at random. Whitehead [2008] model the El Farol Bar problem using a reinforcement learning model. They find that the agents generally minimise their bad experiences and maximise the good ones when learning in this manner but show that the resultant behaviour of this is a partitioning of the agents into two distinct sets of those who always attend and those who never attend rather than the mutually beneficial solution of attending on disparate days. In this regard both show that the El Farol Bar problem is another type of game that would benefit from external cooperation or coordination.

The language coordination game (or naming game [Steels, 1998]) represents the emergence of a communicative capability between agents using conventions and norms. Each agent has a set of words and a set of concepts and an internal mapping between them. They attempt to communicate concepts to other agents using the words they map to them and are rewarded if the other agent correctly understands their communication. In this regard it is similar to the pure coordination game; if agents both have the same mapping they are rewarded positively, if they differ they are rewarded negatively and in both cases the agent is able to learn from the interaction. The difference lies in the fact that the interaction only uses a small part of the lexicon (the mapping) that the agent has, individual parts of the lexicon are updated and changed rather than the whole action choice. This changes the underlying dynamic and means that multiple conventions are more likely as multiple different lexicons are present within the system. Additionally, it introduces the concept of the *quality* of a convention, with better mappings being better quality (the case of mapping each word to a single concept would make communication easier but be a poor quality lexicon for instance). The game also allows for much larger numbers of possible conventions than in other versions of the coordination game (W^C where W and C are the numbers of words and concepts respectively). Nevertheless, widespread convention adoption is possible. Salazar et al. [2010] utilise a spreading mechanism to transfer partial information and show that this can be used to encourage rapid convention emergence which produces high-quality conventions regardless of noise in the system. Whilst this differs substantially from the convention emergence approaches in other work it shows that encouraging conventions can happen even in large convention spaces with minimal adjustments. Lakkaraju & Gasser [2008] similarly show that convergence can be encouraged by using additional information to predict others lexicons. They utilise an additional "Text Observation Game" where agents are able to hear sentences formed using other's lexicons and use various metrics to predict the lexicon. However, both of these approaches require modifications to the agents to allow them to perform these additional tasks. Franks et al. [2013] instead take a different approach and instead insert agents with high-quality lexicons into the system to encourage emergence to them. They show that a relatively small number of these placed at influential locations can facilitate convention emergence to the desired convention.

Both of these games offer a distinct and differing view at convention emergence than that found from the pure coordination game. Whether the findings of this thesis are applicable in these disparate cases is something we discuss in Chapter 6.

2.5 Destabilisation and Meta-Stable Subconventions

Whilst the emergence of conventions in MAS has been previously well studied, there is little to no work in the literature on the notion of *destabilising* and replacing an already existing, established convention. There are numerous scenarios where it may be desirable for a system designer to be able to replace existing conventions:

- The established convention has some salient difference from the convention that is desired by the designer such that, whilst there may be no difference in payoff, they wish to change the convention used by the system. Consider the notion of changing a technological standard such as a communication protocol. The system has little incentive to change if there is no direct benefit to doing so but an external force may desire the change.

- The established convention is sub-optimal in pay-off. This situation can arise in many games that require cooperation to most efficiently solve (see Section 2.4). Left to their own devices, the agents of a population may converge to a sub-optimal but consistent convention. Encouraging them to switch away from this to the more beneficial convention will require removal of the previously established one.

- The established convention was previously optimal but this has changed due to external conditions. In this instance the agents may continue to use the sub-optimal convention due to the level of precedence it has within the system; attempting to change unilaterally is likely to be harmful to agents who do so. Consider the scenario of changing the generally used charging mechanism from USB 2.0 to USB C. Individuals who switch without others doing so are likely to be hindered by a lack of compatibility between their devices and others. Unless large numbers of the population switch then there is a disincentive to doing so despite the fact that the new technology is superior.

As noted by Arrow [1974]: "It may be really true that social agreements ultimately serve as obstacles to the achievement of desired values, even values desired by all or by many. The problem is that agreements are typically harder to change than individual decisions... What may be the hardest of all to change are unconscious agreements, agreements whose very purpose is lost to our minds." Sugawara [2011] explores the stability of conventions in a Markov game and show that agents which act conservatively are unable to adapt to the emergent convention becoming undesirable.

Boyer & Orléan [1992] note that the powerful pressures to conform mean agents, individually or even in small groups, are unlikely to bring about change to the system by themselves. They describe a taxonomy for the limited cases where such a shift in population preference does occur and provide insight into how we might encourage it within MAS. They describe four ways in which a general population change may occur:

General Collapse A situation which indirectly destroys the overall structure of conventions throughout the system. This can take two primary forms: the utility of the current convention falls (though remains positive), allowing a small number of "mutants" to effect change; the utility of the current convention becomes negative or zero in which case the system will rapidly change even without mutants.

External Invasion A second population of individuals whom adhere to a different convention are brought into contact with the original population and the relative proportions are enough to reduce the benefits of the original population continuing to use their convention.

Translation If the two conventions exhibit a certain compatibility between them then the old convention can be "translated" in terms of the new and the population may shift because of this. They note that a series of *cumulative* transformations may lead to quite distinct original and final conventions with intermediate steps. This is similar to our notion of *gradient* that is discussed in Section 2.4.1.

Collective Agreement By intelligent and collective agreement the population may note that the new convention is inherently better than the old and all agree to change behaviour. This requires a central authority.

It is the second of these which is the focus of the work presented in this thesis. We have already discussed how fixed strategy agents can be used to encourage convention emergence to a desire strategy. Expanding on this, we propose using similar agents in scenarios where conventions are already established to destabilise and replace the existing conventions.

Villatoro et al. [2009] introduced the notion of meta-stable subconventions, areas of a topology where the specific structure of the area self-reinforces a convention different to that of the majority of the population. They identified a number of different structures that are likely to exhibit this phenomenon in both scale-free and fully-connected star networks and noted that these meta-stable subconventions were persistent, hindering global convention emergence. Toivonen et al. [2009] independently examined similar notions in random graphs where they showed that small cliques[2] of nodes would form differing conventions to the wider population that were robust and difficult to change, what they term "dynamical robustness against invasion".

Meta-stable subconventions were also shown in ring topologies by Epstein [2001]. They showed that multiple separate regions of distinct conventions were able to emerge with little variation between their boundaries over time. As such the entire population was unable to converge to a single convention, consisting instead of "islands" of the two possible conventions.

Each of these cases show that these meta-stable subconventions can hinder wider acceptance of a majority convention and highlight the need to be able to cause a change in convention within even these parts of the system. Villatoro

[2] A subset of vertices such that every vertex is connected to every other in the subset.

et al. [2011a] utilise observation and rewiring of the edges in these structures in order to eliminate the meta-stable subconventions and show that this is effective in multiple different scenarios. However, the ability to unilaterally change the network topology is not consistent nor likely in many domains. Additionally, they only concern themselves with changing the strategies of the meta-stable subconventions. We widen the problem and look to destabilise conventions in the wider population without making assumptions about the ability to change topological traits. The work of Villatoro et al. here forms the basis of how robust, self-reinforcing conventions can be altered and the work in this thesis is a natural extension of that concept to the wider problem of destabilising conventions throughout a population. We will revisit the work of Villatoro et al. in Chapter 6 to discuss how the work has been expanded and utilised.

2.6 Network Topologies

Many MAS represent agent populations as being connected by an underlying topology that limits interactions to neighbours in the graph. This better models interactions in the real-world as it is rare that all individuals in a population know all others (for instance in social networks) or are necessarily able to communicate with them (due to geographical limitations, lack of knowledge or interest). Convention emergence in regular networks, those which assume total ability to communicate, have been well studied [Sen & Airiau, 2007; Shoham & Tennenholtz, 1992a; Walker & Wooldridge, 1995] and have been shown to have different features than those in topologies [Delgado, 2002; Delgado et al., 2003; Pujol et al., 2005]. As such, in this thesis we assume that MAS exist within a topology. In this section we present the topological models we will be using, both static and dynamic, as well as real-world datasets that provide interaction graphs.

2.6.1 Static Network Topologies

We utilise two types of static networks throughout: scale-free and small-world. These have been shown to have features that are present in many real-world networks [Lewis, 2006; Mislove et al., 2007; Travers & Milgram, 1969; Watts & Strogatz, 1998] and hence make them good substitutes when studying convention emergence. Both types of network are generated using the Java Universal Network/Graph Framework (JUNG)[3].

Scale-free networks

Scale-free networks were first described by Barabási & Albert [1999]. The most prominent feature of these networks is that their degree distributions follow a power-law. That is, there are a large number of lower degree nodes and exponentially fewer higher degree nodes. These high-degree nodes are often referred to as "hubs" and are highly influential locations within the network [Brautbar & Kearns, 2010; Maiya & Berger-Wolf, 2010]. Scale-free networks also exhibit the feature of *preferential attachment* where higher degree nodes are more likely to gain new edges than lower degree nodes, a feature which mimics real-world networks such as transport links (new locations are more likely to be connected to already well-connected locations) and social networks (those with high numbers of friends are more likely to make new ones).

The model of Barabási-Albert uses this in order to generate scale-free topologies. It takes a number of initial vertices, m_0, and a number of edges to attach to each new node, $m \leq m_0$. Nodes are then added to the network one at a time and m edges are attached to them. The other endpoint of these edges is chosen randomly from the existing nodes with probability proportional to their existing degree. This produces a scale-free network with a power-law distribution with $\gamma = 3$ and an average path length that increases logarithmically with the size of the network.

[3] Version 2.1.1, `http://jrtom.github.io/jung/`

Small-world networks

Inspired by the notion of "6 Degrees of Separation" [Travers & Milgram, 1969], small-world networks are those that are characterised by high levels of local connectivity (a node's neighbours are likely to be neighbours of each other) and that every node is able to reach every other in a small number of "hops". As found in Milgram's experiment [Travers & Milgram, 1969], human social links are well-modelled by this type of network as most of your immediate friends and family are more likely to know each other but there will also be distant social contacts that link your cluster to others.

Whilst there are several approaches to generating small-world graphs [Watts & Strogatz, 1998], we utilise the model put forward by Kleinberg [2000a] and Kleinberg [2000b]. In this model we start with a toroidal 2D lattice of nodes, each node being linked to its 4 neighbours in this lattice. We then add a number, l, of additional "long-range connections", from each node to another outside of these initial 4. The probability of which node the edge is connected to is inversely proportional to the distance between them on the lattice and is controlled by a clustering exponent, α. That is, $P(u, v) \propto D^{-\alpha}$. When $\alpha = 2$, such that the long-range connections follow an inverse-square law, Kleinberg [2000b] showed that a uniquely fast greedy algorithm for information propagation in the network existed with any other value of α having asymptotically much larger delivery times. This finding indicates that efficient navigation is only a feature for some values of α and is something to be considered when considering the spread of a convention.

2.6.2 Real World Networks

Whilst synthetic networks are good at generating the general features of many real-world networks (such as scale-free or small-world properties) there are still often failures of the synthetic networks in capturing certain aspects [Delgado, 2002]. In particular, Franks [2013] and Pujol et al. [2005] show that there

	Network		Largest WCC	
	$\|V\|$	$\|E\|$	$\|V\|$	$\|E\|$
CA-CondMat	23,133	93,497	21,363	91,286
Enron-Email	36,692	183,831	33,696	180,811
Twitter	81,306	1,768,149	81,306	1,342,296

Table 2.6: Original and Modified Network Sizes

are large differences in the clustering coefficients of synthetic and real-world networks and that these differences can substantially change the behaviours exhibited within them. To alleviate this we also include real-world networks when considering the nature of conventions in the work of this thesis.

We make use of three real-world networks from the Stanford SNAP datasets [Leskovec & Krevl, 2014]. These datasets represent a number of different methods of social interaction and, as such, each have different features allowing a wide-ranging look at the effects of real-world networks on agent populations. The three datasets chosen are: CA-CondMat [Leskovec et al., 2007], the collaboration network of the arXiv COND-MAT (Condensed Matter Physics) category; Email-Enron [Leskovec et al., 2009], the email communications between workers at Enron; and Ego-Twitter [McAuley & Leskovec, 2012], a crawl of Twitter follow relationships from public sources (for our purposes we ignore the directed nature of the edges). These datasets are used frequently in both convention emergence and influence spread research [Chen et al., 2014; Franks et al., 2014; Pei et al., 2015; Wang et al., 2016] as performance benchmarks.

For the purposes of monitoring convention emergence in these networks, we only want to examine a single, connected component. As such, all 3 networks were reduced to their largest weakly connected component (WCC). Additionally, any self-loops (edges from a node to itself) were removed as such edges artificially inflate a node's degree whilst not increasing its ability to influence others. Table 2.6 shows the number of nodes and edges in each network and the number of nodes and edges (without self-loops) in their largest WCC.

2.6.3 Dynamic Network Topologies

Few studies have explored the notion of convention emergence in dynamic topologies despite the fundamental differences that allowing nodes and edges to be added and removed brings to the network dynamics. Some work has been performed in the related field of norm emergence but is primarily concerned with essential rather than conventional norms [Mungovan et al., 2010; Savarimuthu et al., 2007]. Savarimuthu et al. [2007] show that norms are able to emerge under a number of conditions and settings of dynamic topologies, but their work differs from ours due to the requirements placed on agents. The interaction model used requires agents to maintain an internal norm as well as being able to query other agents. We make minimal assumptions about agent internals or the information available. Additionally, our work investigates the *manipulation* of convention emergence, something not considered by Savarimuthu et al. for norms.

Mihaylov et al. [Mihaylov et al., 2014] briefly consider convention emergence in dynamic topologies using the coordination game. However, their work focuses on a new proposed method of learning, rather than on the emergence itself. In particular, they do not consider fixed strategy agents, or the action that emerges as a convention. In this thesis, we consider both convention emergence in dynamic topologies and the use of fixed strategy agents to understand the impact of network dynamics.

We seek to establish the performance and characteristics in dynamic networks compared to static ones and as such make use of two dynamic topology generators throughout this thesis.

González model

We utilise a particle-based simulation, developed by González et al. [González et al., 2006a; González et al., 2006b], to model dynamic network topologies with characteristics comparable to those observed in real-world networks. Agents are represented as colliding particles and the topology is modified by collisions creating links between the agents. A population of N agents, represented as a set of particles with radius r, is placed within a 2D box with sides of length L. Initially, all agents are distributed uniformly at random within the space and are assigned a velocity of constant magnitude v_0 and random direction.

Each timestep, agents move according to their velocity and detect collisions with other agents. When two agents collide, an edge is added between them in the network topology if one does not already exist. Both agents then move away in a random direction with a speed proportional to their degree multiplied by a speed factor, $\bar{v}$. Thus, higher degree nodes have an increased probability of further collisions, which in turn further increases their degree. In this way, the model exhibits preferential attachment, a characteristic found in static scale-free networks [Barabási & Albert, 1999].

Additionally, all agents are assigned a Time-To-Live (TTL) when created. This is drawn uniformly at random between zero and the maximum TTL, T_l. After each timestep agents' TTLs are decremented by one. When an agent's TTL $= 0$ the agent and all its edges are removed. A new agent is placed at the same location within the simulation with the randomised initial properties discussed above. In this manner, the topology is constantly changing.

Different topologies can be characterised by the value of T_l/T_0 where T_0 is the characteristic time between collisions. This can be expressed as:

$$\frac{T_l}{T_0} = \frac{2\sqrt{2\pi} r N v_0 T_l}{L^2} \tag{2.2}$$

González et al. show that this value dictates key characteristics of the generated topology, primarily the average degree and degree distribution and in [González et al., 2006b] they show that these are good approximations of the actual social networks amongst students in a school.

The concept of a quasi-stationary state (QSS) is discussed by González et al., such that a QSS emerges after a number of timesteps and is characterised by macro-scale stability of network characteristics. Micro-scale characteristics, for individual agents, remain in flux. In [González et al., 2006a] it is shown that the QSS can be described as any timestep, t, where $t \gtrsim 2T_l$.

Ichinose model

Ichinose et al. [2013] build a dynamic network as an extension of the Barabási-Albert model. As such it has the same useful features of that model, namely a scale-free nature and short average path length. The presence of these features will allow comparison between the similar static and dynamic networks in terms of the convention emergence upon them.

The Ichinose model begins by building a Barabási-Albert graph of the requested size with the same parameters, m, m_0. Then, each iteration, a node is removed and a new one inserted into the topology with the number of edges in the system kept the same. In this manner, both the number of nodes and edges in the topology will remain constant with only their arrangement changing.

Ichinose et al. specify 2 different methods of node removal: *targeted*, where the highest degree node in the topology is removed, and *random*, where a node is removed at random. The degree of the removed node is noted as n. A new node is then created with degree m or n, whichever is lower. The edges of these nodes are then attached to others using two other methods: *preferential*, where the node is chosen with probability proportional to their degree (as in the Barabási-Albert model) and *random*, where the node is chosen uniformly at random. If the number of edges in the network is less than it was before, edges are inserted into the system from source nodes (chosen uniformly at random) to target nodes which are selected the same way as the new node's edges, preferentially or at random.

Thus there are four possible modes that the Ichinose model can operate in:

- *Random* removal and *Random* addition **(RR)**
- *Random* removal and *Preferential* addition **(RP)**
- *Targeted* removal and *Random* addition **(TR)**

- *Targeted* removal and *Preferential* addition **(TP)**

In the paper the model is introduced they investigate the effect of their different topologies on the robustness and stability of the cooperation in the Prisoner's Dilemma showing that the various modes have drastically varying effects on the level of cooperation with targeted removal reducing cooperation even with a low benefit to defection. They evaluate the model's effect on numerous graph metrics compared to the original Barabási-Albert topology and show that all four reduce the degree variance in the network and shift the degree distribution with the modes of targeted removal drastically decreasing the degree variance and maximum degree. These different behaviours highlight that the 4 settings produce dramatically different topologies and as such can be expected to have distinct influences on convention emergence.

2.7 Network Metrics

When attempting to encourage or direct convention emergence using fixed strategy agents their location within the underlying network topology is important. Nodes that are located at influential positions, such as nodes with high centrality scores, have been shown to drastically increase the effectiveness of fixed strategy agents placed at those locations as discussed above [Franks et al., 2014; Griffiths & Anand, 2012]. Finding these influential locations has been the subject of wide-ranging research in the graph theory community [Borgs et al., 2012a; Cooper et al., 2012; Lawyer, 2015] and has applications in many fields and problems in MAS.

There are many different network metrics used to find influential nodes under various criteria and we present a non-exhaustive selection of these below.

1. **Degree Centrality (Degree)** Degree centrality (more commonly referred to simply as degree) is the size of the neighbourhood of a given node, $|N(v)|$, the number of edges the node has within a topology. Degree can be easily calculated from only local information and is intuitive in its

measure of influence; the more nodes that v can directly interact with, the more capability it has to propagate its views.

2. Eigenvector Centrality (EC) Also known as eigencentrality, eigenvector centrality measures the influence of a node based on the influence of nodes it is connected to. It assigns relative scores to each node, v, with connections to higher-scoring nodes contributing more to the score of v than connections to low-scoring nodes. Intuitively, eigencentrality assumes that a node is important if it is connected to other important nodes. Variations of eigencentrality include Google's PageRank algorithm [Page et al., 1999] and show the validity of this assumption as the effectiveness of PageRank in determining useful web pages is well-documented.

Eigencentrality can be calculated rapidly using the adjacency matrix of the graph in question, $A = (a_{v,u})$ where each entry is 1 if an edge exists between nodes v and u and 0 otherwise. The eigencentrality score $EC(v)$ can thus be found as,

$$EC(v) = \frac{1}{\lambda} \sum_{u \in N(v)} EC(u) = \frac{1}{\lambda} \sum_{u \in G} a_{v,u} \times EC(u) \tag{2.3}$$

where λ is a constant.

3. Betweenness Centrality (BC) Is a centrality measure that ties influence to the number of times a node, v, is found on shortest paths between any two other nodes. It can be calculated as follows,

$$BC(v) = \sum_{s \neq v \neq t \in V} \frac{\sigma_{st}(v)}{\sigma_{st}} \tag{2.4}$$

where σ_{st} is the number of shortest paths between nodes s and t and $\sigma_{st}(v)$ is the number of these which include v. As such, betweenness centrality thus represents how much information that flows through the shortest paths of the network must pass through v. As the shortest paths are the optimal ones through which to transmit said information, this gives vertices with high betweenness centrality increased capability to act as "gatekeepers", influencing the information that is passed through them. Whilst the calculation is a simple proportion, finding the shortest paths between all pairs of vertices in a graph is computationally expensive, requiring $O(V^3)$ time in the general case. For unweighted graphs, Brandes [2001] have shown that it can be done in $O(VE)$ time but this is still exponentially larger than most other metrics.

4. Closeness Centrality (CC) The closeness centrality (or simply closeness) of a node, v, is a measure of how "close" the node is to all others in the graph. In other words it is a measure of the average shortest path length between node v and all others. It was first defined by Bavelas [1950] as,

$$C(v) = \frac{1}{\sum_{u \in V} d(u, v)} \tag{2.5}$$

where $d(u, v)$ is the *distance* (the length of the shortest path) between nodes u and v. Closeness centrality is often used as a measure of influence [Borgatti, 2005; Lawyer, 2015] due to the ability of a "close" node to readily reach all other nodes within a network. When transmitting information along shortest paths, the maximally close node in the network is best positioned to do so. However, due to the calculation of shortest paths, closeness centrality is similarly computationally expensive as is the case with betweenness centrality.

5. Highest Edge Embededdness (HEE) Edge embeddedness was proposed by Easley & Kleinberg [2010] as a measure of how clustered the endpoints of the edge are. That is, how many common neighbours the two endpoints share. Edges with high edge embeddedness thus represent the most direct path amongst many between the two nodes whilst edges with low edge embeddedness represent pseudo-"bridges", only one of a limited number of paths through the local neighbourhoods of the two nodes. Formally, edge embeddedness is defined as,

$$EE(e_{uv}) = |N(u) \cap N(v)| \tag{2.6}$$

where $N(x)$ is the set of neighbours of node x in the graph.

However, as we are concerned with finding influential *nodes* not edges, we must adjust the use of edge embeddedness to allow this. We follow the approach of Franks et al. [2014] and assign to a node the highest value of edge embeddedness found in its set of edges. We call this highest edge embeddedness (HEE). A node with a large HEE value thus has an edge that is well-embedded in its local cluster and represents a node that should be able to readily influence not just the node at the other end of the edge but those shared neighbours as it has multiple short pathways to do so.

6. **Hyperlink-Induced Topic Search (HITS)** Like PageRank, hyperlink-induced topic search (HITS) is an iterative algorithm designed to exploit the nature of links on the web to find influential nodes, particularly in the use case of trying to find authoritative sources of information in web searches. Designed by Kleinberg [1999], HITS works by iteratively learning two sets of nodes: *hubs* which are well-linked directories of information but are not necessarily authorities on said information and *authorities* which actually contain the information. Hubs are ranked highly if they point to many authorities whilst authorities are ranked highly if they are pointed to by many hubs. The algorithm converges the two scores for each node until a final value is reached, a combined measure of both how well-viewed as a source of influence the node is as well as how well-linked it is to other sources of influence. As such, well-ranked nodes in HITS should be able to propagate information and conventions both from others and to others, making them good candidates for focusing convention efforts.

These metrics and heuristics are but a small subset of those available when analysing influence within graphs but we believe that they represent the most likely candidates for managing node influence. This is backed up by the findings of Franks [2013] and Franks et al. [2014] who utilise these metrics amongst a total set of 14 others. Their findings were that degree, EC, HEE and HITS were those that were most strongly correlated with the influence of a node within the graph.

2.7.1 Temporal Metrics

With the consideration of dynamic network topologies in addition to the static ones, there is the additional consideration of the time-changing nature of these graphs. A highly influential node in one timestep may not be so the next. The notion of shortest paths also changes with the varying nature as paths may appear or disappear as the system progresses. Numerous extensions to the graph theoretic literature have been made in recent years to address these concerns and extend the notions of centrality and other metrics to encompass dynamic networks.

Kim & Anderson [2012] treat dynamic networks as special cases of static networks that exhibit directed flows between different layers of a meta-graph, one in which each snapshot of the dynamic graph at a particular time is treated as a separate sub-graph with links from nodes to later versions of themselves. Reducing the network in this way they are able to extend the notion of shortest paths and degree to those that exist between the temporal instances of the graph using the edges that pass between different times. They extend both betweenness centrality and closeness centrality to make use of these temporal shortest paths and compare them to aggregated or average versions of each metric over the timesteps. They show that the temporal metrics may exhibit different behaviours than the average or aggregate ones but that this is highly dependent on the actual underlying dynamic topology.

Nicosia et al. [2013] similarly address the issues of dynamic networks and additionally introduce the capability of finding connected components in temporal graphs. They extend the notions of connectivity in dynamic topologies taken from Kempe et al. [2002] and note that time-varying graphs can be modelled as graphs with an additional dimension, similar to Kim et al. They additionally provide another variation on the centrality metrics for dynamic networks, distinct from Kim & Anderson [2012].

Pan & Saramäki [2011] provide another model of dynamic graphs, focusing on the time-ordered temporal paths that exist within them. They highlight that spreading dynamic will be different because of the existence of these paths and explore the notion of time-dependent distance between nodes. They show that, whilst correlated to the the static distance between nodes there is large variance in this regard and thus nodes that are close in the static network may nevertheless utilise spreading that follows very different paths. They provide another formulation for the notion of closeness centrality in dynamic graphs based on the idea of temporal paths with time-cutoff.

As can be seen just from these approaches, there is still uncertainty in the literature on how to extend the notions of graph metrics to encapsulate dynamic topologies. Additionally, each of these approaches makes the assumption that the graph is fully observable in its temporal aspect. That is, you can see all timesteps of the graph and hence calculate the centrality metrics based on the shortest paths observed therein. However, in many applications, particularly ours, this is not the case. Convention emergence is monitored and encouraged on a per-timestep basis with no knowledge of what future graph topologies will look like. Because of this, we must instead focus on measures that can be calculated given that singular timestep. Whilst we could utilise all timesteps up until this point and disregard future changes to the graph this then becomes a prediction problem; trying to predict which nodes will continue to exhibit high levels of the metric in question at future times.

In this thesis, we therefore focus instead on the application of the traditional graph metrics within the effectively static graphs created at each timestep. Whilst we make some allowances for the time-varying nature of the metrics (see Chapter 3) we do not attempt to find nodes with high temporal centrality measures explicitly.

2.7.2 Other Metrics

As well as the traditional and temporal metrics discussed above, work in other fields that need to identify influential nodes have produced a number of variant heuristics for finding these nodes.

Lawyer [2015] uses the notion of epidemic spreading potential to identify influential nodes as those that have the highest expected "Force of Infection". He models potential epidemic spreading from each node and enumerates all possible clusters of infected nodes after x transmission events from this start node. The expected Force of Infection is then calculated as the entropy over each of these possible transmission outcomes. He shows that this metric accurately quantifies the spreading power of the target node and does so with better effect than other node metrics. Indeed, many models of contagion spreading offer insights in this field, even if the findings are not directly transferable. Finding which models offer applicable knowledge and using that to influence design decisions is a general problem, beyond the scope of this thesis however.

Chen et al. [2009] introduce the concept of Degree Discount as a way of selecting multiple high-degree nodes such that their joint influence is maximised. In the domain of the influence cascade problem, they show that selecting multiple high-degree nodes that are inter-connected causes redundancy in the overall influence capabilities of these nodes. They show that artificially discounting the degree of nodes when their neighbours have already been selected as influencers will ensure that overall influence in the system is maximised.

Both of these approaches are promising in that they outperform the standard graph metrics in their respective domains. Whether they are generally applicable or too heavily reliant on assumptions about the influence model being used is unknown and beyond the scope of this thesis. However, their performance shows that it is often possible to improve on the simplistic graph metric approach to identifying influential nodes by using specific domain knowledge, an approach we utilise in our extension of metrics for dynamic topologies.

Another set of metrics make use of the concept of network coverage, trying to find a minimal set of nodes that are linked to all others to provide a vertex cover. This problem is NP-hard, although approximation algorithms exist. One method of approaching this is to try provide coverage by partitioning the graph into subgraphs with minimal edges between them [Simon & Teng, 1997]. This requires a general overview of the network topology however, an assumption we try to avoid in this thesis. As such, we leave integration of these metrics and approaches as future work, discussed in Section 6.2.3.

2.8 Conclusion

In this chapter we have explored the related work in the literature on convention emergence. We have examined what a convention is, and the benefits behind allowing them to emerge in MAS and how they can facilitate coordinated behaviour amongst agents that will reduce clashes and increase efficiency. We briefly explored the related notions of cooperation and norms, highlighting the differences between them and the aspects that have implications for convention emergence. We discussed the notion of *manipulating* a convention to a desired outcome using Intervention Agents (IAs) both to encourage initial convention emergence and convention destabilisation. We provided an in depth account on the nature of games that can be used to model coordination and convention as well as other games that have been used in the modelling of conventions. Finally, we provided a summary of the network topologies that are to be used in the rest of this thesis and a description of the graph metrics that will be used to place IAs.

CHAPTER 3

Interventions and Destabilisation

In the previous chapter we presented an overview of the current state-of-the-art with regards to studying convention emergence in multi-agent systems (MAS). As discussed, encouraging rapid and robust convention emergence is of great importance as it can minimise the amount of clashes and wasted resources that agents experience due to miscoordination. Additionally, in cases where undesirable conventions have already emerged, it may be beneficial to replace them with others. In this chapter we use a model of convention emergence based on the coordination game and attempt to manipulate both the emergence and replacement of conventions within the system. We show that a relatively small proportion of agents playing a fixed strategy can both elicit convention emergence and remove established conventions. We apply these insights to both static, real-world and dynamic topologies and examine the differences between them.

3.1 Introduction

In MAS coordinated actions help to reduce the costs associated with incompatible choices and increase the efficiency of a system. However, in many domains such behaviour cannot be enforced, as there is no centralised control and a lack of *a priori* knowledge of which actions clash. In practice, many systems rely on the evolution of *conventions* as standards of behaviour adopted by agents with no, or little, involvement from system designers. Understanding how these conventions emerge, how they can be influenced, and how aspects such as topology affect them is an active research area [Delgado et al., 2003; Franks et al., 2014; Kittock, 1995; Sen & Airiau, 2007; Villatoro et al., 2009].

Conventions have been shown to support high levels of coordination without the need to dictate action choices in a top-down manner. Facilitating the emergence of high-quality conventions in a short period of time, without requiring prior computation, is of particular importance. Much work has focused on the emergence of conventions given only agent rationality and the ability to learn from previous choices. Small numbers of fixed strategy agents (agents who choose the same action regardless of others' choices) have been shown to influence the conventions that emerge and to increase the speed of adoption [Franks et al., 2014; Griffiths & Anand, 2012; Sen & Airiau, 2007].

The ability to remove, as well as establish, conventions allows correction or replacement of adopted actions. In domains where the desirability of actions can change over time, being able to cause such a change is beneficial to the system as a whole. Additionally, understanding how to cause this shift gives insights into what makes a convention robust to outside influence.

Additionally, in many domains, the nature of the relationships between agents is not static. Agents may leave the system, new agents can enter, and the links between agents may change over time. These dynamic interaction topologies induce different system characteristics than those found in static networks. Relatively little work has studied the nature of convention emergence in these types of network.

In this chapter, we examine what is needed to elicit fast (conventions should emerge faster than it would take to dictate action choice to each agent in turn, i.e. less timesteps than the number of agents) and high-quality (once established conventions should be stable unless interfered with) convention emergence. Additionally, we investigate what allows the *destabilisation* of an established convention. We propose temporarily inserting fixed strategy agents, known as *Intervention Agents* (IA), to facilitate these effects. For *initial intervention*, where we wish to use these IAs to encourage conventions to emerge within the system, we desire the strategy assigned to these IAs to be the one that emerges as dominant. In *late intervention*, where we wish to destabilise an already established convention, we provide the IAs with strategies that differ from this established convention to influence a population into discarding it. The insertion of IAs is equivalent to incentivising individuals to take particular actions, for example through continued reward or payment. We show that a small proportion of IAs placed at targeted locations in the population are able to encourage rapid emergence of stable conventions and can destabilise established conventions, replacing it with another of our choosing. We also show that conventions can be destabilised in such a way that we are not required to select a replacement, and instead place the system into a state that allows a new convention to emerge. We study these effects in both static and dynamic topologies. For convention manipulation in dynamic topologies we introduce a new heuristic, LIFE-DEGREE, to support this investigation, which considers features unique to the dynamic nature of the system when placing fixed strategy agents. We examine the importance of dynamic topology characteristics by comparing the performance of LIFE-DEGREE against previously used heuristics based on network metrics.

The remainder of this chapter is organised as follows. In Section 3.2 we provide a brief review of the more salient parts of the literature that apply to this chapter. Section 3.3 introduces the model of interactions that we use to simulate the emergence of conventions. In Section 3.4 we use this model to study the effect of initial intervention in encouraging convention emergence in static topologies, both synthetic and real-world. Section 3.5 does the same for dynamic topologies and explores the nature of unaided convention emergence within these types of networks. We then look at the effectiveness of IAs in destabilising existing conventions, what we term *late intervention*, for both static and dynamic topologies in Sections 3.6 and 3.7 respectively. Section 3.8 introduces the notion of *Passive Destabilisation*, removing an existing convention without necessarily replacing it, and explores this in both static and dynamic networks. In Section 3.9 we extend previous work in the literature on how varying the size of the convention space can affect convention emergence both for initial and late intervention. Section 3.10 investigates the way in which varying the payoffs received by agents can change the nature of manipulating conventions, and we briefly explore convention emergence in the gradient coordination game in Section 3.11. Finally, in Section 3.12 we present our conclusions.

3.2 Background

Whilst the field of convention emergence under numerous models is relatively well-established, there has been little work on manipulating conventions and trying to direct their emergence to one of the several equally beneficial outcomes or equilibria. Our hypothesis is that fixed strategy agents, or IAs, can be used to direct conventions upon their initial emergence as well as being used to destabilise and replace already established conventions.

Airiau et al. [2014] and Sen & Airiau [2007] have shown that fixed strategy agents are able to direct convention emergence to a specific outcome in a population much larger than themselves. They utilise a form of "social learning" where agents use the payoffs they receive from playing a coordination game amongst themselves to inform their future decisions. However, their model is limited due to a small convention space (agents only able to choose between two differing actions) and the lack of an underlying topology (agents are able to interact with all others in the population). The lack of an underlying topology is particularly restricting as this has been shown to have a large effect on convention emergence and in real-world networks individuals are unlikely to be able to interact with all others in the population.

Delgado [2002] is one prominent example of how the topology that agents are situated within can have a large effect on convention emergence. He also utilises the coordination game and shows that complex networks, such as scale-free or small-world, emerge conventions substantially faster than more simplistic networks such as regular graphs and ring networks. His findings indicate that small-world networks are also slower to converge than scale-free, likely due to the underlying degree distributions. Delgado et al. [2003] expands this work and also shows that increases in the population size have the effect of increasing the amount of time taken for convergence in a sub-linear fashion. They argue that the *diameter* of the network (the longest shortest path between any two nodes) is the main driving factor in convention emergence time and this grows logarithmically for many network types.

Griffiths & Anand [2012] build on the work of Sen & Airiau [2007] by including an underlying network topology that limits agent interactions. They show that *where* the fixed strategy agents are placed within this topology can substantially increase their efficiency at bringing about rapid convention emergence, particularly in scale-free topologies. Their findings again indicate a difference between small-world and scale-free graphs, but they only consider a population of between 100 and 1000 agents.

Franks et al. [2013] explore the notion of fixed strategy agents in the language coordination game in both small-world and scale-free networks. They show that increases to the graph density (the number of edges compared to the number of nodes) hastens convention emergence in scale-free networks in this domain and that changing the clustering exponent of small-world networks has no similar comparable effect. However, again, they only consider 1000 agent populations and only up to 5000 edges in their comparison for scale-free networks. They highlight that placing fixed strategy agents by degree rather than randomly substantially increases their effectiveness.

Few studies have considered the additional effect of dynamism on network topologies and the effect it may have on convention emergence. Savarimuthu et al. [2007] have considered the problem for the related work of norm emergence but make underlying assumptions about agent architecture that limit its general applicability, although their results are among the first to show that emergence is possible in dynamic topologies. Franks [2013] briefly considers a simple model where otherwise static networks have small levels of population churn but this does not include changing edges in the network, only the agents themselves and as such is a poor approximation of many real-world dynamic networks. Their model also requires underlying changes to agent capabilities to allow them to incentivise or sanction others, whilst we assume that agents are heterogeneous in their abilities. Mungovan et al. [2010] examine norm emergence in a primarily static network but allow each agent the ability to randomly interact with distant neighbours each timestep, adding a dynamic element to the simulation. Their results show that increasing the dynamic nature increases the rate of norm emergence indicating that the unique interactions available help to increase the spread of the norm between otherwise distant components of the network.

Overall, the current work in the literature indicates that topological artefacts can have a substantial effect on convention emergence and we expect the inclusion of more accurate dynamic networks to reflect this as well.

3.3 Interaction Model & Experimental Setup

Conventions emerge as a result of agents in a population selecting the same action and learning the best strategy (action choice) over time. We assume that a population consists of a set of agents, $Ag = \{1, ..., N\}$, who select from a number of actions, $\Sigma = \{\sigma_1, \sigma_2, ..., \sigma_n\}$. Each timestep each agent selects an interaction partner at random, and both partners choose an action from Σ. The individual payoff for each agent is determined by the combination of action choices, the *joint action*. We adopt the n-action coordination game, such that interaction partners receive a positive payoff if they select the same action and a negative payoff if their actions differ. The 2-action coordination game is often used in exploring convention emergence, but we expand to the n-action coordination game to avoid restricting the number of possible conventions as discussed above. We otherwise utilise the payoff matrix of Sen & Airiau [2007] such that choosing the same action gives a positive payoff ($+4$) and choosing differing actions results in a negative payoff (-1). Sen et al. showed that these values were able to facilitate rapid convention emergence and as such are well-suited for our exploration of other factors that might effect the emergence. We explore different values for these in Section 3.10 to examine the effect that this asymmetry might contribute.

Each agent chooses the action that it believes will result in the highest payoff based on its previous interactions. It does this by making use of a simplified version of the Q-Learning algorithm [Watkins, 1989]. For each action $\sigma \in \Sigma$ an agent maintains an estimate of the payoff it expects to receive from choosing that action in the future (a "Q-value"). The agents update the relevant value after receiving a payoff for choosing an action, σ, in an interaction such that:

$$Q(\sigma) = (1 - \alpha) \times Q(\sigma) + \alpha \times payoff \tag{3.1}$$

where α is a variable in the range $[0, 1]$ that controls the learning rate. For all agents we start with $Q(\sigma) = 0 \quad \forall \sigma \in \Sigma$ so as not to bias any agent towards specific action.

We also assume an element of exploration, such that with probability $p_{explore}$ agents will choose a random action from those available instead of the action they believe to be optimal. This allows agents to avoid local optima in the convention space and facilitates the emergence of global convention. If agents have multiple highest Q-values they will choose randomly between them as well. In this regard our model adopts the approach of Villatoro et al. [2009] by using this Q-Learning algorithm for both partners in an interaction to update their strategies. Airiau et al. [2014] show that populations of entirely Q-Learners emerge conventions faster than the related strategy of "Win or Learn Fast" policy hill-climbing (WoLF-PHC) or of mixed learners and so we adopt this learning method globally.

We assume that agents are situated on a topology that restricts their interactions such that agents can only interact with their neighbours (and hence select randomly from amongst these). The particular topologies used are discussed in each relevant section.

We establish the membership of each convention by querying each agent every timestep on what action they would choose if they were not exploring. The strategy they respond with will thus be the one with the highest Q-value for them or chosen randomly from amongst equal Q-values.

3.3.1 Intervention Agents

As discussed, fixed strategy agents, which we refer to as *Intervention Agents (IAs)*, have been shown to influence convention emergence when introduced at the beginning of a simulation. Building on the work of Franks et al. [2013] and Griffiths & Anand [2012] we propose inserting these IAs at locations within the topology to affect convention emergence.

We generally seek to place these IAs at topologically influential locations as determined by a number of graph metrics. This has been shown to increase their efficacy with placement at both high-degree locations [Franks et al., 2013] and high-betweenness–centrality locations [Griffiths & Anand, 2012] performing better than random placement. Franks et al. [2014] additionally show that placement by eigenvector centrality (EC), highest edge embeddedness (HEE) and hyperlink-induced topic search (HITS) increase efficacy but only consider the case of a single IA being positioned by these metrics.

As such, we generally utilise the following 4 metrics to place IAs in this chapter: degree, eigenvector centrality, highest edge embeddedness and HITS. We also consider random placement as a baseline where appropriate. These metrics are discussed in detail in Section 2.7 and have been shown to good indicators of agent influence [Franks et al., 2014].

We choose to exclude both betweenness centrality (BC) and closeness centrality (CC) for two main reasons: (i) Franks et al. [2014] and Griffiths & Anand [2012] have shown that they offer little if any improvement over the metrics chosen and (ii) they are substantially more computationally expensive than the other metrics, being ill-suited for larger topologies or topologies where the metrics must be recalculated often, as is the case for dynamic topologies. This view has been previously raised in the literature. Kang et al. [2011] argue that these metrics were defined when large-scale networks (such as social networks or the Internet) were uncommon and that they are inherently inappropriate for use on large graphs as they do not consider scalability and are not amenable to parallelisation. Pfeffer & Carley [2012] agrees, stating that both betweenness centrality and closeness centrality are limited in applicability. They introduce approximations based on bounded-distance shortest path calculations but these still represent computational complexity that makes them infeasible for our use. Additionally, both Lawyer [2015] and Šikić et al. [2013] argue that these centrality measures are only good indicators of influence for generally central nodes in the network and severely underestimate the influence of more peripheral nodes. Due to these limitations we exclude both centrality measures from our investigations.

The IAs will always choose to play their assigned strategy and have no ability to explore or deviate from this. They do however, continue to learn, via the same Q-Learning mechanism as all other agents, whilst being used as IAs. As they are unable to explore, they will only learn the value of their assigned strategy but this allows them to establish how well that strategy is performing as a convention. As mentioned in Section 2.3, we query these agents as we would any other, meaning that what convention membership they hold may not be that of the fixed strategy assigned to them if their assigned strategy is not performing well.

This interaction model is the general one used throughout this chapter and indeed throughout this thesis. Any deviations from it will be described in the relevant section. It is our view that this model is well-understood from previous work in the literature and will be particularly applicable to the exploration of destabilisation as it facilitates rapid and robust convention emergence and thus our experimentation can instead focus on the other aspects that affect destabilisation.

3.4 Initial Intervention in Static Networks

We begin by examining the use of IAs in effecting convention emergence within populations where one has yet to emerge, what we define as *initial intervention*. We start by exploring this in synthetic networks in Section 3.4.1 before switching to examine the nature of convention emergence in real-world networks in Section 3.4.2.

3.4.1 Synthetic Networks

We start by investigating convention emergence in our model upon both scale-free and small-world networks. We generate these using the Barabási-Albert and Kleinberg models respectively, as detailed in Section 2.6 using the Java Universal Network/Graph (JUNG) framework. Our initial study is focused on the effect that the network generation parameters have on convention emergence in our model, without the inclusion of IAs. Whilst the topological effects, and the differences they cause, have been studied before these have either been over smaller ranges of parameters [Delgado, 2002; Delgado et al., 2003] or have involved a different model of convention emergence (Delgado [2002] and Delgado et al. [2003] do not utilise social learning in their model and Franks et al. [2013] make use of the language coordination game, which is substantially different). Given that our later work will rely on the rapid emergence of conventions that we can target for destabilisation, understanding the performance of the model without intervention is important.

We initially look at the effect that the population size has on the speed of convention emergence without IAs. We generate graphs with agent populations of $\{1000, 2000, 5000, 10000, 20000\}$, for both scale-free and small-world topologies. Scale-free graphs are generated using the Barabási-Albert model with $m_0 = m = 3$ (m_0 is the initial number of nodes, m the number of edges to attach to each new node) whilst the small-world graphs are generated using Kleinberg's model with $ce = 2$, $l = 1$ (ce is the clustering exponent which in-

Symbol	Meaning	Default Value
m_0	Initial number of nodes in Barabási-Albert model.	3
m	Number of edges to attach each round in Barabási-Albert	3
ce	Clustering exponent in Kleinberg model	2
l	Number of long-range connections in Kleinberg model	1
α	Learning rate in Q-Learning	0.25
$p_{explore}$	Exploration rate used alongside Q-Learning	0.25
payoff	Payoff values for coordination game	+4, -1
-	Number of runs results are averaged over	100
-	The number of actions in the coordination game	10

Table 3.1: Parameters and their default values for synthetic static networks

forms how far long-range connections will reach, l is the number of long range connections). The settings for these models are explained further in Section 2.6.1 and a summary and the default values used are shown in Table 3.1. These produce graphs of comparable size with roughly the same number of edges in both the scale-free and small-world topologies ($|V|/|E| \approx 3$) which allows for easier comparison. The value of ce is chosen based on the initial paper of Kleinberg [2000b] which shows that for this value only there exists a rapid greedy method for information transmission.

To ensure a large enough convention space, agents interact using the 10-action coordination game. With a convention space of this size we would expect our desired strategy, without the presence of IAs to emerge 10% of the time. Simulations were run for either 5000 timesteps (scale-free) or 10000 timesteps (small-world) in order to give sufficient time for conventions to emerge. We use the 90% Kittock criteria such that we deem a convention to have emerged when 90% of agents in the population would choose that action over others.

Figure 3.1 shows the results for both topologies averaged over 100 runs with error bars showing the standard deviation. Consistent with the findings of Delgado et al. [2003], the effect that increasing population size has on convention emergence speed is sub-logarithmic with small-world topologies taking substantially longer to reach convention emergence than scale-free ones. We also find

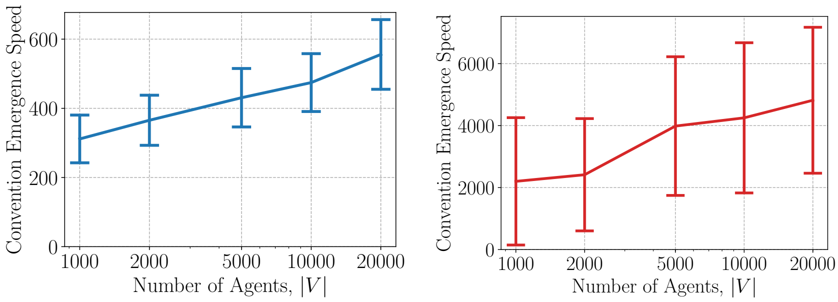

(a) Scale-free

(b) Small-world

Figure 3.1: Speed of unaided convention emergence against varying population sizes.

that the convention emergence speed in small-world topologies is substantially more *varied* than in scale-free much higher standard deviations. This indicates that convention emergence is less consistent in small-world topologies and thus we also expect the inclusion of IAs to increase this as they provide direction within the population. Indeed, whilst 100% of runs on the scale-free topology reached convention emergence (most much faster than the timestep limit) only between 54-92% did for small-world topologies. Allowing the small-world simulations to run for longer increases the number achieving convention emergence somewhat but often still does not guarantee 100%. This is in keeping with similar findings of Franks et al. [2013] and indicates that small-world topologies are more likely to have multiple semi-stable conventions amongst the population. This is most likely due to the lack of "hub" nodes that are present in scale-free topologies and are able to influence large sections of the population. The more localised nature of small-world topologies means they are more robust to external invasion of competing conventions.

We can verify this by changing the clustering exponent for small-world topologies. Using a population of size 5000 and $l = 1$ we vary ce with values of $\{0, 1, 2, 3, 5\}$. We again allow 10000 timesteps for populations to converge and otherwise use the same setup as before. Figure 3.2 shows how the convention membership sizes of the largest convention in the small-world topologies changes

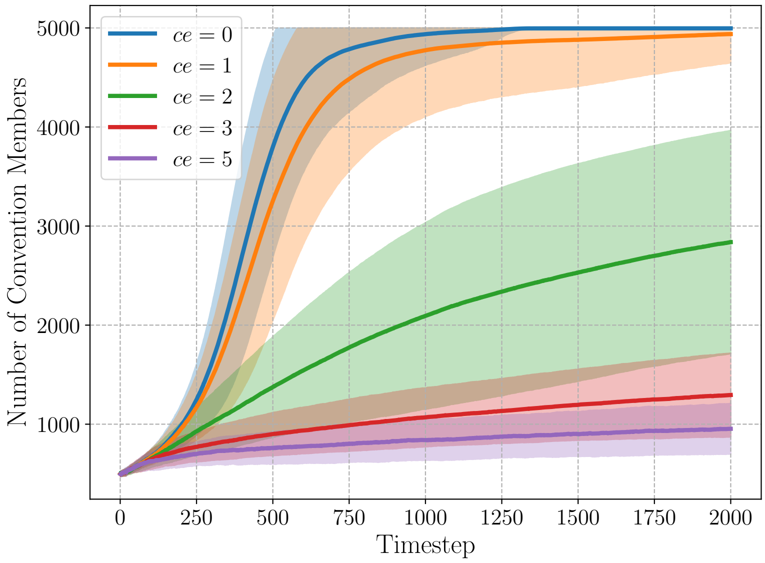

Figure 3.2: Convention membership sizes over time for the largest convention for varying clustering exponent in the small-world network. The shaded areas represent the standard deviation over 100 runs.

during the simulation over the first 2000 timesteps. The values are found by averaging the size of the largest convention each timestep over 100 runs and the shaded areas show the standard deviations. As can be seen, the clustering exponent has a significant and substantial effect on the way conventions emerge within the system with lower clustering exponents facilitating much faster convention emergence than higher values. This contrasts with the findings of Franks et al. [2013] who noted that the clustering exponent had no effect on convention emergence in the language coordination game. The substantial differences here are thus likely due to the different coordination models being used and highlight that different models are susceptible to different aspects of the underlying topologies. Additionally, the assumptions of Franks et al. include a smaller network size (only 1000 agents) and a differing measure of convention emergence which is applicable only to the language coordination game. The manner in which clustering exponent affects convention emergence is consistent with the hypothesis of Delgado [2002], that the network diameter is the major factor in allowing rapid and consistent convention emergence. Decreasing ce makes it more likely that the long-range connections will join otherwise distant parts of the topology, decreasing the shortest paths between all nodes and allowing conventions to spread more easily amongst the network. Conversely, increasing ce limits the effective radius on the lattice that long-range connections will be made to and so increases the clustered nature of the topology, making it consist of localised areas that are resistant to outside change. The difference between $ce = 1$ and $ce = 2$ is the most significant however, with the former causing substantially faster convention emergence and to a higher level of consistency, with 100% of the runs achieving convention emergence within 10000 timesteps. The difference between $ce = 1$ and $ce = 2$ is much more so than between either $ce = 1$ and $ce = 0$ or $ce = 2$ and $ce = 3$ indicating that there is a prominent shift between these values, likely accounted for by the unique nature of $ce = 2$ as espoused by Kleinberg [2000b]. However, none of these values invalidate the small-world feature of having a richly connected local area with a small diameter.

Similarly, increasing the number of edges will, in general, decrease the diameter of the network and so we would expect this to have a similar effect on convention emergence. We can readily control the number of edges in both scale-free and small-world topologies and so we vary $m_0 = m = \{3, 5, 7, 9\}$ and $l = \{1, 3, 5, 7\}$ to produce values of $|V|/|E| \approx \{3, 5, 7, 9\}$ in each topology. In the small-world network, we leave $ce = 2$. As before, we use a population of 5000 agents and 5000 or 10000 timesteps for scale-free and small-world topologies respectively. We perform 100 runs for each setting and find the average performance of the largest convention over time. Figures 3.3 and 3.4 show the results of this for the small-world and scale-free graphs respectively. As can be seen in Figure 3.3, increasing the number of edges has a dramatic effect on the nature of convention emergence in small-world topologies. This can primarily be attributed to the extra edges necessarily taking the form of additional long-range connections as in the Kleinberg model the edges in the lattice are already

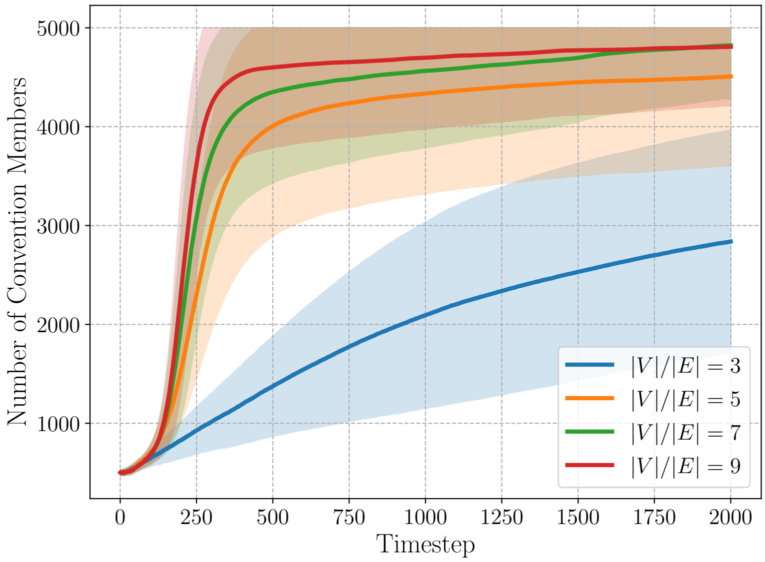

Figure 3.3: Convention membership sizes over time for the largest convention for different values of $|V|/|E|$ in the small-world network. The shaded areas represent the standard deviation over 100 runs.

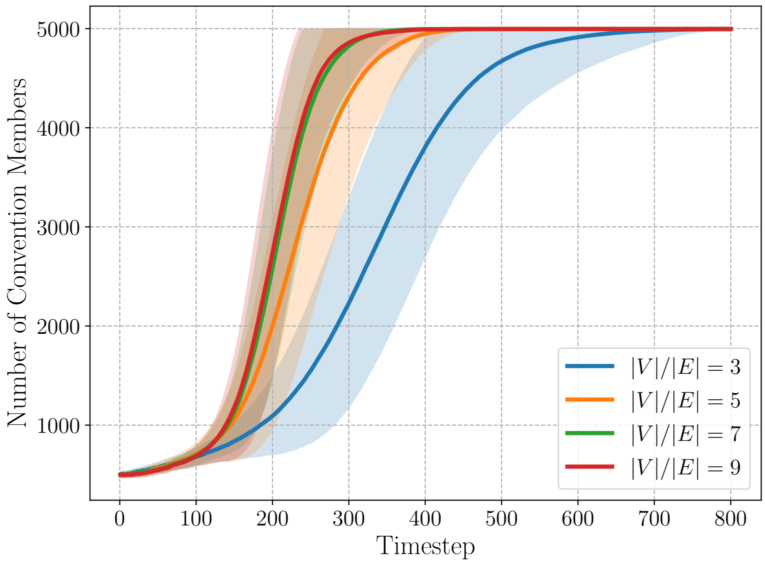

Figure 3.4: Convention membership sizes over time for the largest convention for different values of $|V|/|E|$ in the scale-free network. The shaded areas represent the standard deviation over 100 runs.

present. However, it again highlights that increasing the connectedness of the local areas can dramatically affect the rate of convention emergence, with even the increase from $|V|/|E| = 3$ to $|V|/|E| = 5$ allowing conventions to emerge to the 80% level almost as rapidly as in scale-free topologies. The increase from the 80% Kittock level to the 90% Kittock level is still noticeably slower than in scale-free networks however, with wider ranging behaviour prevailing in the small-world network (shown by the standard deviations in the shaded areas). In contrast, Figure 3.4 shows that whilst increasing the value of $|V|/|E|$ does increase the convention emergence speed in scale-free networks the difference is much less marked than in small-world topologies. Additionally, there is negligible difference between the performance at $|V|/|E| = 7$ and $|V|/|E| = 9$ indicating that there are diminishing returns of increasing the number of edges. For the scale-free networks these findings are corroborated by Franks et al. [2013] who found a similar relationship for scale-free networks in the language coordination game. The effect in small-world networks extends this and shows that, for the coordination game at least, better connectedness improves convention emergence in both topologies.

Introducing IAs

Having established that conventions can emerge naturally within our model we now focus our attention on the role and functionality that IAs can play in effecting convention emergence to a desired convention.

To investigate this we introduce a number of IAs into the topologies at time $t = 0$. The IAs are placed at influential locations as determined by the metrics discussed in Section 3.3.1 (degree, EC, HEE and HITS) and are placed at the highest value locations available according to each metric. As we have no particular preference over which strategy emerges as a convention, the strategy is chosen uniformly at random from amongst the 10 available and assigned to each of the IAs so that they all play the same fixed strategy. We explore the case where IAs are assigned different strategies in Section 3.8. We utilise the

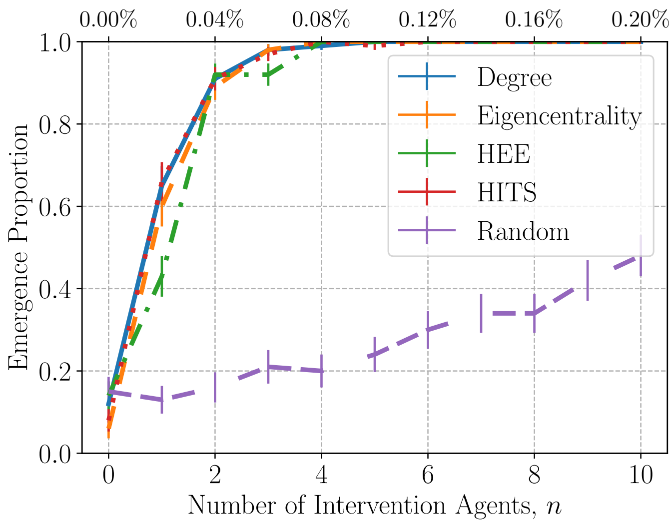

Figure 3.5: Proportion of runs emerging the target convention for varying numbers of IAs in the 5000 node scale-free network.

90% Kittock criteria as before so that a convention has emerged when 90% of agents, when queried, would all choose the same action.

We create scale-free networks with $m_0 = m = 3$ and small-world networks with $ce = 2$, $l = 1$, as before. Both, initially, consist of 5000 agents. The simulations run for 5000 timesteps in the scale-free networks and 10000 timesteps in the small-world networks. As previously discussed, this may prematurely stop some small-world simulations that would otherwise have emerged conventions but we hold that the varied nature of convention emergence speed in small-world networks is a facet that we hope the IAs will help mitigate and speed up. Additionally, as the average speed of unaided convention emergence is well below this limit, we believe it will only affect some outliers and, as it is applied consistently across all runs, does not affect the conclusions we can draw about the relative performance of the metrics. Each setting was performed over 100 runs.

Figure 3.5 shows the proportion of runs that emerged the target convention (that of the IAs) for increasing numbers of IAs in the scale-free network. As we are concerned with the *proportion* of runs that emerge to the desired convention there is no averaging occurring and no standard deviations. The error bars thus represent the standard error in proportions, $SE_p = \sqrt{p(1-p)/n}$, to show the level of uncertainty in the results.

The immediate thing to notice is that, apart from random placement, all of the metrics perform nearly identically, with very little to differentiate them. Indeed, the only point at which there is a statistically significant difference between degree and another targeted metric is for the performance of HEE at 1 IA (two-proportion z-test, $p < 0.05$). We use degree as the baseline comparison due to its previously found benefits [Franks et al., 2013; Griffiths & Anand, 2012] and so we are primarily concerned with relative behaviour to it. Additionally, the number of IAs needed to cause full convention emergence, where 100% of runs converged to the desired convention is 4-5 in each of the targeted metrics. This represents a small fraction of the agent population (shown at the top of the chart) and shows that small numbers of targeted IAs can affect much larger populations as was found by Griffiths & Anand [2012] and Sen & Airiau [2007]. The similar performance between each of the metrics is likely due to the strong correlation between the metrics [Franks et al., 2014] and means they are likely selecting nearly the same type of nodes.

Figure 3.6 shows the same type of results for the 5000 node small-world topology and has a number of distinct characteristics compared to the scale-free results. Firstly the number of IAs required to elicit the same level of desired convention emergence is substantially higher in the small-world topology with 100% emergence not occurring until 30-40 IAs. Whilst this is still a small percentage of the overall network (<1%) it indicates that the small-world topologies are more resilient to the effects of IAs, an aspect that can be attributed to the resilient local clusters discussed previously. The effectiveness of the HEE metric is also lesser here with it performing significantly worse than the degree metric from 10-26 and 30 IAs ($p < 0.05$). HITS also performs worse though it is only significant at 16, 20 and 24 IAs. Most importantly, random placement performs comparably to all other metrics (worse than degree at 12-20 and 26, $p < 0.05$),

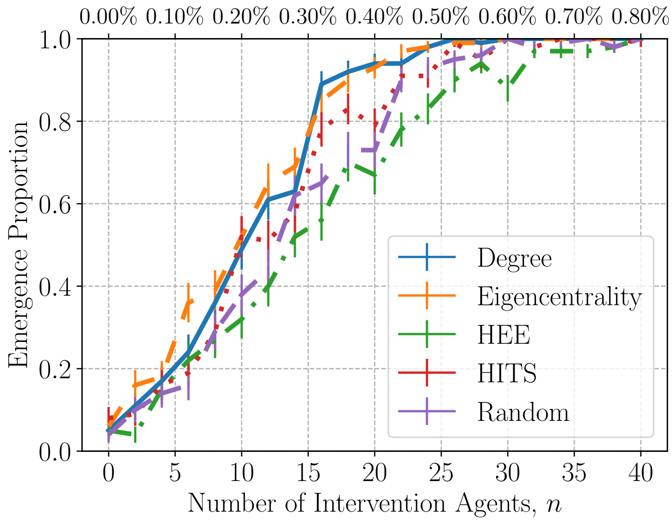

Figure 3.6: Proportion of runs emerging the target convention for varying numbers of IAs in the 5000 node small-world network.

a stark departure from its performance in scale-free topologies. This is likely due to the nature of small-world topologies, with them lacking the power-law degree distribution that gives the "hub" nodes of the scale-free topology their influencing power. Instead this indicates that as long as resilient clusters are converted, the level of influence the node exhibits according to these metrics is less important in small-world topologies at encouraging initial emergence of conventions.

Another commonly used metric by which to evaluate convention emergence is the speed with which the convention emerges. We consider this also for the use of IAs to study not just how the convention emerges but how quickly. Figure 3.7 shows the effect of IAs on the speed of convention emergence in the 5000 node scale-free network discussed previously. The time for the convention to emerge, regardless of whether it was the desired one or not, was averaged over the 100 runs and is presented here. As can be seen, the inclusion of IAs dramatically reduces the time taken for a convention to emerge for all targeted placement metrics. Whilst HEE still performs worse than the others the difference is marginal and overall shows that even a small number of IAs can help to guide

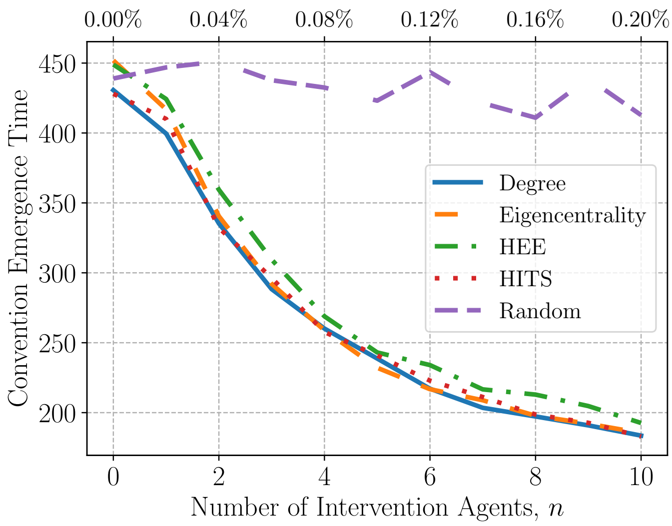

Figure 3.7: Convention emergence time against the number of IAs in the 5000 node scale-free network.

rapid convention emergence. The relationship is one of diminishing returns however, with larger numbers of IAs increasing the speed less and less.

Figure 3.8 shows the same for the small-world topology. It is even clearer here that HEE performs markedly worse than the other targeted placement metrics but the effect is otherwise the same; IAs reduce the time for the convention to emerge. These findings corroborate those of Griffiths & Anand [2012] but the fact that the additional metrics we study have the same effect is of interest.

Having previously shown that the population size, number of edges and, in the case of the small-world network, clustering exponent can have a substantial effect on convention emergence in the systems, Figures 3.9 and 3.10 show how varying these parameters affects convention emergence when utilising IAs. For the sake of clarity, these graphs only show the performance of degree placement (that is, placement by the degree metric) under each setting, as it performs as well or better than all other placement metrics. Figure 3.9 shows the results for scale-free networks with varying population sizes and densities. We note that there is very little difference here in the performance of IAs in eliciting the desired convention emergence with ~4-5 IAs still causing 100% of runs to converge

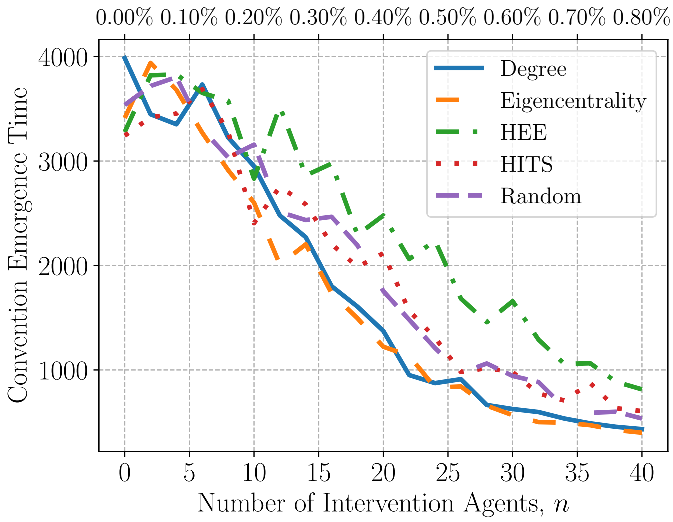

Figure 3.8: Convention emergence time against the number of IAs in the 5000 node scale-free network.

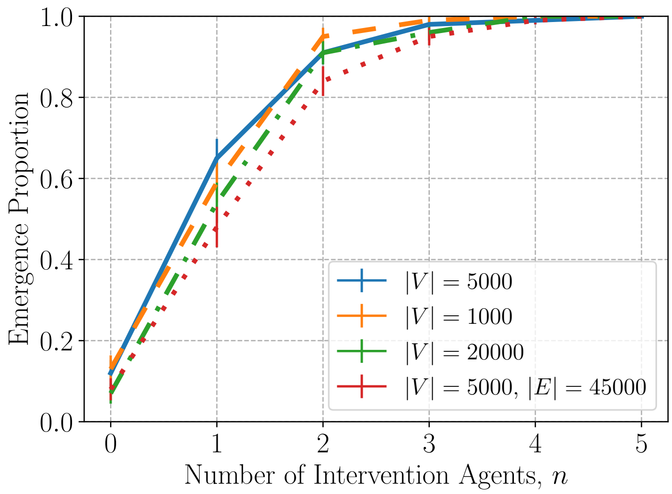

Figure 3.9: Proportion of runs emerging the target convention for varying numbers of IAs using degree placement in scale-free networks with different parameters.

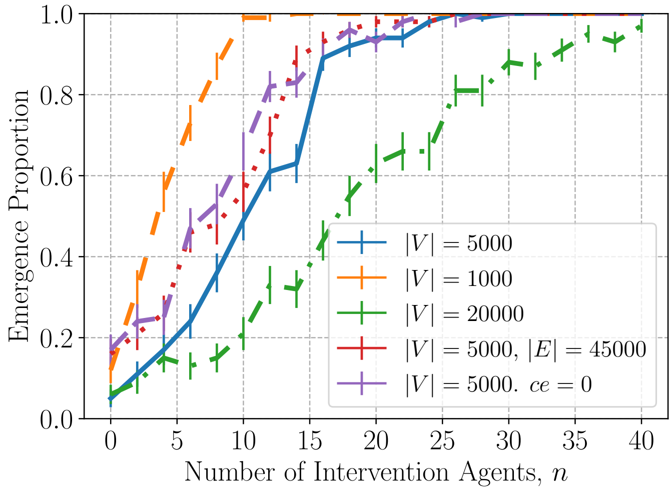

Figure 3.10: Proportion of runs emerging the target convention for varying numbers of IAs using degree placement in small-world networks with different parameters.

to the desired convention, regardless of network size and density. Including just a few IAs at the correct points in scale-free networks seems to allow consistent and robust convention emergence in a range of settings.

There is a more marked difference in the performance of degree placement in the various types of small-world topologies, as shown in Figure 3.10. In these topologies, increasing the number of agents changes the number of IAs required to cause the same amount of change, with larger populations requiring more IAs. However, even at 20000 agents, almost 100% convention emergence is achieved using only 40 IAs, despite them representing a much smaller proportion of the population than at 1000 or 5000 agents. This, coupled with the results for scale-free networks, indicates that the number of IAs required to effect the desired convention emergence is mostly independent of network size and highlights the effectiveness of this approach. Varying the number of edges or the clustering exponent has a much lesser effect than population size, despite the manner in which they affected unaided convention emergence. Whilst they do change the number of IAs required, the change is minimal compared to the base case of

5000 agents with standard settings.

We have shown that IAs can be highly effective at producing the desired convention to emerge rapidly over a range of graph settings and population sizes in synthetic networks, expanding the previous scales at which these had been examined and utilising a number of previously mostly unexplored placement metrics.

3.4.2 Real-World Networks

We now turn our attention to using these same metrics in real-world topologies. Real-world networks have been shown to have many different characteristics to those generated synthetically [Franks, 2013; Pujol et al., 2005] but have been mostly ignored in the study of convention emergence, particularly for the use of IAs.

Franks [2013] uses a number of real-world topologies in his work but utilises sampling of the networks, reducing them to much smaller sizes (around 1000 nodes). As we have seen, population size can have a noticeable effect on convention emergence and so reducing these networks removes some of the benefit of studying them. Additionally, as noted by Franks and Gjoka et al. [2011], there are a number of issues with sampling networks. Gjoka et al. note that a number of the potential sampling methods are heavily biased towards high-degree nodes, creating samples that are unrepresentative. Whilst they explore others, such as Metropolis Hastings Random Walk, which perform better and are able to replicate a number of macro-scale features of the sampled topologies, they note that there are none that replicate local structure whilst also replicating these macro-scale features. Local structure has been shown to be important when it comes to convention emergence [Villatoro et al., 2011a; Villatoro et al., 2009] as well as the macro-scale features and losing either of this reduces the generalisability of the results.

Thus, in this thesis, as Villatoro et al. do, we focus on the entirety of the real-world networks, only reducing them to their largest weakly-connected component (WCC) for ease of allowing global convention emergence. This allows us to benefit from both the macro-scale and local features of the topologies, as well as providing an insight into population sizes beyond even those examined in Section 3.4.1. We utilise 3 real-world networks as discussed in Section 2.6: CondMat, Enron and Twitter.

To explore convention emergence within them we perform similar simulations as before, locating a population of learning agents within the topologies. The settings are as discussed in Section 3.3 although we allow the simulations to run for 30000 timesteps due to the larger sizes of the topologies. We vary the number of IAs introduced and place them as before, measuring their effect over 100 runs. Initial simulations showed that emergence of conventions to the 90% Kittock criteria was unlikely, even over longer time periods and so we change to the 80% Kittock criteria which still represents a substantial proportion of the network adhering to a single strategy.

Figure 3.11 shows the results of these simulations for the 3 real-world topologies. In each of the topologies the percentage of the population that must be IAs to cause full emergence to the desired convention is smaller even than in the 5000 node scale-free topology examined before. In these topologies, 0.02%, 0.06% and 0.06% are sufficient in the CondMat, Enron and Twitter networks respectively to cause full emergence to the desired convention. This lends further credence to our hypothesis that the percentage of IAs required is nearly independent of the population size, even in these larger and more complex networks. The more noticeable difference between these and the synthetic networks is the relative performance of each of the placement metrics. Whilst the 4 targeted metrics are comparable in the CondMat network (with both HEE and HITS actually being statistically significantly better than degree or eigencentrality for 1 IA ($p < 0.05$)), both Enron and Twitter exhibit significantly worse performance for these two metrics with HEE and HITS not causing a single run to change to the desired convention in the Twitter network over the same range where degree and eigencentrality cause 100% of runs to do so. In the Enron network

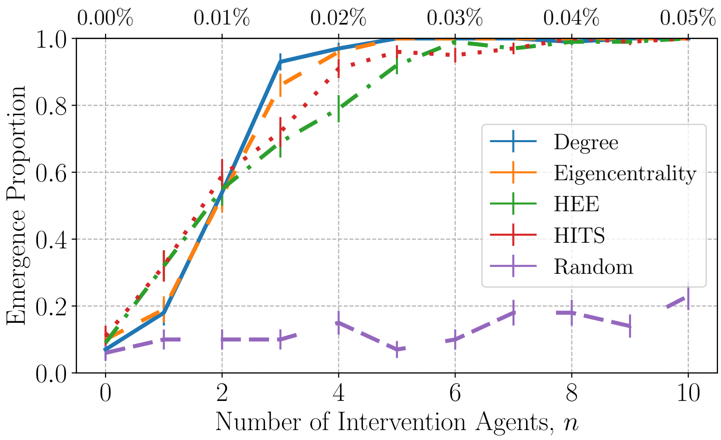

(a) CondMat

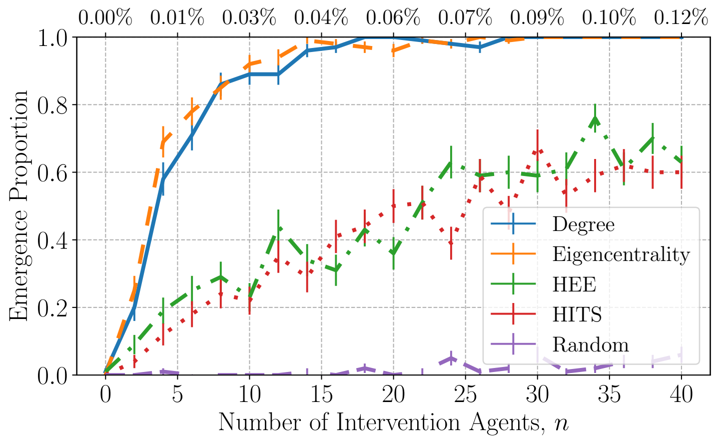

(b) Enron

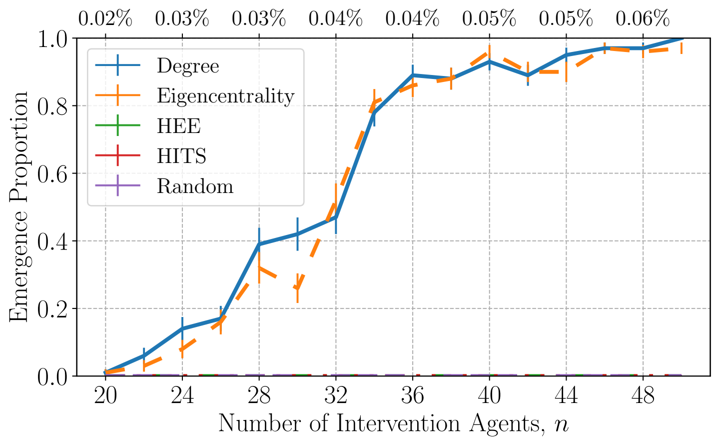

(c) Twitter

Figure 3.11: Proportion of runs emerging the IA convention in real-world networks.

these metrics are able to elicit convention emergence but at a significantly worse performance than both degree and eigencentrality. In all topologies, degree and eigencentrality continue to perform effectively indistinguishably from one another and we believe this to be due to them selecting mostly the same nodes given the correlation between them found by Franks et al. [2014]. These results further our belief that degree placement is the most effective at allowing IAs to influence convention emergence and informs later decisions on placement and evaluation.

3.5 Initial Intervention in Dynamic Networks

Having shown that small proportions of IAs can affect convention emergence in much larger populations for both synthetic and real static networks, we now turn our attention to their effectiveness within dynamic topologies. These have been shown to have different system dynamics than static topologies [Brandt & Sigmund, 2005; Franks, 2013; Savarimuthu et al., 2007] and as such it is unclear how effective convention emergence might occur within these topologies. We seek to explore that notion in this section for initial interventions.

3.5.1 Interaction Model within Dynamic Networks

In the formulation proposed by Kittock [Kittock, 1995], a convention is considered to have emerged when a high proportion (90%) of the population would all choose the same action given an open choice. Due to the dynamic nature of the topologies in this section, whilst we adopt this definition of a convention, as we have done previously, we must modify it to better fit within the dynamism of the network topologies. Instead of considering the entire population, we monitor adoption within the largest connected component. Whilst this could in theory result in consideration of multiple, much smaller, populations than intended, depending on dynamic network topology, our initial simulations show this to not be a concern in the two models we utilise: González and Ichinose as described

Symbol	Meaning	Default Value
m_0	Initial number of nodes in Ichinose model	3
m	Maximum number of edges to attach to new nodes each round in Ichinose, rest rewired. During initial creation, number of edges to attach to each new node.	3
r	Radius of particles in González model. Set as part of T_l/T_0.	0.01
v_0	Initial speed of particles in González model. Set as part of T_l/T_0.	0.3
$\bar{v}$	Speed boost from collision of particles in González model.	0.3
T_l	Maximum life of agents in González model. Set as part of T_l/T_0.	500
L	Arena size in González in model. Default calculated based on desired $\|V\|$ to ensure density of 0.625.	-
α	Learning rate in Q-Learning	0.25
$p_{explore}$	Exploration rate used alongside Q-Learning	0.25
payoff	Payoff values for coordination game	+4, -1
-	Number of runs results are averaged over	100
-	The number of actions in the coordination game	10

Table 3.2: Parameters and their default values for dynamic networks

in Section 2.6.3 whose default parameter values are shown in Table 3.2.

In the González model we find that in most simulations a giant cluster consisting of nearly all agents will emerge. Agents not within this cluster are likely to be recently created agents and, as such, should not be included in the adoption rate calculation as they have not interacted. This is reinforced by our simulations which showed that most agents not within the largest connected component had degree zero. This follows the findings of both González et al. and Savarimuthu et al. [2007]. However, whilst there is a large connected cluster, the number of agents outside of this is not insubstantial. In order to ensure that the population of this giant cluster is close to what we want to consider, we adjust the size of the González model in all experiments such that the giant cluster itself consists of roughly the right number of agents. Through experimentation we find this adjustment to be ~10% such that a requested population of 1000 agents will be created as a González population of 1100 so that the main cluster is ∼1000 agents at all times.

In the Ichinose model, due to the methods in which nodes and edges are removed (ensuring that both remain at constant size), there is less risk of the main population not being of the size intended. Due to the rewiring, there is a chance that individual nodes may become isolated but we find this to be in the order of 0.01% of the population at worst and so does not have an appreciable effect.

IAs will be placed within the network to study the effect on convention emergence. These agents are selected from the population rather than replaced and as such have all the edges and knowledge they had at selection time, as was the case in the static topologies. Such agents will all be assigned the same fixed strategy (determined at the start of the simulation) and their placement will be done via the relevant metric as discussed below. Unlike in static topologies, we must consider what happens if an IA is removed from the graph. In this case, a new IA will be selected using the same metric and assigned the same fixed strategy as all the others.

The González model is known to require "burn-in" before the giant connected cluster emerges and the network topology settles into what González et al. [2006b] call a "quasi-stationary state" or QSS. This is known to occur by time $t \approx 2T_l$ where T_l is the lifespan of agents in the González model and as such we perform this many steps of graph simulation before introducing agents to the topology. The Ichinose model has no such requirements as it is built initially as a Barabási-Albert graph of sufficient size.

Otherwise, the interaction model is the same as in the static topologies with agents learning via Q-Learning and only able to interact with their neighbours in the graph.

3.5.2 Placement Heuristics

The dynamic nature of the topologies introduces a number of ways to apply the metrics already discussed. As well as only considering the metrics with respect to the largest connected component we must also consider whether the metrics will be *static* or *updating*.

Static versions of the metrics correspond to the equivalent metrics for static networks. At the time of insertion, agents are chosen to be IAs in descending order of the metric in question. This selection is static once chosen, only being modified upon agent expiration as detailed above. This simplistic approach is computationally cheap, a factor of importance in settings where gathering or computing this information is expensive. However, this risks selected agents potentially becoming sub-optimal choices as the simulation progresses. The static nature of this manner of placement means that if another agent acquires a larger value of the metric it will not be selected until one of the current IAs expires. Depending on the model and expected lifespans of the current IAs, this could be a substantial period.

To address this issue we propose another version of each metric: *Updating*. This approach is sensitive to the dynamic nature of the topology and reselects the IAs each timestep, based on highest current metric value. Whilst this offers a solution to the potential sub-optimality of the Static metrics it suffers from two problems. Firstly, the ability to acquire this information each timestep in a timely manner may be infeasible in many domains. Secondly, there is the potential that the IAs will not remain in a given location long enough to influence the local area before being replaced.

The Static and Updating metrics do not fully consider the dynamic network context. Whilst high metric agents are likely to be influential due to their ability to interact with many others or their centrality within the network, additional dimensions may affect their applicability. Agents close to expiring may be less desirable than younger agents as their expected number of interactions before replacement is lower. However, the youngest agents, those that are newly created, cannot be guaranteed to become influential later on. Hence, the age of an agent adds an additional consideration. We propose a new metric, LIFE-DEGREE, that allows exploration of the effect of age in addition to degree on a the efficacy of an IA.

In many settings it may be impossible to *know* an agent's expected lifespan. However, we can estimate an agent's remaining life (or their "youthfulness") in both the González and Ichinose models based on a comparison with the maximum age possible in the networks. In the González model this is a setting of the graph itself and hence is easily calculable. In the Ichinose model, the expected lifespan of an agent depends on its own degree and the mode in which the Ichinose network model is operating. We can however find the *current* maximum age in the network and compare all others to it. Thus, we can calculate the expected remaining time-to-live, E_{rTTL}, for a node n as,

$$E_{rTTL}(n) = 1 - \frac{age(n)}{\max_{n' \in LCC}(age(n'))} \tag{3.2}$$

We can also calculate the normalised degree of a node n within the largest connected component as:

$$deg_{norm}(n) = \frac{deg(n)}{\max_{n' \in LCC} deg(n')} \tag{3.3}$$

The Life-Degree heuristic is then defined as:

$$\text{Life-Degree}(n) = \omega \times deg_{norm}(n) + (1 - \omega) \times E_{rTTL}(n) \tag{3.4}$$

In this, $0 \leq \omega \leq 1$ is a weight, determining the relative contributions of degree and expected TTL.

Life-Degree allows combination of the relevant information, normalised against theoretical maximums, in a manner that allows exploration of the importance of both. Two variations of Life-Degree will be used, Static and Updating, to compare against the versions of the metric discussed above.

3.5.3 Characterising Topology

We initially consider convention emergence without external manipulation in the dynamic topologies. This gives insight into the impact of network dynamics on convention emergence and provides a baseline. Additionally, it allows us to quantify the point at which a stable convention will have emerged for later experiments that focus on destabilisation.

The features of the dynamic topologies can be manipulated by varying the parameters of the network models. Due to its uniqueness, we begin by considering convention emergence in the González model to show that it is feasible. In the González model the combination of parameters is encapsulated in different values of:

$$\frac{T_l}{T_0} = \frac{2\sqrt{2\pi} r N v_0 T_l}{L^2} \tag{3.5}$$

González et al. [2006b] show that the features of the topology thus only depend on the ratio T_l/T_0 and the density, $\rho \equiv N/L^2$. Additionally, they show that the average degree is a non-linear function of T_l/T_0 that depends on the chosen ρ. As such, for all experiments we use a constant $\rho = 0.625$ (e.g. $N = 1000$, $L = 40$) to allow meaningful comparisons of the T_l/T_0 values. This is automatically calculated based on the chosen population size (once adjusted).

Parameter settings were chosen that generated values of T_l/T_0 between 0 and 20. These were rounded to the nearest integer to combine similar T_l/T_0 values, with each bucket containing 10 randomly chosen values/settings. The average time taken, over 30 rounds, for convention emergence to occur was measured on the generated González topologies and the average time over the bucketed values was then calculated. Values which did not result in convention emergence after 10000 timesteps were discounted from the second average as they were unlikely to result in conventions emerging. Only runs with $T_l/T_0 \lesssim 4$ are affected by this. Simulations with a higher T_l/T_0 exhibited convention emergence for all runs. With $T_l/T_0 \lesssim 4$ as much as 60% of the runs for a given simulation did not result in convergence. The transition is notable and is discussed below.

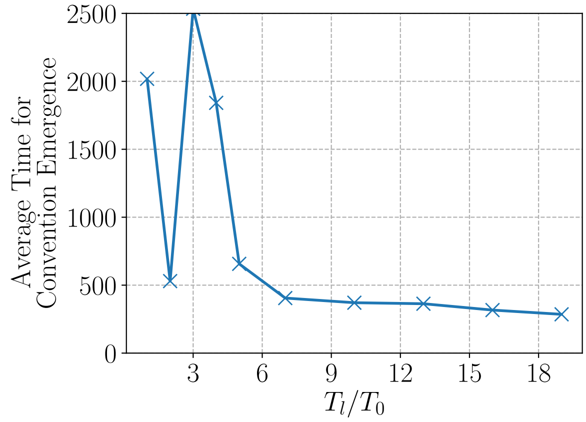

Figure 3.12: Average convention emergence time for different values of T_l/T_0 with no IAs in the González network.

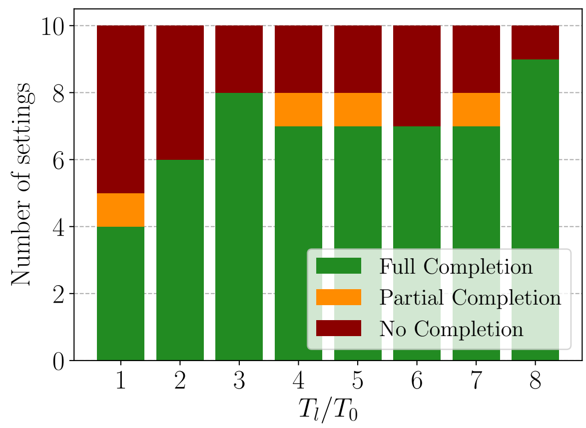

Figure 3.13: Convention emergence performance per setting for different values of T_l/T_0 with no IAs in the González network.

It is clear that convention emergence is successful in the González topology, and for most values of T_l/T_0 there is little variation in the average time for convention emergence as shown in Figure 3.12. Values of $T_l/T_0 \gtrsim 5$ all have a convention emergence time of around $t = 500$ with little variation between runs. However, values of $T_l/T_0 \lesssim 4$ displayed significant variation and, in general, much more time was required for convention emergence to occur if it occurred at all. Higher values of T_l/T_0 did not exhibit this. Figure 3.13 examines the breakdown of convention emergence at the lower values of T_l/T_0 showing the number of settings for each value that resulted in full completion (30 out of 30 runs reaching convention emergence), partial completion (between 1 and 29 runs reaching convention emergence) and no completion (0 runs reaching convention emergence).

At low T_l/T_0 values the topology either did not generate a giant cluster or agents were found to expire before meaningful convention emergence could occur. This follows from the parameter settings required to give a small T_l/T_0 and means that there is a lower threshold for the González topology to experience convention emergence. In particular, there is a minimum level of connectedness and lifespan that must be present. Below this threshold the network will be partially disconnected and not representative of real-world topologies. However, once this is achieved the time required for convention emergence is mostly independent of T_l/T_0. As such, we select parameter settings that are used for all following simulations that give $T_l/T_0 \approx 4.7$ which was found to provide stable convention emergence times. These are: $T_l = 500$, $v_0 = \bar{v} = 0.3$, radius $= 0.01$. The arena size, L, is calculated based on the adjusted number of agents to ensure that $\rho \equiv N/L^2 \approx 0.625$ as discussed above. For completeness, additional T_l/T_0 values in the range 20 to 200 were also examined. There was a slight decrease in the average time at higher values, although the low variation remained. As the real-world networks examined by González et al. [2006b] had T_l/T_0 values around 5-6 these results were purely to determine the impact of high T_l/T_0 values rather than to be used in actual simulations.

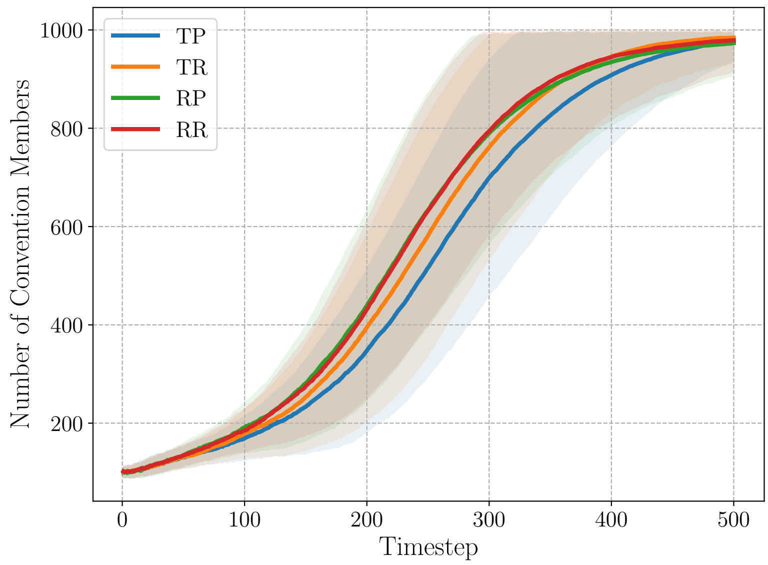

Figure 3.14: Convention emergence in the Ichinose models with no IAs.

For the Ichinose models, given their resemblance to Barabási-Albert scale-free networks, we would expect them to be quite consistent in emerging robust conventions unaided. Figure 3.14 shows this to be the case for all 4 different modes of the Ichinose model generated with $m_0 = m = 3$ and $|V| = 1000$. There is little variation between each of the 4 modes, despite the fundamental differences between them, indicating that the initial starting state of the topology may be beneficial in allowing conventions to emerge. Even with the rapid change that occurs in the TP and TR settings, a convention nearly always emerges by $t = 500$ with TP and TR, on average only slightly behind RR and RP.

3.5.4 Results

Having established that convention emergence occurs in dynamic topologies, we now examine the effect of IAs. We start by considering the scenario where IAs are introduced early in a system's lifespan to manipulate convention emergence: *initial intervention*.

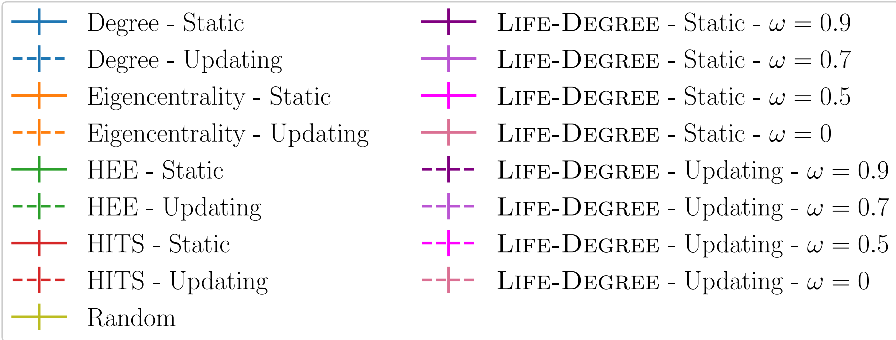

Figure 3.15: Legend for all metrics used in initial intervention in dynamic networks.

Given the large number of metrics being used (Static and Updating versions of all placement metrics except random as well as various weightings of both Static and Updating Life-Degree) we provide a legend that is applicable to all runs and plots for both the González and Ichinose graphs to avoid redundancy in the plots themselves. This is shown in Figure 3.15 and follows the general rule that updating versions of each metric will use dashed lines whilst static versions are solid.

We begin by considering the versions of the traditional metrics discussed in Section 3.5.1: Static and Updating metrics. We also consider random placement of the IAs as a baseline, which is done in a static manner. The IAs were inserted into the system at $t = 0$, after any required burn-in had occurred and the simulation allowed to run for 5000 timesteps. Prior simulations showed that conventions always emerged well before this time even without the presence of IAs. The number of IAs inserted into the system was varied and the proportion of simulations in which the IA strategy emerged as the convention was monitored over 100 runs, as in the static networks.

Figure 3.16 shows the results of this approach for the González model. As can be seen, much as in static networks, a small proportion of the agent population acting as IAs is able to direct the convergence to the desired convention with nearly all metrics, static and updating, doing so by $n = 10$ or 1% of the

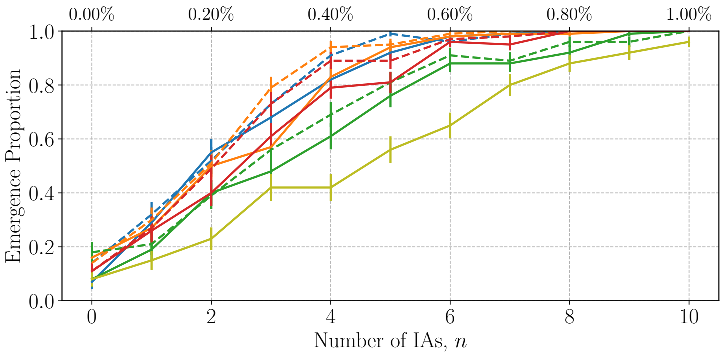

(a) Traditional Metrics

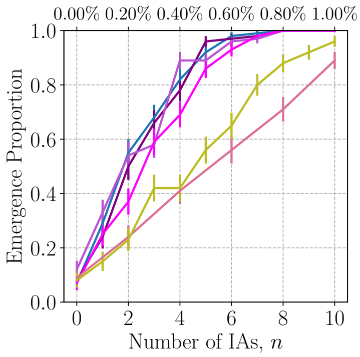

(b) Life-Degree - Static

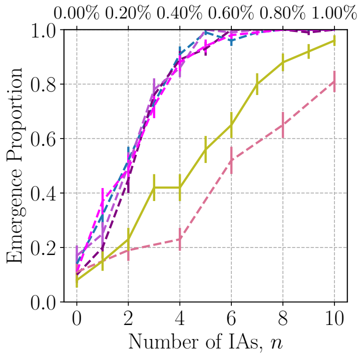

(c) Life-Degree - Updating

Figure 3.16: Proportion of runs emerging the IA convention in the González network.

population. This proportion is however, much larger than that required for the equivalent effect in the static networks, being almost a factor of 10 larger. This is to be expected given the constantly changing nature of dynamic networks. As links are consistently made and lost between distinct portions of the topology we would expect other conventions to find it easier to spread and hence more IAs are needed to counteract this.

As we can see from Figure 3.16a there is little difference between the traditional metrics with both degree and eigencentrality consistently doing well, though not significantly better than HITS. HEE is yet again found to perform worse than all of the other metrics being statistically significantly worse than the equivalent degree metric at nearly all points ($p < 0.05$). Most importantly, there is little difference between the static and updating versions of each metric with none of them performing consistently better or worse than the other. This indicates that, for initial intervention at least, whether the information is up to date or not is of less importance. Given this, and the additional complexity and resource requirements for calculating the Updating metrics, Static metrics are likely sufficient in most cases.

Figures 3.16b and 3.16c show the results for the static and updating versions of the LIFE-DEGREE metric in the González model. Overall we find that there is little difference between the performance of LIFE-DEGREE and the equivalent degree metric, with $\omega = \{0.9, 0.7, 0.5\}$ not being consistently different over the range of n. This indicates that consideration of agent age is not enough to affect the efficiency of the placement one way or the other. When we consider *only* agent age, such that $\omega = 0$, we find that performance is substantially worse, doing as bad or worse than random placement. This is due to the fact that, for this value of ω, the placement metrics are essentially trying to choose which agents they believe will live the longest. Their poor performance highlights the difficulty in accurately choosing agents that will become influential later on. We include it as a comparison baseline for later runs. These results show that an agent's connectivity, indicated by its degree, is a much larger contributor to its

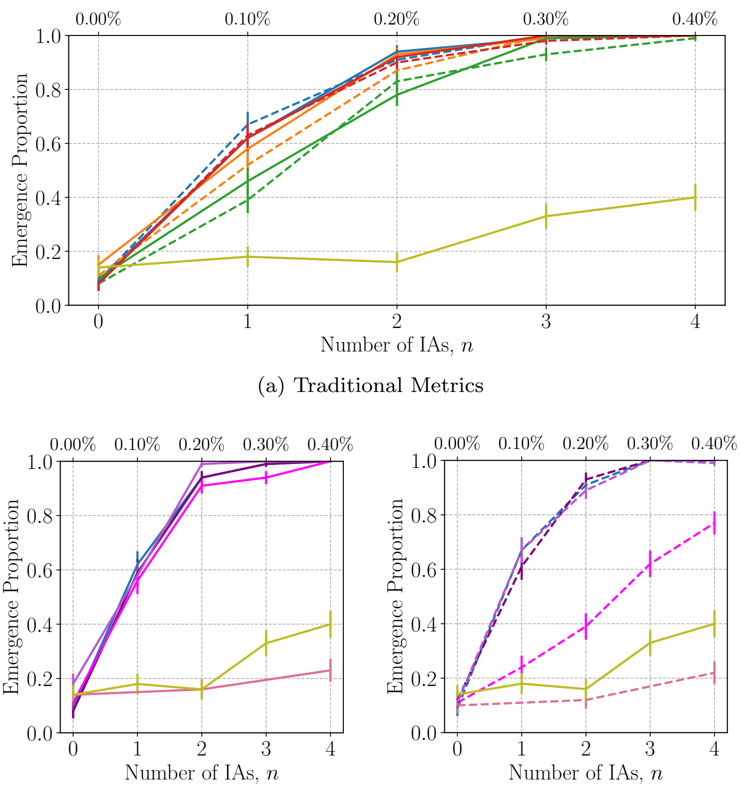

(a) Traditional Metrics

(b) LIFE-DEGREE - Static

(c) LIFE-DEGREE - Updating

Figure 3.17: Proportion of runs emerging the IA convention in the Ichinose RR network.

ability to influence others than how long that agent will remain in the system. The fact that considering age can only decrease the effectiveness of the chosen agents indicates that agents' short-term influence is a larger factor in convention emergence than choosing long-term targets.

Figure 3.17 shows the equivalent plots of the Ichinose RR model where nodes are removed at random and edges attached randomly as well upon node replacement. As in the González model, there is little difference in the efficacy of most of the metrics with the exception of HEE which performs noticeably worse. The major difference however is that the number of IAs required to cause 100% convention emergence is much lower, with nearly all metrics doing so by $n = 4$. This indicates that the Ichinose RR model is easier to influence, likely a factor

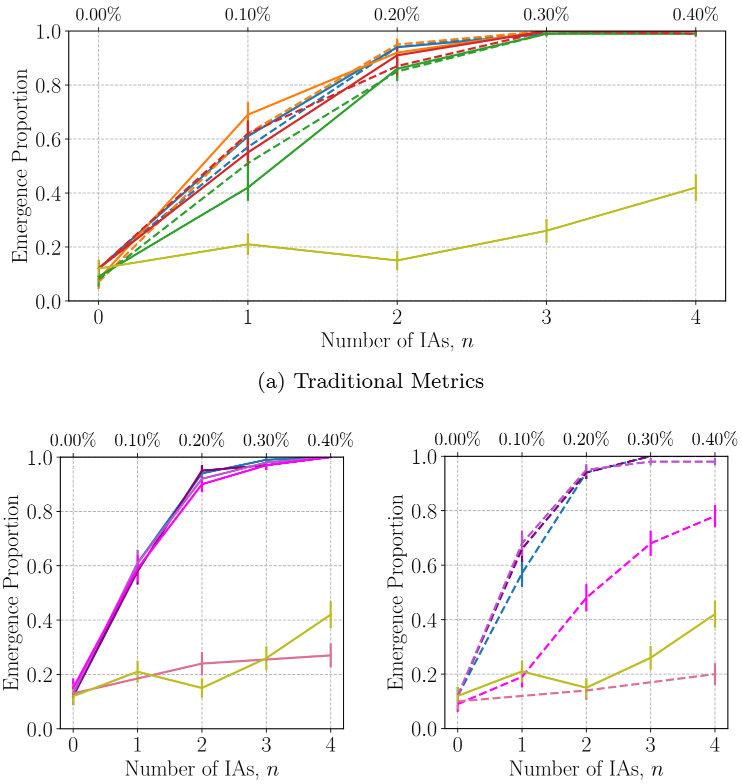

(a) Traditional Metrics

(b) Life-Degree - Static

(c) Life-Degree - Updating

Figure 3.18: Proportion of runs emerging the IA convention in the Ichinose RP network.

of the influential locations (high-degree nodes) not being removed as readily as they will be in the González model where they are likely to be older and closer to expiration.

Additionally, whilst most weightings of Life-Degree have no significant difference between their static and updating forms, $\omega = 0.5$ where agent age and degree have equal weighting, performs significantly worse in this topology than in González. As this doesn't appear in the static Life-Degree equivalent, this indicates that constantly using up-to-date age information is detrimental in the Ichinose RR model. We hypothesise that this is because of the random removal nature of topology meaning that agent age has little bearing as they may be removed at any time. Figure 3.18 shows similar behaviour in the Ichinose

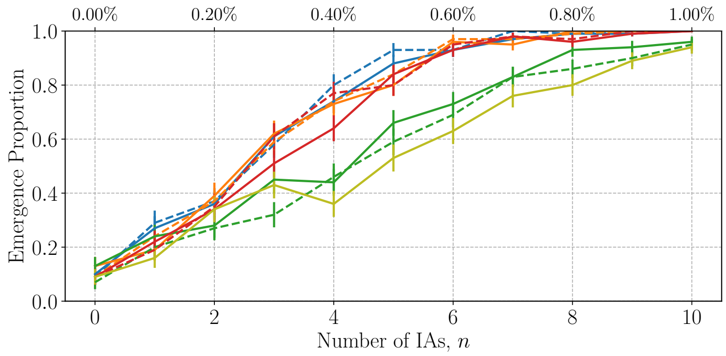

(a) Traditional Metrics

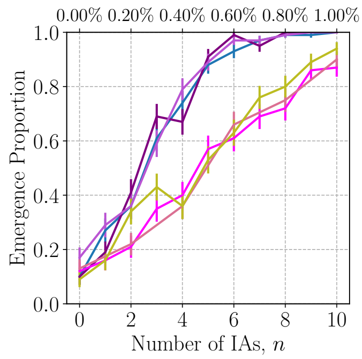

(b) Life-Degree - Static

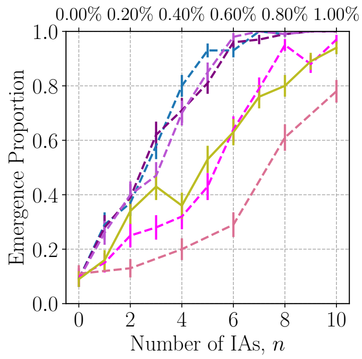

(c) Life-Degree - Updating

Figure 3.19: Proportion of runs emerging the IA convention in the Ichinose TR network.

RP model and hence adds weight to this theory as the major shared feature is the random nature of removal.

Figure 3.19 shows the effect of IAs in the Ichinose TR model, where the highest degree node is removed each timestep. This feature is a major difference in network dynamics between this version of the Ichinose model and the previous two and we see the effect of this in the results. The primary difference is that the number of IAs required is closer to that of the González model than either of the previous two Ichinose models with $n = 10$ being needed to consistently cause the desired convention emergence. Due to the removal of the high-degree node each timestep, which is likely to be ranked highly by all other metrics, the system is constantly losing IAs that have only been placed there for

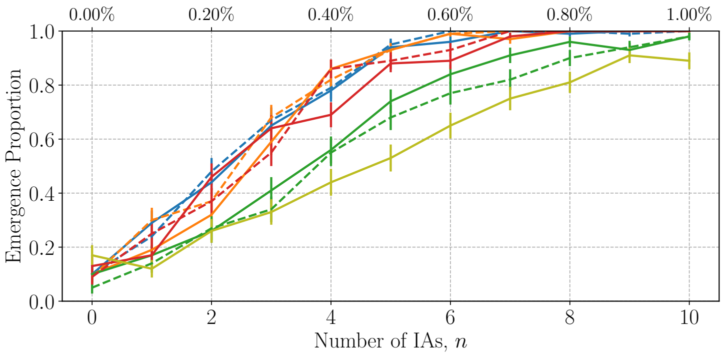

(a) Traditional Metrics

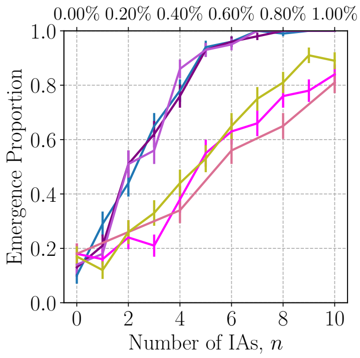

(b) Life-Degree - Static

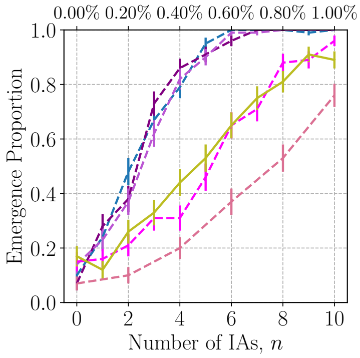

(c) Life-Degree - Updating

Figure 3.20: Proportion of runs emerging the IA convention in the Ichinose TP network.

a little while meaning that more IAs are required to constantly and consistently spread the desired convention and thus counteract this. Also of interest is that HEE performs poorly even by the previous standards, with a marked decline in efficacy compared to the other metrics, making it nearly as poor as random placement.

When examining the performances of Life-Degree we find that the poor quality of $\omega = 5$ that was present only in the updating metric in the other Ichinose models is also present in Static Life-Degree here with performance comparable to random placement. This is likely due to the fact that age is now tied strongly to degree as those vertices in the graph that have been around for longer are (i) more likely to be targeted for removal sooner and (ii) more likely to have accrued edges. As such, there is a redundancy in the information contained within LIFE-DEGREE that acts to its detriment. This is corroborated by the results of the Ichinose TP model shown in Figure 3.20 where similar behaviours are observed. The two sets of otherwise similar behaviour (RR and RP, TR and TP) indicate that the removal method in the Ichinose model is by far the most prominent feature that determines network dynamics, rather than the attachment method.

Overall, we have shown that convention emergence is possible in dynamic topologies of various types and that, as in static networks, small proportions of the population being used as IAs can dramatically affect this emergence, directing it to the desired outcome. We have shown that the age of the information used to select IA locations is of little importance when considering initial intervention and that consideration of agent age, choosing agents that will remain in the system for longer, is of less benefit than simply choosing highly influential agents and replacing them as necessary.

3.6 Late Intervention and Destabilisation in Static Networks

Having shown that rapid and robust convention emergence is possible using IAs in a range of networks, both static and dynamic, we now change our focus to the matter of *destabilising* an existing convention and replacing it with another of our choosing. These *late interventions* would allow system designers to resolve suboptimal conventions without having to change the internals of agents and instead focus on the use of IAs to facilitate this.

We begin by considering synthetic static networks as before. We generate scale-free topologies of sizes 1000 and 5000 using the Barabási-Albert model with settings $m_0 = m = 3$ and additionally generate small-world topologies of the same sizes with $ce = 1$ and $l = 1$. Both of these models with these settings have been shown to allow conventions to emerge unaided within the population and as such allow us to focus our efforts on the destabilisation of these once they have emerged. We utilise the same placement metrics and interaction model as before. The work thus far in this chapter has shown that the coexistence of multiple stable conventions under this model does not occur, with a single strategy dominating almost the entire population. As such we do not seek to address the nature of multiple conventions in this thesis to any great effect.

Initially we examine the effect of introducing a varying number of IAs into the population indefinitely. In Chapter 5 we will see the effect that only including them temporarily has on destabilisation. To establish a baseline for the minimum number of IAs required to cause destabilisation, we introduce a set of IAs at a time after conventions have already emerged. This was found to have occurred by $t = 1500$ in the scale-free topologies and $t = 2500$ in the small-world topologies. The IAs then remain in the system until the end of the simulation which is set for 5000 timesteps after their introduction to allow comparisons between both topologies.

For these results we use focus on what we term *aggressive destabilisation* such that IAs are all assigned the same fixed strategy which is chosen uniformly at random from all strategies that are not the current dominant convention with the intention of causing the chosen strategy to replace the dominant convention. As before, IAs are placed at influential locations as determined by the metrics used: degree, eigencentrality, HEE and HITS. We also consider random placement as a baseline for performance.

We start by studying the effect that late intervention has on the convention memberships within the systems. We run 100 simulations for each setting and plot the average membership sizes for both the dominant and selected IA strategy for each timestep. IAs are placed by degree initially.

Figure 3.21 shows these results for scale-free graphs with 1000 nodes and different numbers of IAs. Figure 3.21a introduces 20 IAs into the system at $t = 1500$ and we can see that they have an effect almost immediately. The membership of the dominant convention falls, but more than can be accounted

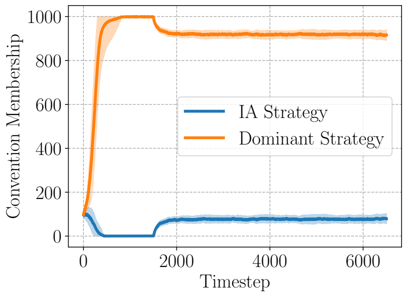

(a) 20 Intervention Agents

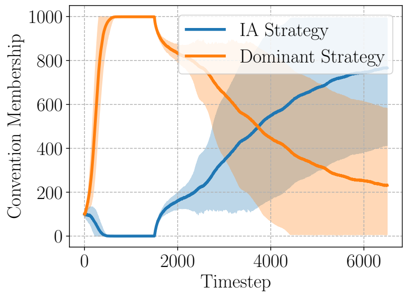

(b) 30 Intervention Agents

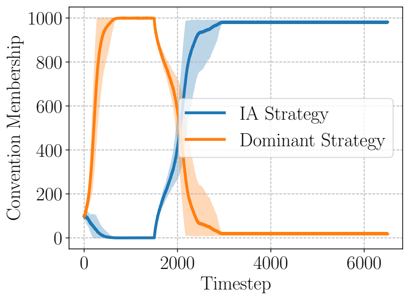

(c) 40 Intervention Agents

Figure 3.21: The effect of late intervention IAs on scale-free graphs. The shaded regions represent the standard deviations.

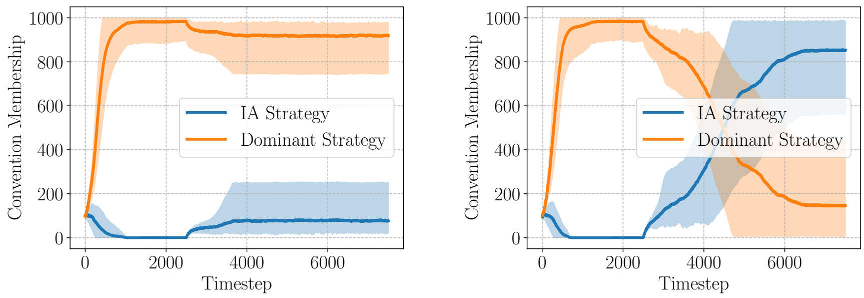

(a) 60 Intervention Agents

(b) 80 Intervention Agents

Figure 3.22: The effect of Intervention Agents on small-world graphs

for by the IAs themselves indicating they are successfully causing agents in their local area to switch away from the dominant convention. However, this reduction soon stabilises with little variation between all runs, indicating that they are unable to cause further agents to switch. Increasing the number of IAs even slightly to 30, as in Figure 3.21b, we see a drastic shift in behaviour with widely varying performance between runs. The destabilisation continues beyond the initial "dip" on average with a steady decline in the average membership size of the dominant convention and a similar climb for the IA convention, showing that the IA convention is switching members of the dominant one away from it. This is not guaranteed however and takes a substantial time. In comparison, Figure 3.21c shows that insertion of 40 IAs causes the entire membership of the dominant convention to switch, within only 1-2000 timesteps. The variation between runs is substantially reduced with the behaviour mostly consistent across all of them. These results show that there is a minimum number of IAs required to induce destabilisation and the relatively small range of IAs from no permanent effect to guaranteed destabilisation indicates that there is a "critical value", a tipping point beyond which the number of IAs guarantees effect, much like in initial intervention. Increasing the number of IAs beyond this minimum was found to accelerate the destabilisation further.

Results for 1000 node small-world networks are shown in Figure 3.22. Whilst the overall behaviour is similar, in that there is a critical number of IAs after

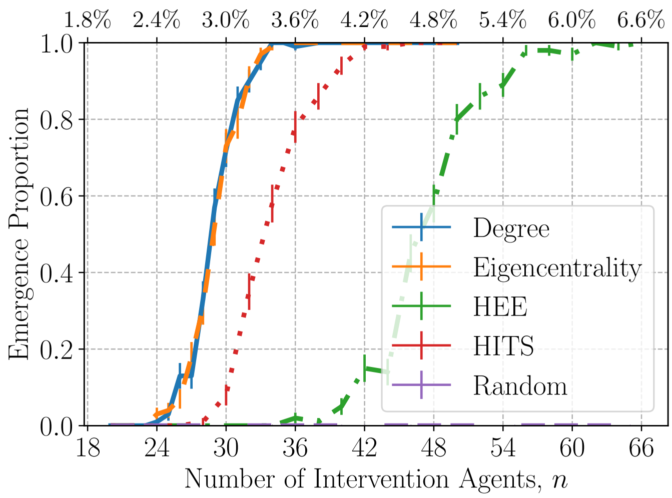

Figure 3.23: Proportion of runs where the IA strategy emerges as a replacement convention for the 1000 node scale-free network.

which destabilisation will occur, the behaviour pre-transition is less well-defined and there are some distinctions to highlight. In particular, the characteristic "dip" that occurs in scale-free topologies is much more variable in the small-world topologies with the level to which the reduction occurs varying more between runs. Additionally, the number of agents required is substantially higher with 80 agents, shown in Figure 3.22b, still not causing as rapid destabilisation as 40 agents do in the scale-free topologies. Whilst we have previously seen disparities in the behaviour of scale-free and small-world graphs, this highlights that it is present for destabilisation as well and is likely primarily due to the lack of "hub" nodes that act as highly influential individuals.

Having shown that destabilisation is possible, and that there appears to be a sharp transition between no effect and guaranteed destabilisation in the two topologies, we now apply these findings to the other metrics. We now place varying numbers of IAs using all the metrics previously mentioned and monitoring the proportion of 100 runs where the IA strategy emerges as a new convention to the 90% Kittock level, hence destabilising and replacing the previously dominant strategy.

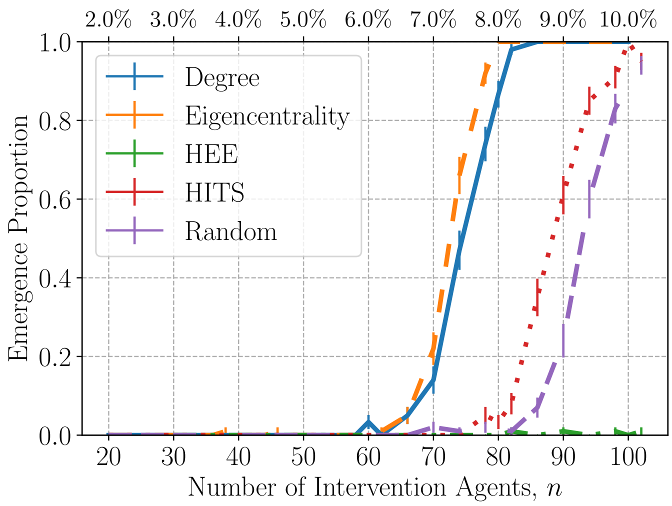

Figure 3.24: Proportion of runs where the IA strategy emerges as a replacement convention for the 1000 node small-world network.

Figure 3.23 shows the results of this for the 1000 node scale-free network. As expected, for each metric there is a sharp transition, or phase shift, between having no effect (emergent proportion = 0) and guaranteeing destabilisation (emergent proportion = 1), that occurs over a range of approximately 10 additional IAs for degree and eigencentrality, 16 for HITS and 24 for HEE. Additionally, unlike in initial intervention, there are distinct performance differences between the different metrics with degree and eigencentrality performing markedly better than HITS and even more so over HEE. This indicates the general applicability of degree and eigencentrality in both types of intervention with the other metrics being less so.

Figure 3.24 shows the similar results for the 1000 node small-world graph. Though the number of IAs required is substantially larger (both in absolute and relative terms), the same pattern emerges with transitions from no effectiveness to full effectiveness occurring over small ranges of IAs compared to the number required. The performance of HEE is even worse in this setting with random placement performing better and HITS only marginally better than that. There is a slight difference in the performance of eigencentrality and degree here though

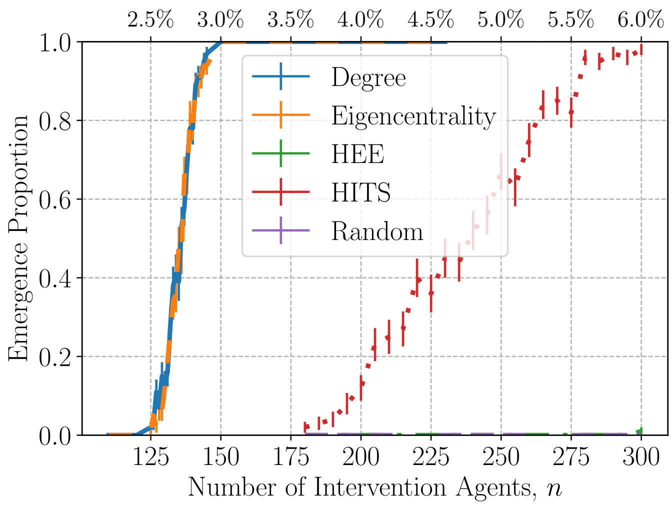

Figure 3.25: Proportion of runs where the IA strategy emerges as a replacement convention for the 5000 node scale-free network.

eigencentrality is only statistically significantly better at 74, 78 and 80 IAs ($p < 0.05$).

We also consider different sizes of agent population, as before. Figure 3.25 shows destabilisation in the 5000 node scale-free topology. Whilst the absolute number of IAs required increases, unlike in initial intervention where the absolute number was relatively invariant, the number of IAs as a proportion of the population is still around 3% indicating that the relative number of IAs may be a deciding factor in the nature of destabilisation. Degree and eigencentrality continue to outperform all others with the performances of HITS and HEE both decreasing, requiring a larger relative number of IAs than was needed in the 1000 node graph.

Figure 3.26 shows the same for the 5000 node small-world network. Again, the performance of both degree and eigencentrality against the relative number of IAs is consistent with the difference between them even more marked (statistically significant between 365-395 IAs, $p < 0.05$). Interestingly, HITS and random placement both maintain their relative performance as well, indicating that in small-world topologies these placement metrics may still be effective.

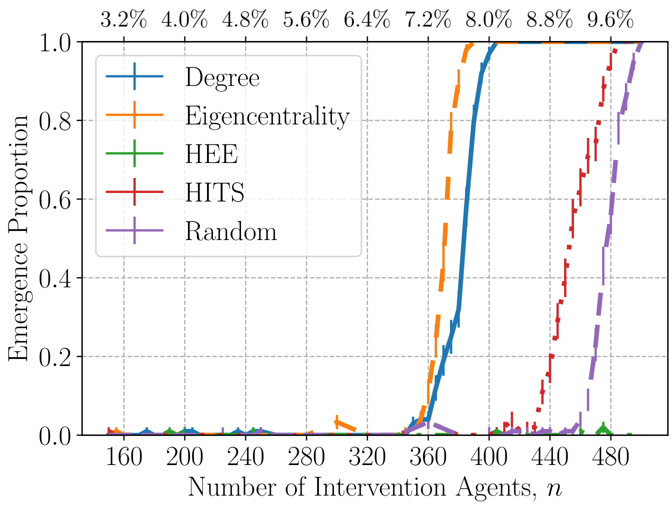

Figure 3.26: Proportion of runs where the IA strategy emerges as a replacement convention for the 5000 node small-world network.

3.6.1 Late Intervention in Real-World Topologies

As we did for initial intervention we also consider the performance of the metrics for eliciting destabilisation in the real-world topologies of CondMat, Enron and Twitter. Given their differences compared to synthetic networks even for initial intervention the complex and dense nature of their topologies is likely to induce different performances than those just shown.

The results are shown in Figure 3.27 for the proportion emergence over 100 runs in each topology. Unlike in synthetic static or dynamic networks, our experiments show that conventions are unlikely to emerge unaided in the real-world topologies, even to the 80% Kittock level within any reasonable timeframe. As such, and as we are concerned with destabilisation of existing conventions, we artificially *saturate* the population using 500 IAs to induce an initial convention emergence. This then allows us to focus on the destabilisation. Whilst this is less preferable than allowing the conventions to emerge naturally we do not believe there to be any salient difference between the two types. These saturation IAs are kept in the system until an appropriate convention is guaranteed to have emerged which occurred by timestep 500 in the Twitter and Enron networks

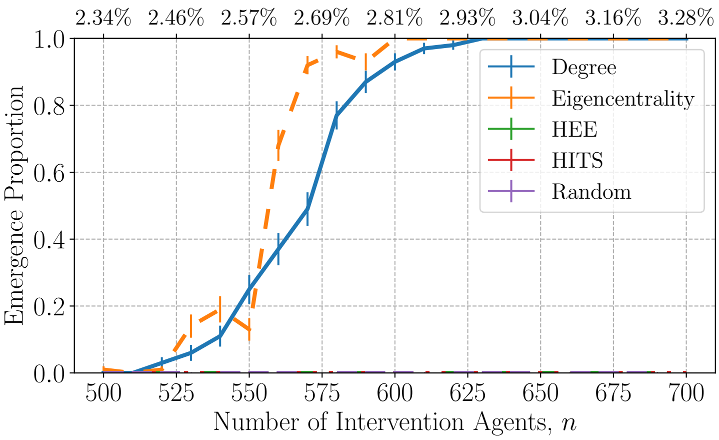

(a) CondMat

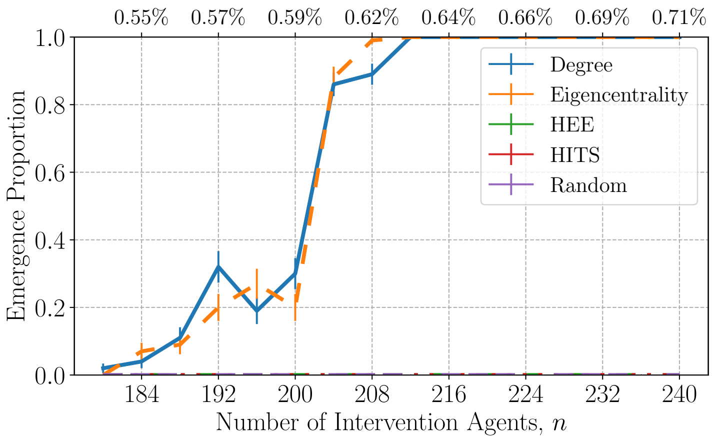

(b) Enron

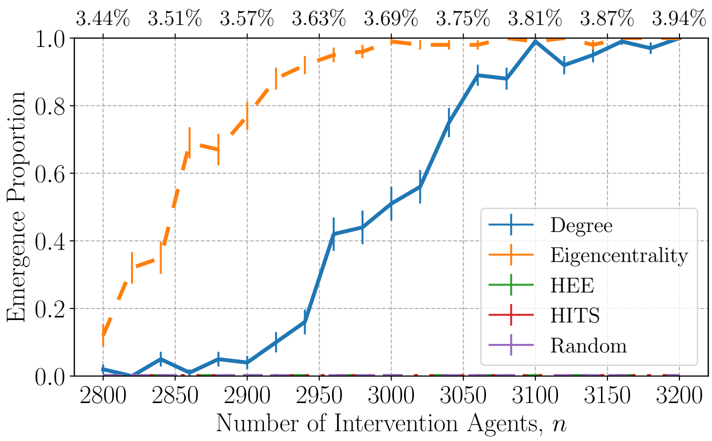

(c) Twitter

Figure 3.27: Proportion of runs where the IA strategy emerges as a replacement convention in real-world networks.

and by timestep 1500 in the CondMat network. At this time the saturation agents were removed and the destabilisation intervention was allowed to begin. The simulations all ran for 15000 timesteps to ensure that likely destabilisations would occur. As in the case with initial intervention, we use the 80% Kittock threshold rather than the 90% due to the networks being unlikely to emerge conventions at that level.

As can be seen in Figure 3.27, there are marked differences between destabilisation in each of the networks. Both Twitter and CondMat require relative proportions of IAs comparable to scale-free topologies, needing about 3-4% of the population to be IAs before destabilisation is guaranteed to occur. By comparison, the Enron network needs substantially less than all of the other topologies, needing only 0.64% to guarantee destabilisation when placing by degree or eigencentrality. This highlights the fact, observed in the difference between scale-free and small-world topologies, that destabilisation is more sensitive to topological particulars than initial intervention was. Underpinning this is the stability of the convention that we are trying to destabilise and these results show that stability is directly tied to topology, as was argued by [Villatoro et al., 2011a].

HEE and HITS both perform substantially worse in these networks with none of the ranges of IAs investigated responding to their use. This adds additional weight to our hypothesis that degree and eigencentrality are the best metrics of those investigated due to their general nature. Indeed, within these topologies we see eigencentrality clearly outperforming degree for the first time in both the CondMat and Twitter networks, and to a larger extent in the latter. This implies that there are not as interchangeable as initially suspected and instead may have substantially different performances depending on topology.

3.7 Late Intervention and Destabilisation in Dynamic Networks

We have thus far shown that destabilisation of existing conventions is possible in static networks and that, whilst the absolute number of IAs required to do so might change, the relative proportion of the population is mostly invariant with all topologies observed so far, regardless of size and underlying differences, requiring at most 8% in order to destabilise the convention.

We now look to investigate the nature of destabilisation in dynamic topologies. Given the fundamental differences between them and the static topologies, we expect the system dynamics to differ here too. Due to the higher numbers of IAs needed to cause destabilisation, the fact that they may be removed or expire at any moment or that their ability to influence may change over time increases in importance when considering destabilisation. Up-to-date information on node metric values, so that a suboptimal node is not left as an IA will be paramount.

Similar to our approach in synthetic static networks, and so that the results are representative of the general case, we allow a convention to naturally emerge without the use of IAs to encourage it. It was found that conventions always emerged before timestep $t = 1500$ in all topologies and, as such, insertion of IAs occurs at this time. This also means that the topology will have had longer to stabilise. Whilst this is less likely to have an effect in the González model due to the already required burn-in, the Ichinose models will have had more time to diverge and so we can expect even more pronounced differences in their behaviour than was observed in initial intervention. As before, the fixed strategy assigned to the IAs is chosen uniformly at random from all actions excluding the established convention and assigned to all agents. Our model is otherwise identical to that used for initial intervention and we perform 100 runs of each setting to calculate the proportion. Each simulation is run for 5000 timesteps as this was found to be long enough for likely destabilisation to occur.

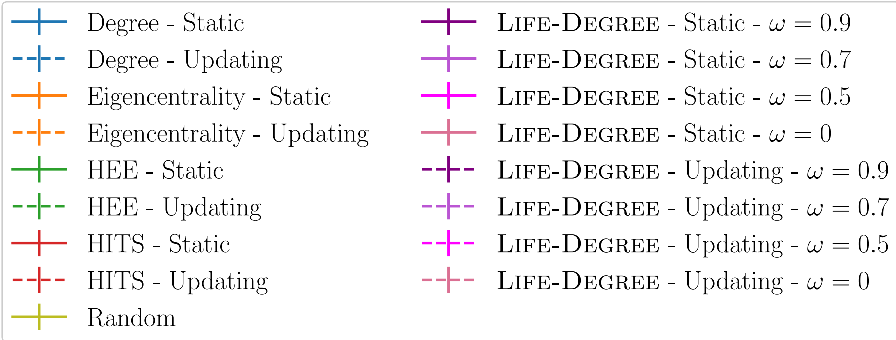

Figure 3.28: Legend for all metrics used in late intervention in dynamic networks.

The same method of identifying the large number of different metrics is used as before and we reproduce the supplied legend here for ease of access in Figure 3.28.

We begin with a consideration of the effectiveness of destabilisation in the González network as shown in Figure 3.29. As can be seen, as was found in static networks, the number of IAs needed to cause destabilisation is much larger than the number of IAs needed to elicit initial convention emergence. Indeed, the relative amount is also larger than equivalently sized static networks with even the best performing metric requiring ~5% of the population as IAs in order to cause destabilisation (compared to ~3.5% in the scale-free networks). This is despite the González model sharing a number of features with the scale-free network such as a power-law degree distribution and preferential attachment. We can conclude from this that it is the dynamic nature itself that requires additional efforts to destabilise conventions, affording them additional stability. However, it is worth noting that the transition from no effect to full effect occurs much faster in the González model than the scale-free topology with a difference of just 6 IAs between the two states. We can conclude that the "critical value" of IAs needed is over a narrower range than in the static topologies.

As was found in the static networks, there is now a marked difference in the efficacies of the various placement metrics with HEE and HITS performing much

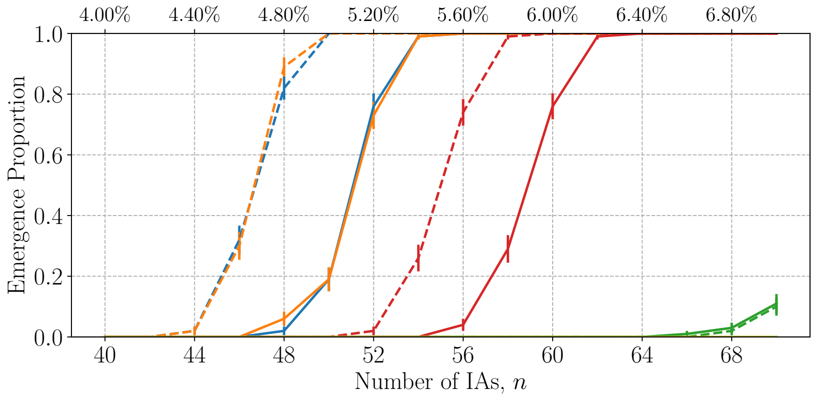

(a) Traditional Metrics

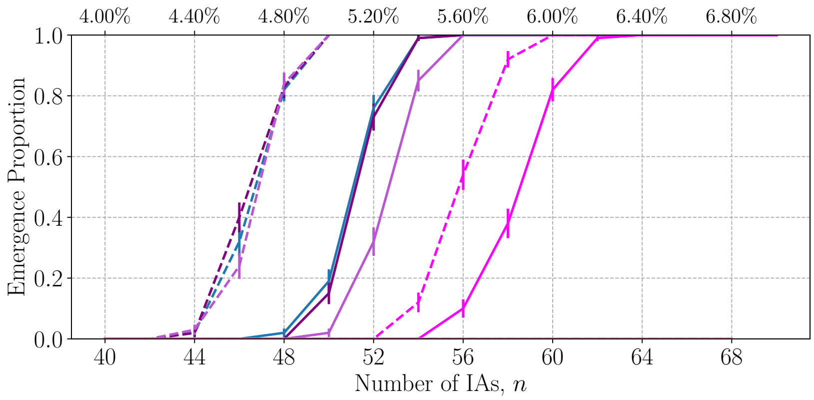

(b) LIFE-DEGREE

Figure 3.29: Proportion of runs where the dominant convention is replaced with the IA convention in the González network.

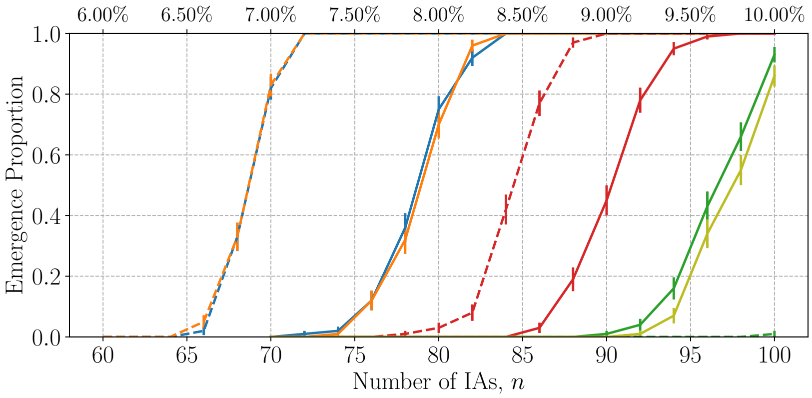

(a) Traditional Metrics

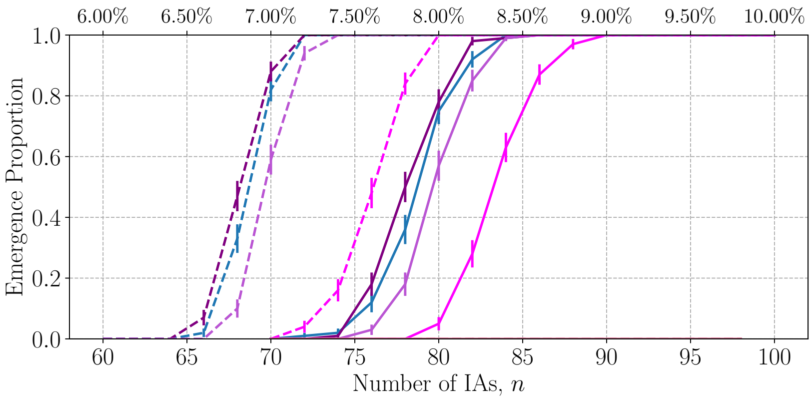

(b) LIFE-DEGREE

Figure 3.30: Proportion of runs where the dominant convention is replaced with the IA convention in the Ichinose RR network.

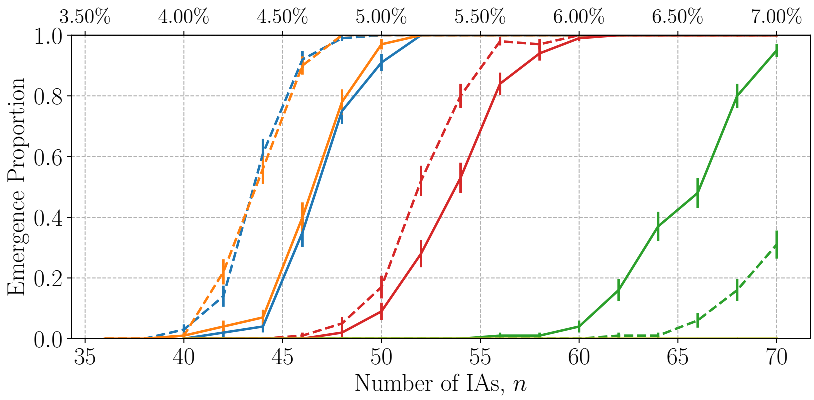

(a) Traditional Metrics

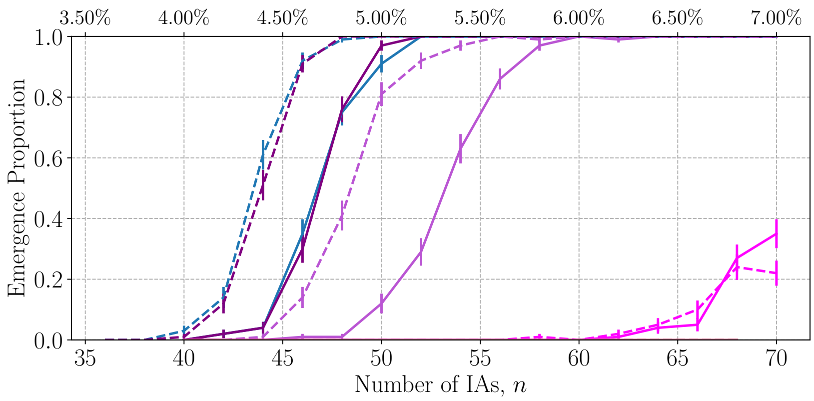

(b) LIFE-DEGREE

Figure 3.31: Proportion of runs where the dominant convention is replaced with the IA convention in the Ichinose RP network.

worse than degree or eigencentrality. Additionally, the updating versions of each metric now perform significantly better than their static equivalents across all metrics with updating degree/eigencentrality performing best overall. This indicates that having up-to-date information on the level of influence a particular agent is capable of is much more important when attempting to destabilise an existing convention. Ensuring that the desired convention has as much reach as possible is necessary to overcome the precedent of the existing convention and avoid it self-reinforcing due to suboptimal agent choice.

Additionally, the LIFE-DEGREE weightings have differing levels of performance compared to their efforts in initial intervention with $\omega = 0.5$ performing

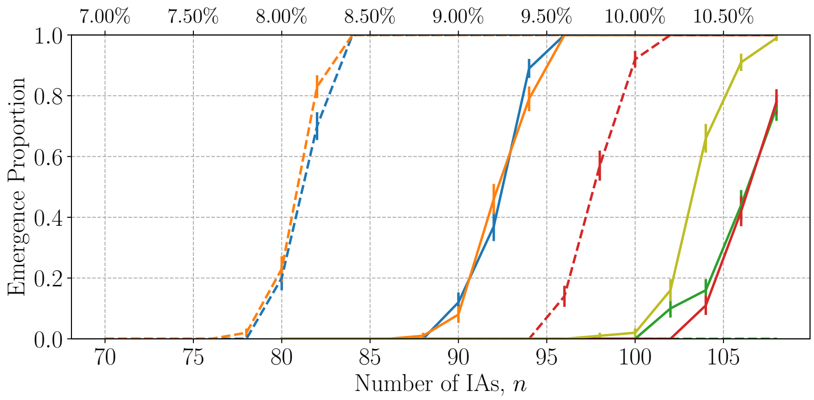

(a) Traditional Metrics

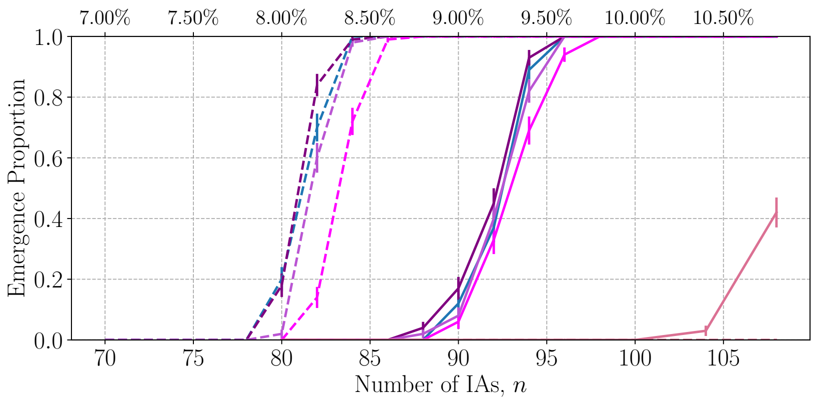

(b) LIFE-DEGREE

Figure 3.32: Proportion of runs where the dominant convention is replaced with the IA convention in the Ichinose TR network.

markedly worse than the other weightings which in turn only approach the equivalent degree metric in efficacy. This indicates that consideration of agent influence is much more important than ensuring agent longevity in our selection and is backed up by the fact that $\omega = 0$ performs so poorly it does not cause any effect in the ranges shown.

Figures 3.30 and 3.31 show destabilisation in the Ichinose RR and Ichinose RP topologies respectively. In contrast to their performances in initial intervention, these two topologies exhibit noticeable differences when considering destabilisation. Ichinose RR, with the best performing metric, requires ~7.25% of the population to be IAs before destabilisation is guaranteed whereas Ichinose

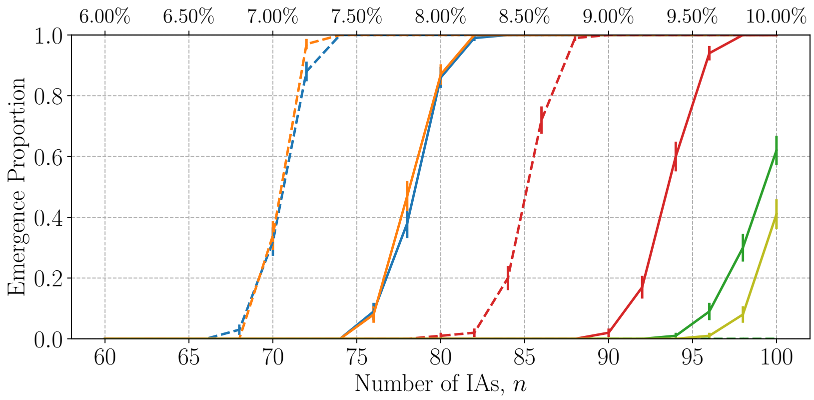

(a) Traditional Metrics

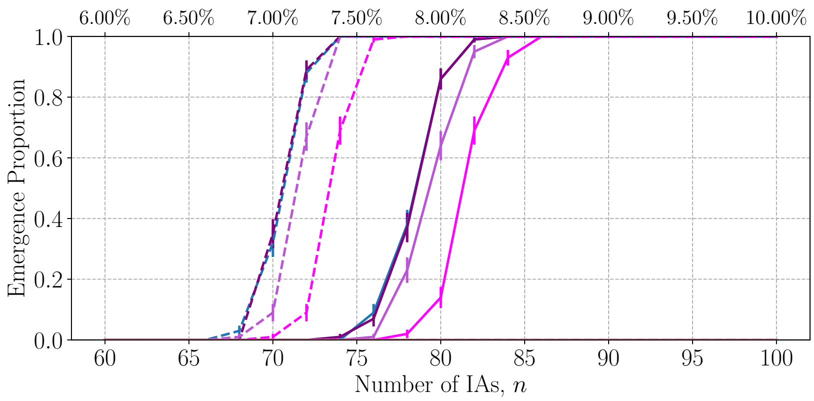

(b) Life-Degree

Figure 3.33: Proportion of runs where the dominant convention is replaced with the IA convention in the Ichinose TP network.

RP only requires ~5%. Additionally, the difference between the performance of the updating metrics compared to their static equivalents is much less substantial in the Ichinose RP topology than in the Ichinose RR topology. Whilst not important when considering initial intervention the method of edge attachment used has a noticeable effect when it comes to destabilisation. The performance of Life-Degree is similarly affected by these differences with performance in RR being much better than in RP, though still offering no improvement over the pure degree equivalent.

Similarly, Figures 3.32 and 3.33 show destabilisation within the Ichinose TR and Ichinose TP topologies respectively. As in each of the other topologies, the updating metrics perform markedly better than their static equivalents, requiring ~12% fewer agents in the Ichinose TR topologies for degree and eigencentrality placement. As was found in the RR and RP topologies, this improvement is lessened in the Ichinose TP topology by comparison with updating degree and eigencentrality requiring ~9% fewer agents, though this difference is less than it was in the random removal topologies. This lends further weight to the notion that the attachment method of the Ichinose model is the primary driver that affects the supremacy of the updating metrics. It indicates that topologies lacking preferential attachment are more susceptible to the effects that out of date information can have as the selected nodes are more likely to decrease in influence in these topologies. Both TR and TP topologies require more IAs to cause destabilisation than the equivalent RR and RP topologies, likely due to the consistent removal of influential nodes, as discussed when considering initial intervention.

Overall, these results show that destabilisation is possible in dynamic topologies as well as static although the proportion of IAs needed to do so is higher. The constantly in flux nature of both the nodes and edges lend additional stability to the established convention likely due to the ability of the established convention to "reinforce" areas of falling support as new links are made. In all topologies we also find that having up-to-date information and selection of influential nodes is substantially more important than when trying to elicit initial convention emergence with updating metrics performing markedly better in all scenarios. Our findings also indicate that concerns of agent longevity, tied to the fact that up-to-date information is preferable, are more detrimental in this domain. Attempting to utilise agents that will be in the system for longer is of much less importance than maximising the influence of the selected IAs. As with the findings for destabilisation in static networks, both HEE and HITS are substantially worse when being used as placement metrics compared to degree and eigencentrality, indicating further than these are the best generally applicable metrics to use.

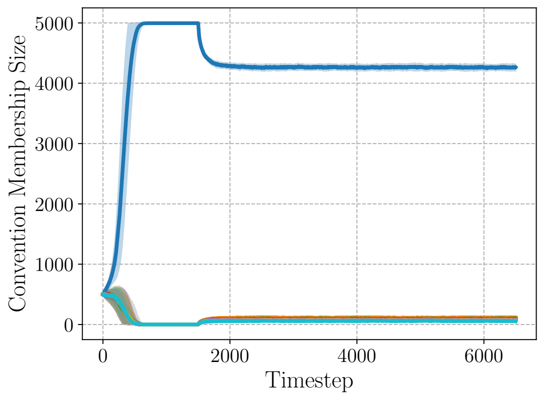

(a) 250 IAs

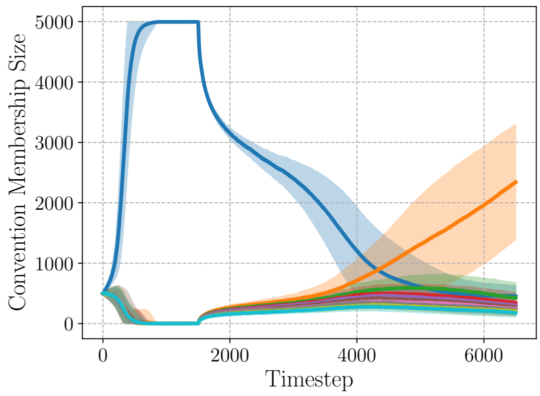

(b) 450 IAs

Figure 3.34: Effect of passive destabilisation on convention membership size in the 5000 node scale-free topology.

3.8 Passive Destabilisation

Having established that destabilisation is possible and highlighted the various factors in both dynamic and static networks that affect it, we now turn our attention to the concept of *passive destabilisation*.

We can categorise destabilisation efforts into two main types: *aggressive destabilisation* where we wish to cause the collapse and active replacement of the current dominant convention, and *passive destabilisation* where we are not concerned with the replacement of the dominant convention but rather only seek to cause it to collapse. All work thus far has utilised aggressive destabilisation but we now explore passive destabilisation as a means to return the system to a state where a convention can emerge naturally once more, without preference for what they convention may be. This is of use in scenarios where the system designer may have some equilibria they actively *do not* want as conventions but otherwise do not wish to favour one over others. In order to facilitate passive destabilisation we utilise IAs as before but instead of assigning them a uniform strategy we instead assign fixed strategies to them uniformly at random from the set of actions that are not the dominant one.

We begin as we did when investigating aggressive destabilisation by examining the effects of the destabilisation efforts on convention membership size and this is shown in Figure 3.34. IAs are placed by degree and the average membership size of each convention (combined by rank) at each timestep is calculated over 100 runs. IAs are introduced at timestep $t = 1500$ after a dominant convention has emerged and their strategies are assigned as discussed. Figure 3.34a shows the results when 250 IAs are placed within the system. Despite being 100 IAs larger than the number required to guarantee aggressive destabilisation in the same network, we can see that this set of IAs is insufficient to cause the dominant convention to collapse, instead simply reducing the number of members. Whilst the size of the dip indicates, as before, that the IAs are successful in switching others away from the established convention the unfocused nature of passive destabilisation means that the established convention is more resilient. In particular, the decrease is very stable with little deviation after the initial fall indicating that, as with aggressive destabilisation, there is likely a critical number of IAs that must be reached to prevent the dominant convention from stabilising.

Figure 3.34b shows the results when 450 IAs are placed into the same system. Despite requiring much larger numbers of IAs we can see here that passive destabilisation is possible with the dominant convention collapsing with little variance in all runs. The process is slower than in aggressive destabilisation however, as is the growth of the replacement convention that begins to emerge.

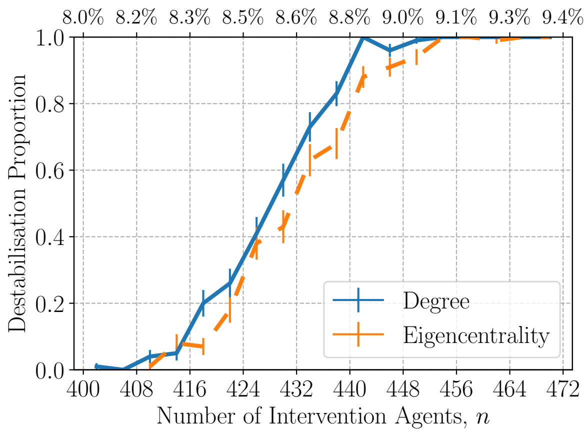

Figure 3.35: Passive destabilisation in the 5000 node scale-free topology. The proportion is the proportion of runs that caused the dominant convention to fall below the 30% Kittock count.

Indeed, longer runs show that this emergent convention takes a long time to grow and that its maximum size is much lower than that achieved by the original convention due to the suppressing effect of the passive IAs. Similar patterns are observed in both small-world and dynamic topologies.

Due to this, we instead focus on the number of IAs that is required to consistently destabilise the dominant convention without concerning ourselves with the level to which a new convention emerges afterwards. To facilitate this we monitor the proportion of runs in which the dominant convention, after the intervention begins, falls below the 30% Kittock level, viewing this to be representative of a substantial collapse in the level of support for the established convention. We refer to this as the *destabilisation proportion*.

Figure 3.35 shows the results of this for the 5000 node scale-free network with the same intervention model used before: IAs are introduced at $t = 1500$ and we monitor the proportion of these that cause a collapse in the dominant convention in the 5000 timesteps following. We utilise both degree and eigencentrality placement of the IAs as these were found to be the most effective in aggressive destabilisation.

As can be seen in the figure, both the absolute number and hence population

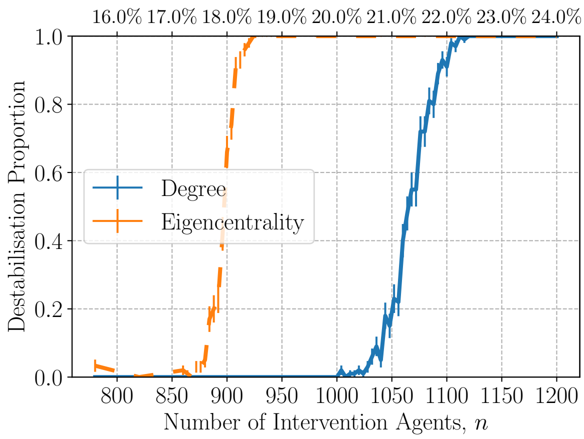

Figure 3.36: Passive destabilisation in the 5000 node small-world topologies. The proportion is the proportion of runs that caused the dominant convention to fall below the 30% Kittock count.

proportion of IAs needed to cause passive destabilisation is much higher than that needed for aggressive destabilisation, requiring almost 2.5 times more IAs. The transition between no effect and full effect is also affected with it occurring over a range of about 1% of the population as opposed to 0.5% seen before. Most importantly, there is a significant difference between the performance of degree and eigencentrality placement with eigencentrality performing worse to a statistically significant level at nearly all points of interest. Whilst degree and eigencentrality are strongly correlated in scale-free graphs, at the number of IAs being utilised here using one metric over the other will select different nodes. Nodes with high eigencentrality are linked to other nodes with high eigencentrality but are not necessarily linked to many nodes overall; being linked to by a few, very important, other nodes will produce high eigencentrality. For our purposes, this means there may be nodes ranked highly by eigencentrality but who actually have a small set of other nodes they are able to influence (and the members of that set have likely already been selected themselves). This is less likely to have an effect at lower numbers of IAs but will be noticeable at higher levels and this is what we see here.

Figure 3.36 shows the same but for small-world topologies. Here IAs insertion occurs at $t = 2500$ and we similarly monitor the destabilisation proportion achieved by runs in the 5000 timesteps following. Passive destabilisation in small-world topologies similarly requires ~2.5 times as many IAs as aggressive destabilisation which results in needing up to 22% of the population before full effectiveness is achieved. Of most interest is that, in this setting, eigencentrality actually outperforms degree to a significant margin requiring almost 4% less of the population (or 200 fewer IAs) in order to cause the same level of collapse. This shows that passive destabilisation can benefit from exploiting different aspects of the underlying topology more so than aggressive destabilisation and the differences between high-degree and high-eigencentrality nodes in the small-world topology facilitate this. The lack of hub nodes in the small-world topology is again likely the cause of this disparity but it is still of interest that the high-eigencentrality nodes are more effective targets due to their connections to other high-influence nodes. This is the other side of the scenario discussed above, as there are many nodes with the same degree (due to the lattice nature of the small-world generation) but those well-linked individuals (those with high eigencentrality) act like the hub nodes in scale-free graphs. Their high eigencentrality makes them different to the myriad of nodes with the same degree and benefits the convention emergence by allowing links to disparate parts of the network.

The distinction between the performances of eigencentrality and degree does not appear in the dynamic topologies and so we only plot the degree metrics in Figure 3.37 for clarity. Figure 3.37 shows the performances of both static and updating degree for destabilisation as the best performing metrics investigated for all dynamic topologies. Again, a factor of ~2.5 times as many IAs compared to the best performing metrics for aggressive destabilisation is required in order to cause passive destabilisation. A major difference however is that, for passive destabilisation, it is the *static* metric placements that perform best in all topologies. Whilst the differences between the static and updating metrics differ in each topology, this consistency indicates that it is more important in passive destabilisation to leave IAs in a location consistently rather than potentially

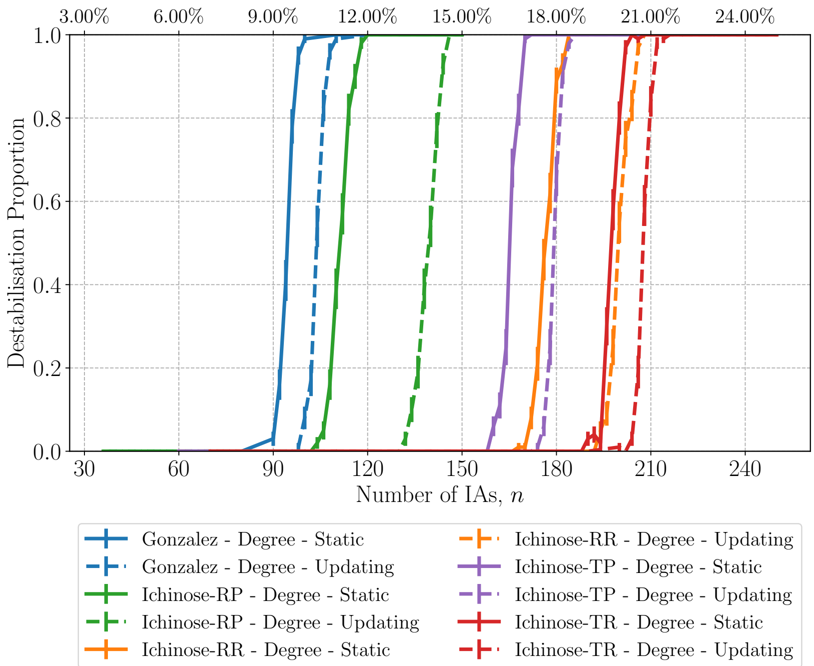

Figure 3.37: Passive destabilisation in dynamic topologies. The proportion is the proportion of runs that caused the dominant convention to fall below the 30% Kittock count.

changing them every timestep as this gives the IAs enough time to cause their local agents to switch away from the dominant convention without the work being undone as soon as they are removed. This is highlighted by the static metrics needing approximately twice as many IAs to elicit the same destabilisation as in aggressive whilst the updating metrics require approximately three times as many. This runs contrary to the requirements for aggressive destabilisation in these topologies and highlights the differences between them.

Overall however, passive destabilisation of existing conventions has been shown to be possible. Whilst it requires larger numbers of IAs to cause the same level of effect, being able to destabilise an established convention without having to actively choose a successor is of great benefit in a number of domains.

3.9 Differing Convention Spaces

The size of the convention space available is known to have marked effect on the nature of convention emergence within MAS. Griffiths & Anand [2012] showed that larger convention spaces slowed the rate at which conventions emerged initially across all topologies and Franks et al. [2013] and Salazar et al. [2010], when using the language coordination game which has a exponentially larger convention space showed results that also exhibited slower convention emergence than we have found here, often requiring on the order of 10^6 timesteps.

As such, we seek to explore the effect that larger, and smaller, convention spaces have on both initial and late interventions in both static and dynamic topologies. This also will show that our results are general and not expressly tied to the 10-action coordination game. We vary the number of coordination actions such that we can explore these effects by setting the number of actions from $\{2, 5, 10, 20, 100\}$. This provides a range of convention space sizes across multiple orders of magnitude. Our experimental setups are otherwise the same as used in the previous sections and we measure the proportion of 50 runs that come to the desired outcome. In each of the following figures, unless indicated

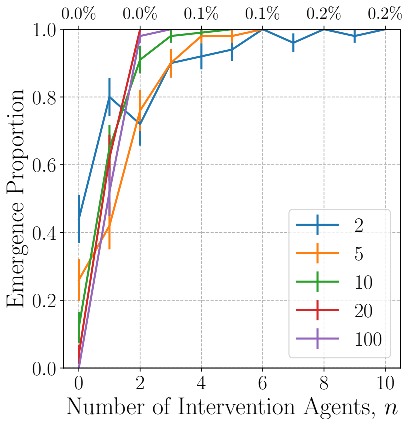

(a) Scalefree

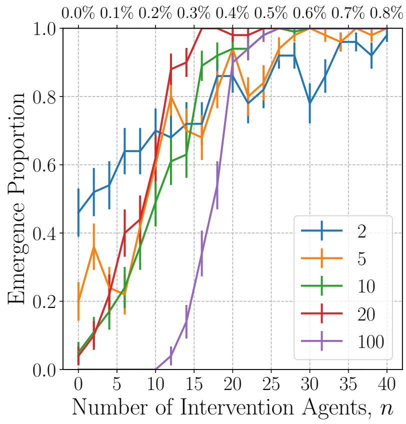

(b) Smallworld

Figure 3.38: The effect of convention space size on initial intervention in synthetic networks. The number of actions in the coordination game is shown in the legend.

otherwise, we show the performance of the best metric in each scenario: degree in static networks and updating degree in the dynamic networks. Similar effects were observed on each of the other metrics but did not change their relative rankings. For larger convention spaces[1], we find that conventions are unlikely to emerge naturally within the system and so, when considering late intervention, we artificially saturate these populations with a convention so that we may focus on destabilisation. This is done using the same methodology as Section 3.6 for the real-world networks.

Figure 3.38 shows the effect of convention space size for initial intervention in synthetic networks. In the scale-free topologies, there is little effect with conventions still being affected by roughly the same number of IAs. Whilst there are slight increases in efficacy from 5 to 10 to 20 action choices, this increase is one of diminishing returns with 100 indistinguishable from 20 and the differences not statistically significant at most points. It does however indicate that the use of IAs is indifferent to the convention space size in this setting with the agents able to direct convention emergence with low levels of IAs at all sizes.

[1]Convention spaces of size 100 in all topologies, and size 20 in the González, scale-free and small-world networks.

In contrast, small-world networks exhibit counter-intuitive and noisy changes in efficacy with varying convention space size. The amount of IAs needed to guarantee convention emergence is perhaps the best method of evaluating performance in the small-world networks and we find that there is little difference between the various convention space sizes with all but size 2 guaranteeing convention emergence with between 0.3% and 0.6% of the population. On the other hand, the number of IAs required to achieve smaller proportions of emergence increases dramatically with convention size, unlike in scale-free networks. This follows from the fact that, with larger numbers of action choices, we cannot rely on the times where the convention emergence would have occurred anyway to bolster the proportions and hence the effects of smaller numbers of IAs is reduced comparatively. For instance, whereas with a convention space of size 10 we can expect roughly 10% of runs to emerge our target convention regardless of the presence of IAs, with size 100 this will only occur 1% of the time and hence more IAs are needed to overcome this initial inertia. However, whilst larger convention spaces require more IAs to elicit small proportions of runs to emerge to the desired convention, the numbers needed to guarantee it are relatively invariant beyond size 2, as was the case in scale-free topologies.

Figure 3.39 shows how the effectiveness of IAs when used for late interventions changes with convention space size. The patterns here are more distinct as they are not affected by the natural convention emergence that occurs in initial intervention. In both scale-free and small world networks increasing the size of the convention space results in more IAs being required in order to cause destabilisation. This follows from the fact that the exposure of other agents to the strategy of the IAs, such that they learn they will be rewarded if they choose it, will be rarer with more action choices available for them to explore. However, the relationship is non-linear with little difference between the number of IAs needed at different convention space sizes; for instance in the scale-free network an increase in convention space size from 2 to 20 only requires 17% more IAs despite having 10 times as many options available to the population.

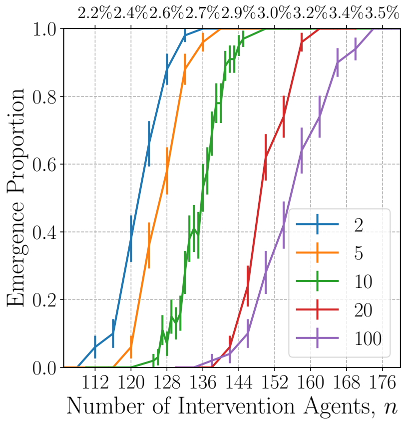

(a) Scalefree

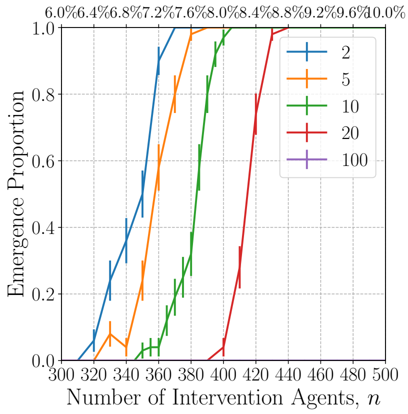

(b) Smallworld

Figure 3.39: The effect of convention space size on late intervention in synthetic networks. The number of actions in the coordination game is shown in the legend.

This indicates that the critical number of IAs needed to cause destabilisation is not directly tied to the convention space with a large component being independent of it. Of particular note, however, is that whilst in the scale-free network the performance with 100 possible conventions is not much worse than with 20, in the small-world topology a size of 100 results in no destabilisation with replacement occurring even at much higher numbers of IAs than the other convention space sizes require. This shows that in some topologies sufficiently large convention spaces make destabilisation much harder.

We also consider convention space size in initial intervention for dynamic networks. Figure 3.40 shows the effect of this for González networks. Whilst mostly similar to the effects found for scale-free topologies, the González network shows clearly that increasing convention space size makes initial intervention easier in this topology with the number of IAs needed to guarantee convention emergence decreasing with larger sizes. This does not occur continuously however, with a convention space size of 100 actually exhibiting no convention emergence at all over the ranges shown. Indeed, even increasing the number of IAs as high as 70, no convention emergence to the desired convention occurs. This demonstrates

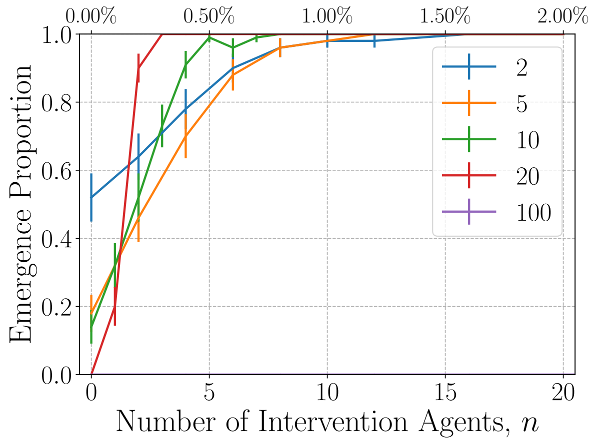

Figure 3.40: The effect of convention space size on initial intervention in the González network. The number of actions in the coordination game is shown in the legend.

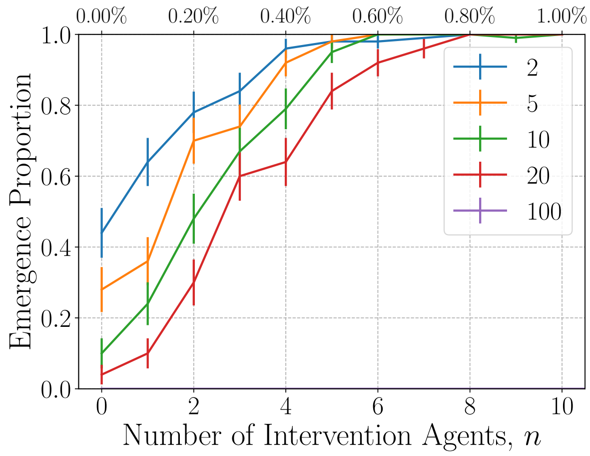

Figure 3.41: The effect of convention space size on initial intervention in the Ichinose TP network. The number of actions in the coordination game is shown in the legend.

that dynamic networks are more sensitive to convention space size than static ones, with a shift from beneficial to detrimental at some point between 20 and 100. Ichinose RR and Ichinose RP exhibit broadly similar behaviour and so aren't shown here.

Ichinose TR and Ichinose TP exhibit differing behaviours however and so Figure 3.41 shows an indicative example of the effects of convention space size for these networks. Unlike González and Ichinose RR and RP, increasing convention space size is actually detrimental in these networks, although the relationship is similar to that observed in small-world topologies where the number of IAs required to guarantee convention emergence is relatively invariant and only the lower proportion behaviours are affected. This is due to the node removal methods in these topologies, with the high-degree nodes constantly changing. Hence the boost in emergence proportion in smaller convention spaces, due to the desired convention emerging more frequently due to random chance, with no or minimal intervention, provides a greater effect here; larger convention sizes mean that the removal of the locations frequently used by IAs results in the "nudge" that IAs can exert, even if these locations are temporary, having less of an effect in directing convention emergence.

Late intervention in the dynamic networks exhibits similar features to initial intervention. Figure 3.42 shows the effect for the González network. Unlike in static networks, there is no clear effect from increasing convention space sizes, with 2, 5 and 10 indistinguishable from one another. A further increase to 20 is actually beneficial requiring slightly fewer IAs to elicit destabilisation, a pattern matching that which occurs for González models in initial intervention. These patterns lend further support to the notion that the network dynamics play a large part in the nature of destabilisation with González in particular benefiting from the extra variance in agent actions that larger convention spaces bring. However, as before, much larger convention spaces are actually detrimental, with IAs unable to cause destabilisation with a convention space of 100 actions. At this scale we posit that the variance in agent choices is detrimental to the

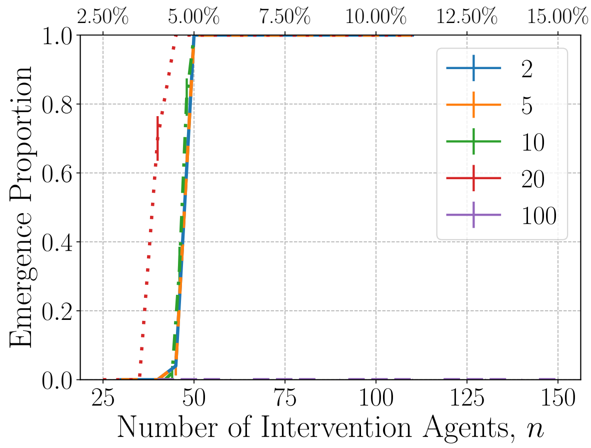

Figure 3.42: The effect of convention space size on late intervention in the González network. The number of actions in the coordination game is shown in the legend.

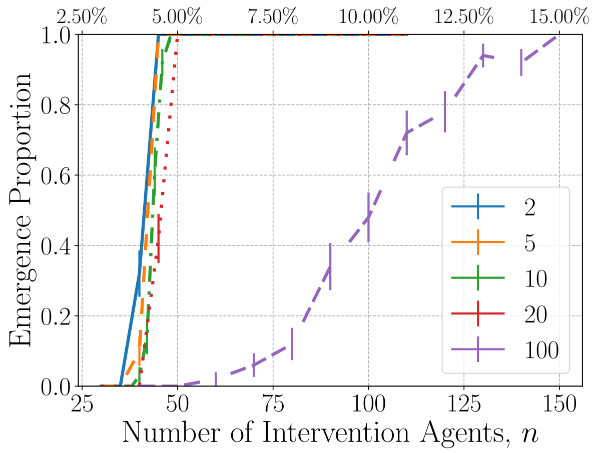

Figure 3.43: The effect of convention space size on initial intervention in the Ichinose RP network. The number of actions in the coordination game is shown in the legend.

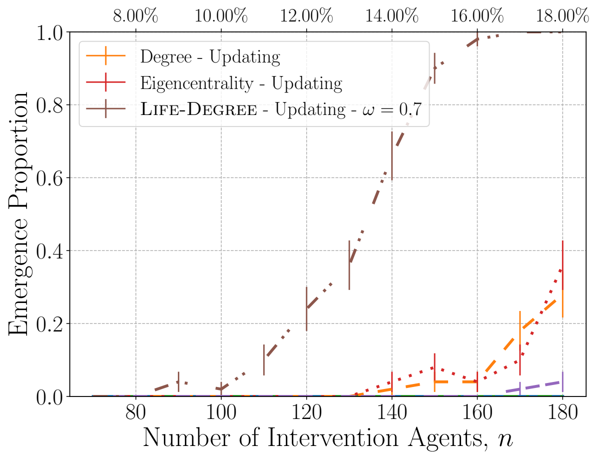

Figure 3.44: The effectiveness of placement metrics in a convention space of size 100 for late intervention in the Ichinose RR network.

efforts of the IAs, making them unable to get other agents to agree with them whereas at size 20 the variance instead works against the dominant convention, making the work of the IAs easier.

Each of the Ichinose models exhibits broadly the same features for late intervention and so we only show Ichinose RP here as a representative sample. Unlike in the González model, increasing convention spaces are never beneficial although the performance difference is minimal requiring only a few additional IAs to overcome. Of more importance is that the Ichinose models are able to consistently cause destabilisation even with 100 possible conventions, in contrast to the González model. Whilst it requires substantially more IAs this again highlights that the network dynamics play a large part, alongside the interaction model, in determining the viability of destabilisation.

Of additional interest is the performance of the other metrics when considering destabilisation with a convention space of size 100 in the Ichinose models. In particular, we find that updating LIFE-DEGREE (with $\omega = 0.7$) outperforms all other metrics in causing this destabilisation. Figure 3.44 shows an indicative example of this in the Ichinose RR model but similar behaviour is observed to varying levels in each of the Ichinose models. This only occurs with the largest convention space but highlights that the longevity of the IAs plays a major part in being able to destabilise the established convention under these conditions.

We have shown that the size of the convention space can have a significant effect on both the emergence and destabilisation capabilities of IAs as well as the natural emergence of conventions within the system. Whilst a further exploration is beyond the scope of this thesis, this highlights the need to consider the domain that the agents are working in when attempting to manipulate convention emergence.

3.10 Alternate Payoffs

We now turn our attention to the effect the payoff matrix has on intervention effectiveness. In particular, we examine whether the positive and negative rewards the agents receive (and the symmetry or asymmetry of these) changes the relationship or relative performance of the various placement metrics when used for IA placement.

This exploration uses 3 different variants of the normal coordination game: $(4, -1)$ (positive reinforcement), $(1, -1)$ (neutral reinforcement) and $(1, -4)$ (negative reinforcement) where the first number represents the payoff for coordinated strategy choice, the second the payoff for conflicting strategy choice. $(4, -1)$ is the payoff structure that has been used in all previous experiments and represents situations where coordination is more beneficial than conflict is harmful, or where coordination is more encouraged. For example, attempting to find a mutual radio channel over which to communicate; whilst there is an expenditure of time for each failure, it is not necessarily very harmful whilst correctly communicating is very beneficial. This structure has been used in previous work [Sen & Airiau, 2007] and has been shown to allow rapid and thorough convention emergence. $(1, -1)$ can instead represent situations where there is symmetry between the benefit and harm, such as choosing which side of a corridor to walk on; there are both minor inconveniences and minor benefits

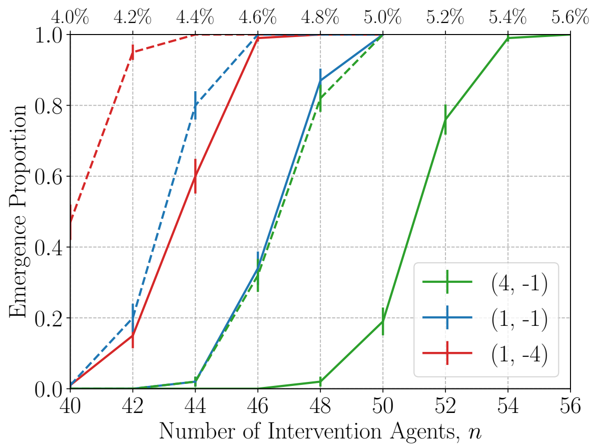

Figure 3.45: Effect on destabilisation of different payoff matrices in the González network.

but neither of a larger scale than the other. Finally, $(1, -4)$ represents situations where conflicting action choices could be very detrimental and should be discouraged rapidly. An example of this is which side of the road to drive on (although this is often described using a symmetric payoff); the negative effects of a crash are substantial.

Figure 3.45 shows the effect of the different payoff matrices during late intervention in the González model. The experimental settings used are the same as in Section 3.7. The figure shows the effects on degree placement of IAs, both static and updating, with updating being shown as dashed lines. The other metrics were found to behave similarly, indicating that this represents a global shift in behaviour, and as such have not been included. As can be seen, the choice of payoff matrix can have a dramatic effect on the number of IAs that are needed to elicit destabilisation with the $(4, -1)$ payoff matrix requiring more than the $(1, -1)$ which in turn requires more than the $(1, -4)$ payoff matrix. We believe this to be due to the fundamental nature of what each payoff matrix is "teaching". With $(4, -1)$ when an agent finds the IAs action through exploration they are rewarded, teaching them this is the "correct" action. However, due to the nature of exploration the chances of them choosing the desired ac-

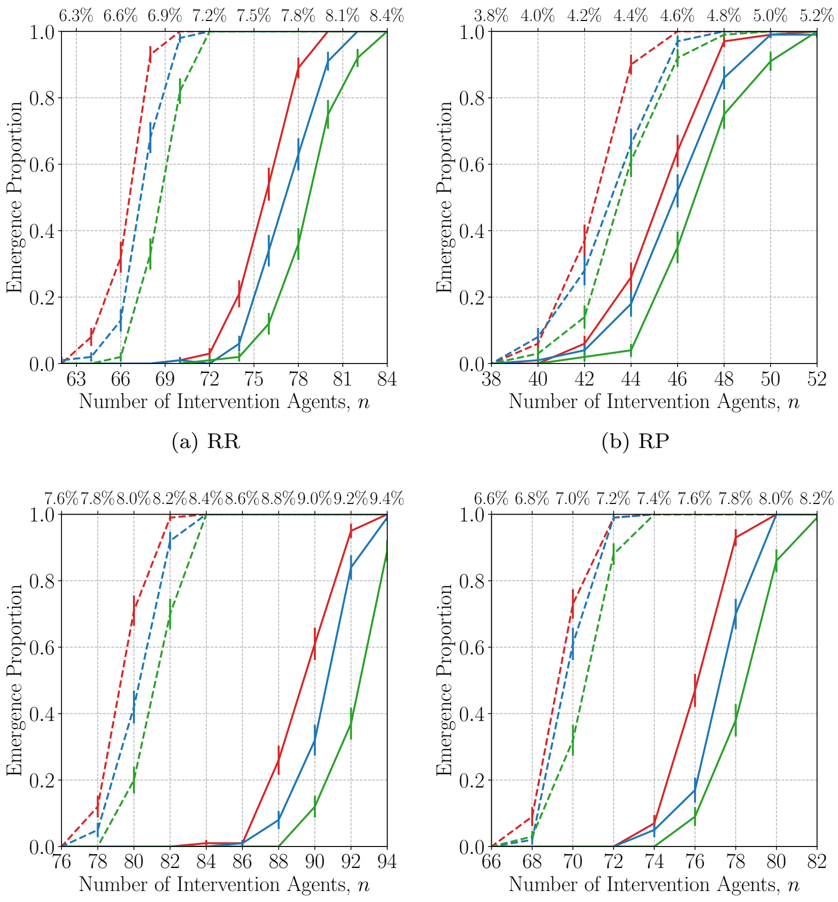

(a) RR (b) RP (c) TR (d) TP

Figure 3.46: Effect on destabilisation of different payoff matrices in the Ichinose networks.

tion is relatively small particularly in the 10-action coordination game as used here. We are reliant on chance to enable the IAs to teach others the correct choice. With $(1, -4)$ by comparison agents are punished each time they make an incorrect choice which will happen far more frequently. As such all other actions except the desired one will have low Q-values meaning that the agent will end up selecting it. In effect, this payoff matrix is teaching agents what *not* to do. The symmetry of $(1, -1)$ represents the middle ground and is the middling result, lending credence to this notion. Thus, when attempting to destabilise systems, those with negative asymmetry will be easier to affect than others.

Figure 3.46 shows the same effect in the Ichinose topologies. Of particular interest is that the effect is nearly identical in all of the 4 models, representing shifts of only a few agents either way, despite the differences between them as noted when exploring late intervention. This indicates that the effect is fundamental to the convention emergence rather than something affected by the nuances of each individual dynamic topology.

We also explored the effect of these payoff matrices in initial interventions in dynamic networks as well as both late and initial interventions in the synthetic networks. However, in each of these settings there was no significant difference in performance, indicating that (i) the effect primarily changes destabilisation efforts as it allows focus on the contrast between the established and desired convention and (ii) that the dynamic nature of the topology is a catalyst for this effect, presumably due to the ability of distant nodes to interact with each other as time changes.

Overall, changing the payoff matrix, either from positive asymmetry to neutral symmetry or negative asymmetry caused no change on the relative effectiveness of the various placement metrics with the same rankings as found in Section 3.7 being found here. However, the absolute performance change, whilst relatively small, highlights the underlying nature of the coordination game as a model for convention emergence when we are looking to destabilise. Whether we are showing agents the correct path or punishing them when they deviate causes slower and faster convention destabilisation respectively.

3.11 Gradient Payoff and Convention Emergence

Finally, we now briefly consider how the gradient coordination game introduced in Chapter 2 affects the dynamics of convention emergence. The graduated nature of the payoffs means that conventions that are similar to one another in terms of their distance in the ordering of action choices will still receive high payoffs. We hypothesise that this will lead to multiple stable subsets of

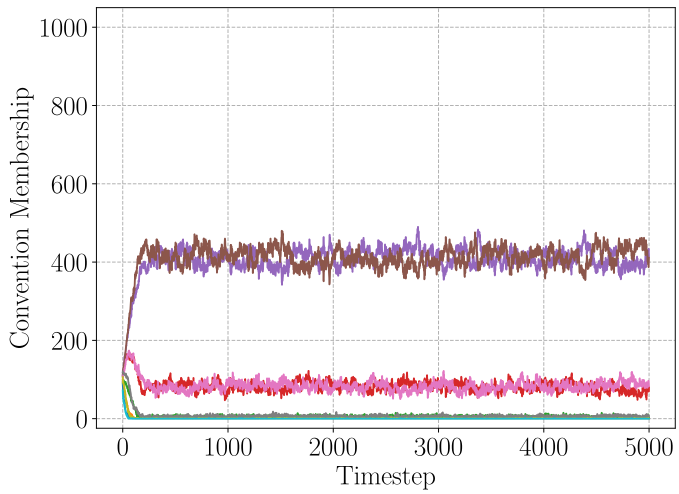

Figure 3.47: Natural convention emergence in the gradient coordination game.

Figure 3.48: Legend for the gradient coordination game simulations.

the population adhering to "similar enough" conventions that allow them to consistently receive high payoffs even though the actions are not identical.

To explore this gradient game we begin by investigating the nature of the convention emergence over time. We utilise a 1000 node scale-free network with otherwise the same settings as before. We use a 10-action gradient coordination game with $pay_{max} = +4$ and $pay_{min} = -1$. However, we find that the results shown generalise to all other networks explored thus far, showing the robustness of the gradient coordination game.

Figure 3.47 shows the nature of convention emergence in the gradient game without the presence of IAs and Figure 3.48 provides a legend for ease. As can be seen, the system very rapidly converges to a stable state where multiple conventions coexist. These exist at 3 primary levels, what we term high, mid and low. The high conventions are those corresponding to strategies 4 and 5 and it is our finding that this is always the case; the system naturally stabilises to have these as the highest conventions. Given their position in the centre of the payoff matrix, this is to be expected as they have the shortest distances between themselves and any other strategy and hence will maximise payoff for any agents choosing them. Interestingly, the stability between these two conventions seems to be intertwined with an almost sinusoidal nature to any changes in membership number in one reflecting a contrasting change in the other before returning to a state of equilibrium. This again is likely due to their location as any agents exploring will naturally tend back to one or other of these strategies as the only stable locations.

The mid level conventions represent strategies 3 and 6, the next strategies on either side of the high-level ones. Again this is to be expected as these are the next best strategies for maximising payoff regardless of interactions with others and the stability of these likely represents agents whose exploration causes them to switch to them. All other strategies (0,1,2,7,8,9) are those we consider "low" and have membership approaching zero. These strategies are those that are most likely to cause clashes in the gradient payoff and hence agents rapidly learn to avoid them. The stability of the two primary conventions, with one not winning out over the other is interesting however and represents a new aspect of the coordination problem as we now have an issue of *cooperation* where it would be maximally beneficial for all agents to use the exact same strategy but exploration and precedence cause them to conflict.

We turn our attention to how IAs might be used to break this deadlock. We consider the insertion of 200 IAs at time $t = 0$ in three different scenarios where the strategy assigned to them is one of the high, mid or low strategies respectively. These different approaches are shown in Figure 3.49. In Figure 3.49a the IAs have been assigned strategy 4 and cause it to increase in size compared to the unaided simulation by more than the number of IAs themselves can account

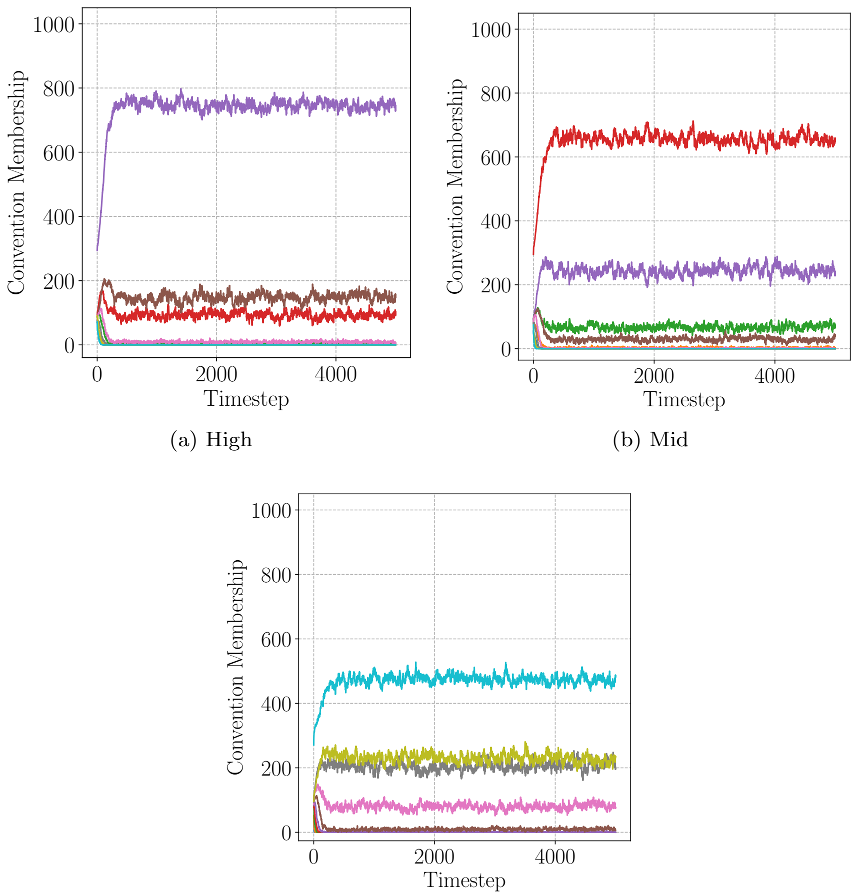

(a) High

(b) Mid

(c) Low

Figure 3.49: Late intervention in the gradient coordination game by promoting the different categories of established convention.

for, indicating that they do influence others. However, the population still stabilises with multiple conventions existing despite the number of IAs being more than sufficient for convention supremacy in all other domains explored in this chapter. The promotion of strategy 4 has decreased the membership size for strategy 6 to close to zero, as the strategy is now further away and hence less beneficial. In effect, the inclusion of IAs has shifted the "centre" of the payoff matrix by artificially increasing the likelihood that the chosen strategy will be used in interactions. This is a useful insight into the nature of IAs that extends beyond this use case.

When promoting mid or low strategies we see two sets of similar effects where the previously high strategies are both reduced in number, down to almost zero in the case of Figure 3.49c. We also see a commensurate rise in other strategies, those that are adjacent to the chosen IA strategy. This effect is most noticeable when utilising a low strategy and results in an entirely different set of conventions. However, in all these cases, the system stabilises with multiple conventions, consisting of these closely related strategies, highlighting again that this can be viewed as a shifting of the centre and artificially changing the equilibria in the payoff matrix. Overall, the gradient game is resistant to emerging a single global convention, instead allowing multiple closely related conventions to coexist.

Figure 3.50 shows the effect of introducing IAs after the conventions have become established. The IAs are inserted at time $t = 500$ and we present the three similar cases again where the IA strategy is from high, mid or low conventions. Whilst it is easier to see the effect to which the increase in one strategy comes at the decrease of another, allowing us to see how the conventions interact, the final levels and types of conventions that exist within the system are almost identical to the case where agents were inserted at time $t = 0$. This indicates that the interaction model created by the gradient coordination game is quite distinct from the pure coordination game with the conventions that emerge dictated strictly by the IAs presence rather than when they are placed

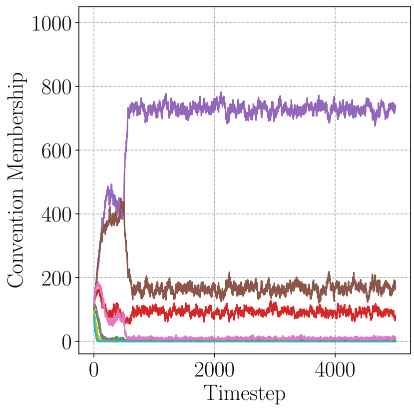

(a) High

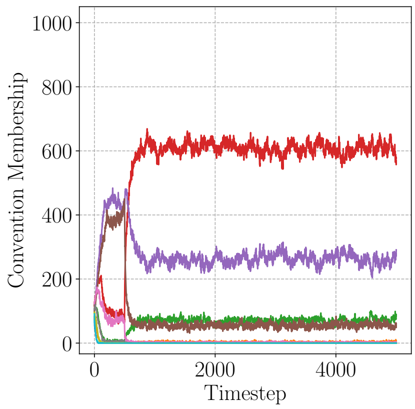

(b) Mid

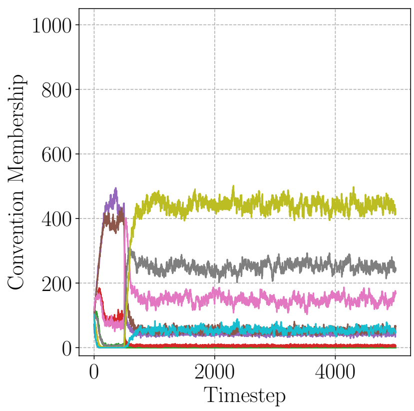

(c) Low

Figure 3.50: Late intervention in the gradient coordination game by promoting the different categories of established convention.

in the system.

The gradient coordination game represents a major shift in the dynamics of the system, actively facilitating the coexistence of multiple conventions and providing robustness and stability even at the cost of global convention emergence. Further exploration of the nature of the conventions within this system and how they may be manipulated is left as future work.

3.12 Conclusions

The use of conventions as lightweight mechanisms to enforce and encourage coordination in multi-agent systems is an ongoing area of research. Having the ability to reduce agent clashes and wasted resources with minimal assumptions about agent capabilities or architecture allows the efficient function of such systems even when agents are controlled by multiple distinct parties. Conventions, as a form of self-imposed social constraints, facilitate this in a decentralised manner by allowing agents to utilise expected behaviour to guide their interactions. The rapid emergence of robust conventions is paramount in reducing the amount of time agents are conflicting with one another. However, the decentralised manner of this emergence means that suboptimal or undesirable conventions might emerge due to their stability and precedence amongst the agents. Being able to remove and replace conventions in this scenario is necessary for their use as a general problem solving tool.

In this chapter we have explored the nature of convention emergence in MAS. We have shown that small numbers of fixed strategy *Intervention Agents (IAs)* can be used in a variety of topologies to facilitate convergence of the agent population to a desired convention. We have explored the use of a number of metrics to inform *where* these agents might be best placed within these topologies to maximise their effectiveness. When placed within a system at the beginning of a simulation these *initial interventions* are shown to be highly effective. A summary of the contributions and results is shown in Table 3.3 and these are

Result	**Section**
Degree placement is best for initial intervention in static networks and $<$ 1% of population as IAs is needed to ensure convention emergence.	3.4
Degree or eigencentrality placement is best for initial intervention in dynamic networks and $<$ 1% of population as IAs is needed to ensure convention emergence. No significant difference between Static or Updating metrics.	3.5
Consideration of agent longevity is not as important as ability to influence more individuals. LIFE-DEGREE shows that high degree is the more important aspect.	3.5
Degree and eigencentrality substantially outperform other metrics for destabilisation in static networks. $<$ 8% of the population as IAs guarantees destabilisation.	3.6
Degree and eigencentrality placement are best for causing destabilisation in dynamic networks and the Updating versions of these metrics are significantly better than the Static versions. $<$ 8.5% of population needed to destabilise in all examined topologies.	3.7
Passive destabilisation is best effected by degree or eigencentrality in both static and dynamic networks and requires 2.5 times as many IAs as aggressive destabilisation.	3.8
Using a Gradient Payoff matrix allows creation of a stable system where multiple conventions exist simultaneously. This system is resilient to external influence with IAs able to shift the effective equilibrium of the system.	3.11

Table 3.3: A summary of the major results and contributions from this chapter.

discussed in more detail below.

We began by considering initial intervention in static topologies, both synthetic and real-world. We showed that degree or eigencentrality-based placement offers the most effectiveness at encouraging convention emergence and that, across a range of network topologies and sizes, utilising substantially less than 1% of the population as IAs was always enough to effect the desired convention emergence 100% of the time. We explored the effect on convention emergence of a number of different features of the underlying topologies and showed that, whilst they may change the level of intervention necessary, the use of IAs is able to cause robust convention emergence within them. We extended previous work in the literature by considering a range of placement metrics and showing the effectiveness of these. Our work here showed that IAs could be used to rapidly (within a few hundred timesteps) produce high-quality conventions that did not fluctuate once they had become established. Chapter 5 takes this concept further and explores the stability of the established conventions when IAs are removed.

We additionally extended this study into dynamic networks, something not actively considered in the literature, where the nodes and edges of the topology are able to change over time. We introduced the notion of different applications of the traditional metrics in dynamic topologies in two forms: static and updating and showed the efficacy of these approaches in causing initial intervention in these dynamic networks. We introduced a new metric Life-Degree which utilises information unique to dynamic topologies and allowed us to study the importance of agent longevity in the selection process for IAs. The results of Life-Degree showed that, nearly universally, consideration of agent longevity was detrimental to the effectiveness of IAs. Instead, efforts should focus on maximising the agent's ability to influence others, even over short time frames. We showed that convention emergence was indeed possible in these networks and studied the nature of convention emergence within them, establishing that the required levels of interventions were higher here than in static networks but that our application of IAs is still able to facilitate it.

Having shown that convention emergence is possible in both static and dynamic topologies, we then introduced the notion of *destabilisation*; using IAs to remove and replace an already established convention. In static networks we showed that degree and eigencentrality substantially outperform HEE and HITS as methods for selecting optimal locations to maximise destabilisation potential. We showed that, across all static topologies observed, no more than 8% of the population have to be utilised as IAs in order to guarantee the removal and replacement of the dominant convention despite it being used and established amongst >90% of the population.

In dynamic topologies we showed that our updating method of placing IAs substantially outperforms the naive static approach, allowing destabilisation with far fewer IAs. We found that, whilst more IAs were needed than in equivalently-sized static networks, the number required still represents a small portion of the total population but can guarantee destabilisation in dynamic topologies as well as static.

We next introduced the notion of *Passive Destabilisation*, using IAs to remove a target convention but without specifically replacing it with another. We showed the efficacy of our previous approaches in causing this type of destabilisation in both static and dynamic topologies and found that, consistently, ~2.5 times as many IAs were needed to elicit passive destabilisation.

We then explored the effect of different sizes of convention space on both initial convention emergence and destabilisation across static and dynamic topologies. We showed that, dependent on topology, the convention space could have a sizeable effect on the nature of conventions and the ability to destabilise them. We followed this with an exploration on how the symmetry or asymmetry of the payoff matrix in the coordination game could affect manipulation of conventions highlighting the difference between positive and negative reinforcement and the impact this has on destabilisation efforts.

Finally, we introduced the notion of a gradient coordination game and showed the distinctive properties that this introduces to the nature of convention emergence, showing that the presence of IAs in the system can effectively shift the equilibrium of desirable actions.

Overall, we have expanded the state-of-the-art in using fixed strategy agents to elicit convention emergence in MAS. We have shown that destabilisation of existing conventions is possible and explored some of the criteria that affects the efficacy of this. In the following chapters we consider additional constraints on the use of IAs to elicit these changes and gain a deeper understanding of how they affect conventions within the system.

CHAPTER 4

Interventions Under Partial Observability

In the previous chapter we have shown that both initial and late convention emergence can be manipulated using Intervention Agents (IAs) to encourage and direct the convergence of agent choices and destabilisation of existing conventions. We showed that placing IAs in topologically influential locations aided their ability to do this and increases their efficacy for these purposes. In this chapter, we investigate how this knowledge can be utilised when topological information is restricted and limited to observations of a restricted number of individual nodes in the network. We propose two algorithms to address these restrictions and still find topologically influential locations. We analyse the performance and limitations of these algorithms and show their efficacy in convention emergence and destabilisation. To do this we make use of the graphs and generators discussed in Chapter 2 as well as the agent interaction model introduced in Chapter 3.

4.1 Introduction

In many modern instantiations of multi-agent systems (MAS) the agents are typically constrained in their ability to interact with one another by an underlying network topology. These connecting networks limit and shape the system dynamics and change how agents at the macro level interact with each other [Delgado et al., 2003; Griffiths & Anand, 2012; Sen & Airiau, 2007]. In particular, when trying to coordinate the actions of independent agents, these limitations affect any potential homogeneity of agent influence and elevate some to positions of higher influence [Easley & Kleinberg, 2010; Franks et al., 2013; Villatoro et al., 2009].

Previous work on convention emergence often assumes that the topology constraining agent interactions is fully observable, allowing highly influential locations to be found easily [Kittock, 1995; Salazar et al., 2010; Sen & Airiau, 2007; Villatoro et al., 2009]. However, in many real-world applications such information is not always readily available. This can be due to factors such as the problem size or external limitations such as restricted access to network information or a network's API as is the case with Twitter [Twitter Developers, 2017a] or Facebook [Facebook Developers, 2017b].

In this chapter we explore the effect of the restrictions placed on Intervention Agent placement in partially observable topologies. We propose an algorithm, PO-Place, to find influential locations within such static topologies given a highly limited number of network queries and another, DynaPO, for dynamic topologies. We show the effectiveness of the algorithms at finding approximations of the highest degree locations for real-world, synthetic and dynamic topologies under a number of restrictions on available information. We then apply the algorithms to select IAs within these networks and examine the effect on convention emergence compared to placing with full topological knowledge. This approach allows an interested third party, with limited access to the system, to find the appropriate locations to target their influence efforts.

The remainder of this chapter is arranged as follows. Section 4.2 explores the background and related work in the literature regarding finding influential nodes with limited or local information. In Section 4.3 we introduce PO-Place and analyse its performance in several real-world networks. Then, in Section 4.4, we use PO-Place as a placement mechanism for initial intervention convention emergence in these real-world networks before extending and expanding our analysis into synthetic topologies. Section 4.5 explains the design behind DynaPO and then analyses its effectiveness at finding influential locations before investigating its performance in encouraging convention emergence in dynamic topologies. In Section 4.6 we turn our attention to using PO-Place to cause destabilisation of existing conventions in static topologies and in Section 4.7 we do the same for DynaPO and dynamic topologies. Finally, in Section 4.8 we present our conclusions and considerations for future work.

4.2 Background

As we established in the previous chapter, convention emergence can be facilitated by placing Intervention Agents (IAs) at highly influential locations in the underlying network topology. However, previous work assumes full visibility of the network topology to inform this placement. Indeed, little work on partial observability for convention emergence has been done. Related work exists in the fields of graph algorithms and influence spread, the latter sharing many qualities with convention emergence. For instance, Brautbar and Kearns present a novel model [Brautbar & Kearns, 2010], *Jump and Crawl*, motivated by operations commonly available in networks such as Facebook. Their model consists of two aspects: *Jump* which moves to a randomly selected node in the network and *Crawl* which searches all neighbours of the selected node for high-degree nodes. They provide bounds for many different types of network but, for an arbitrary network, finding the guaranteed highest degree node approaches $O(n \log n)$, a large factor for even medium-sized networks. In full, they show that $O(n^{\beta} \log n)$ steps finds an $O(n^{1-\beta})$ multiplicative approximation of the highest degree which is generally too inaccurate for our purposes and scale without simply observing the whole graph.

Borgs *et al.* [Borgs et al., 2012a; Borgs et al., 2012b] propose a polylogarithmic algorithm for finding the root node (the initial node) in Preferential Attachment topologies, a quality which is present in many real-world social networks. The algorithm runs in $O(\log^4 n)$ and works by continuously selecting the highest degree node from a growing fringe. However, the constant factors associated with it are non-trivial for many network sizes. Additionally, the algorithm outputs a set S of size $O(\log^4 n)$ which contains the root node with high probability. This means that most nodes in $N(S)$ must have been examined to know which to select each iteration. For our purposes this means they all must have been observed, something Borgs et al. does not allude to. Finally, when used to find high degree nodes, their algorithm is only proven to return a node of degree at least $1/log^2(n)$ of the maximum degree in the graph in time $O(\log^4 n)$, a poor guarantee for our use case.

Building on this, [Avrachenkov et al., 2014] propose a simple algorithm for rapidly finding a selection of high-degree nodes in directed graphs using only local information. Their approach works by treating the graph as effectively bipartite and using a two-stage approach to find likely candidates in one partition by randomly sampling from the other and noting which nodes have the highest number of in-edges from the sample. They then choose the top candidates and find their actual degree. Importantly theirs is some of the earliest and only work that treats the number of steps as finite, explicitly calling out API limitations as the justification. However their work is only applicable to directed graphs and relies heavily on the graph density.

The influence maximisation problem [Chen et al., 2014; Chen et al., 2009] attempts to find a selection of nodes such that the spread of influence (often modelled as single chance 'cascades') from them is maximised. As such, this is intrinsically tied with finding influential locations. Zhuang et al. [2013] propose a "growing fringe" method that works in both dynamic and partially observable topologies to maximise the influence spread. It works by constantly adding from the set of fringe nodes the one that will increase the potential spreading area the most. The performance of this algorithm lends further credence to the idea that probing and expanding node sets is an applicable method for finding usable nodes in both static and dynamic topologies.

As in this chapter, Mihara *et al.* [Mihara et al., 2015] assume the network is initially unknown and show that influence maximisation effectiveness of 60-90% with 1-10% network observation is achievable. This work also uses a 'growing fringe' approach with priority based on degree estimation. As influence maximisation and convention emergence are similar in aim, this indicates that results are achievable under partial observability constraints.

Whilst many of these approaches are similar in application, they differ in that our investigation focuses on the often encountered scenario of limited, finite observations. Making optimal use of these is paramount and so necessitates a different set of considerations. The work here differs from that of Avrachenkov et al. [2014] in that we consider undirected networks, as well as dynamic networks, rather than only directed, bipartite networks. Additionally, we consider the application of finding influential nodes under partial observability to the field of convention emergence rather than for purely graph theoretic purposes.

4.3 PO-Place Algorithm

For this thesis, the partial observability problem for networks can be described as any scenario where a network's topology is initially unknown and is revealed incrementally within a local neighbourhood of nodes already explored [Borgs et al., 2012b]. As a solution to the partial observability problem for Intervention Agent selection we propose a heuristic algorithm, PO-Place. This section describes the function of the algorithm as well as the justification for the design choices.

The placement strategy is presented in Algorithms 1 and 2 and has the following aim: Given a network, $G = (V, E)$, a desired number of locations, n, and a limited number of observations, o, find a selection of nodes $S = \{v_1, ..., v_n\} \subset V$ such that

$$\textit{deg-sum(S)} = \sum_{v \in S} \textit{deg}(v) \tag{4.1}$$

is maximised. We define an observation as a query that retrieves the list of neighbours, $N(u)$ for a given node, u. This functionality is frequently available in real-world network APIs (such as Twitter [Twitter Developers, 2017b] or Facebook [Facebook Developers, 2017a]) and so we assume that such information is available. This assumption is later relaxed to allow the algorithm to explore situations with only limited neighbour information. We assume that the set of nodes, V, is known but the set of edges, E, (and hence neighbours and degree of a node) is not. Finding the highest degree nodes is desirable since IA placement by degree consistently produces effective convention emergence [Franks et al., 2013; Griffiths & Anand, 2012; Marchant et al., 2015a; Marchant et al., 2015b] but without requiring computationally expensive metrics such as betweenness centrality. The degree of nodes can be entirely derived from local information and, as such, is an applicable heuristic within partially observable networks.

Other heuristics and approximation algorithms exist for computing various graph metrics using only local information. For instance, Andersen et al. [2007] provide a technique for calculating an approximation of PageRank for a given node (itself a variant of Eigenvector Centrality) using only the local nodes around the target location. It requires $O(1/\epsilon)$ observations to produce an ϵ-approximation of the PageRank score and does so in a manner that grows out the local area, which is the approach taken below and by many similar approaches in the literature [Brautbar & Kearns, 2010; Maiya & Berger-Wolf, 2010; Mihara et al., 2015]. However, using this approach would produce a set of scores that are good approximations for those at the centre of the growing area but very poor approximations for those at the fringes who have not had their neighbours examined yet. This is likely to lead to ignoring these fringe nodes despite them having potentially having better scores as they must wait for their neighbours to be examined to provide a good approximation. In comparison, using the degree of the node provides accurate information for all nodes examined regardless of their position within the area being searched. Indeed all approaches that rely on local area knowledge beyond the immediate node will exhibit this problem including simple and entirely computable metrics such as Average Neighbour Degree [Franks et al., 2013; Mislove et al., 2007], Local Clustering Coefficient [Watts & Strogatz, 1998] and Edge Embeddedness [Easley & Kleinberg, 2010]. To address this problem would require only selecting nodes from *within* the fringe itself (effectively creating two fringes of

Algorithm 1 Partial Observability: Placement

Symbols: $G = (V, E)$, a graph of V vertices and E edges; n, the number of locations to find; o, the number of observations; s, the number of concurrent starting locations; p, the proportion of neighbours that can be retrieved in an observation; f, the number of nodes to expand simultaneously; o_{rem}, the number of observations remaining; o_{local}, the number of observations assigned to a starting location.

procedure PO-PLACE($G = (V, E)$, n, o, s, p, f)
    Create empty node set, S
    Create empty mapping, N
    $o_{rem} \leftarrow o$

    while $o_{rem} > 0 \wedge |S| < |V|$ **do**
        Select v uniformly at random from $\{V \setminus S\}$
        if $o_{rem} \mod s \neq 0$ **then**
            $o_{local} \leftarrow min(\lceil o/s \rceil, o_{rem})$
        else
            $o_{local} \leftarrow min(\lfloor o/s \rfloor, o_{rem})$
        $o_{rem} \leftarrow o_{rem} - o_{local}$
        $o_{unused} \leftarrow Traverse(G, o_{local}, v, p, f, S, N)$
        $o_{rem} \leftarrow o_{rem} + o_{unused}$

    return n highest-degree nodes in S

exploration) which, given the highly limited nature of the number of observations would reduce the potential nodes substantially. Because of this (and the fact that the potential gains from using other metrics is small at best as seen in Chapter 3) we focus on and utilise only node degree for the purposes of choosing influential nodes under partial observability.

The algorithm begins by creating an empty set, S, to monitor which nodes have already been explored and an empty mapping, N, that maps a node v to $N(v)$, its set of neighbours. By storing this information we can avoid using observations redundantly but this approach will need to change when considering situations where the neighbours of a node may change over time (see Section 4.5).

Many of the other approaches [Borgs et al., 2012b; Maiya & Berger-Wolf, 2010; Mihara et al., 2015] to finding high-degree nodes select a random starting node and then 'grow' outwards, selecting the highest degree nodes from the neighbourhood surrounding those already explored. However, this is not desirable in IA placement since, with limited observations, it is likely to produce a single cluster of well-explored nodes. Selecting from this cluster will then mean that all IAs are close together, making some of their influence redundant. Instead, we build on the notion of *Jump and Crawl* [Brautbar & Kearns, 2010]. We incrementally explore the local area around a given seed node (a 'crawl'), by making graph queries and adding the explored nodes to the set of already queried nodes and then repeating this process using the neighbourhood around this set. We do this for a defined amount of queries, and then 'jump' to another, unexplored, location and exploring around this new point.

The use of 'jumps' helps to minimise the risk of overlap between high-degree nodes, as well as ensuring that a bad initial random selection does not hinder the final selection by counteracting the effect of local maxima. To facilitate this, we introduce a parameter, s, which dictates the minimum number of separate local area explorations that will take place. The observations are split, as evenly as possible, between each of these explorations with the earlier ones receiving any spare observations (this is achieved between Lines 7 and 10 of Algorithm 1). This subset of observations is then passed to the local area traversal which is presented in Algorithm 2. If any observations are unused by the local area traversal (for instance if it finds a local maxima) they are returned to the pool of available observations and used in later, additional local traversals.

Algorithm 2, Traverse, describes the local area traversals. It is aware of both S and N, to avoid redundant exploration, as well as the initial start node of the local area, v. It is also passed its own local limit of observations and two parameters from outside, p and f, which are explained below. It maintains a max-priority queue to determine which node(s) it should next explore by highest degree and begins by adding v to this queue. Continuously choosing the unexplored node with the highest degree (Line 10) allows exploration up the gradient of higher degree nodes and has shown effectiveness in the work of Borgs et al. Throughout Algorithm 2, observation of a node's neighbour list is stored in N to avoid additional queries. The algorithm then performs the following, until either the queue is empty or all assigned observations have been

Algorithm 2 Partial Observability: Local Area Traversal

Symbols: $G = (V, E)$, a graph of V vertices and E edges; o, the number of observations available; v, the starting node in the graph; p, the proportion of neighbours that can be retrieved in an observation; f, the number of nodes to expand simultaneously; S, the set of already explored nodes; N, a map of nodes to their list of neighbours.

1: **procedure** TRAVERSE($G = (V, E)$, o, v, p, f, S, N)
2:     Create max-priority queue, Q
3:     $count \leftarrow 0$
4:     **if** v **not in** N **then**
5:         $N[v] \leftarrow N(v)$
6:         Add v to S
7:         $count \leftarrow count + 1$
8:     Add $(v, |N[v]|)$ to Q

9:     **while** $|Q| > 0 \wedge count < o$ **do**
10:         $Fringe \leftarrow$ top $min(f, |Q|)$ elements of Q
11:         **for all** u in $Fringe$ **do**
12:             $Avail \leftarrow \{N[u] \setminus S\}$
13:             $num \leftarrow min(|Avail|, max(f, \lfloor p \times |Avail| \rfloor))$
14:             $Chosen \leftarrow$ uniformly at random select num members of $Avail$
15:             **for all** w in $Chosen$ **do**
16:                 $N[w] \leftarrow N(w)$
17:                 Add w to S
18:                 $count \leftarrow count + 1$
19:                 **if** $count = o$ **then**
20:                     **return** 0
21:                 Add $(w, |N[w]|)$ to Q

22:     **return** $o - count$

used up:

1. Take the top f (highest-degree) nodes from the queue (or all elements, if fewer). [Line 10]
2. For each of these nodes, find the set of unexplored nodes in its neighbours. [Line 12]
3. Choose a proportion, p, of these (or up to f if this proportion would be less than f) uniformly at random. [Lines 13 and 14]
4. Add these nodes to the queue after finding their neighbours. [Line 15 to Line 21]

Parameter f is the 'fringe size', the number of nodes that are expanded

	Network		Largest WCC	
	$\lvert V\rvert$	$\lvert E\rvert$	$\lvert V\rvert$	$\lvert E\rvert$
CA-CondMat	23,133	93,497	21,363	91,286
Enron-Email	36,692	183,831	33,696	180,811
Twitter	81,306	1,768,149	81,306	1,342,296

Table 4.1: Original and Largest Weakly-Connected Component Network Sizes

simultaneously before their neighbours are queued. This acts as a control over how ‘breadth-first’ or ‘depth-first’ the local traversal approach will be. This allows the algorithm to be avoid the pure depth-first approach of the likes of [Borgs et al., 2012b] which may lead to undesirable local optima faster.

Parameter p is the proportion of the node’s neighbours that should be queried. This allows the algorithm to simulate situations where a node’s full neighbour list is either not fully available (for instance, an API that only returns a subset) or where doing so incurs additional cost. In the latter case we seek to explore the effect that only querying p proportion of neighbours has on the performance of PO-Place. Whilst it is hypothesised that it will reduce the effectiveness, establishing the extent of this reduction, and whether the results are still close enough to degree placement, allows PO-Place to be effective over a wider range of scenarios.

The algorithmic complexity of PO-Place is bounded by the number of observations it can make. Each expensive query is accompanied by the use of an observation. Assuming efficient graph representation this means that the priority queue is the primary bound resulting in the complexity of PO-Place being $O(o \log o)$.

4.3.1 Networks

We performed simulations of PO-Place on the real-world network previously used in Chapter 3. As before, we concern ourselves only with the largest weakly-connected components of the networks and the sizes of these are shown again in

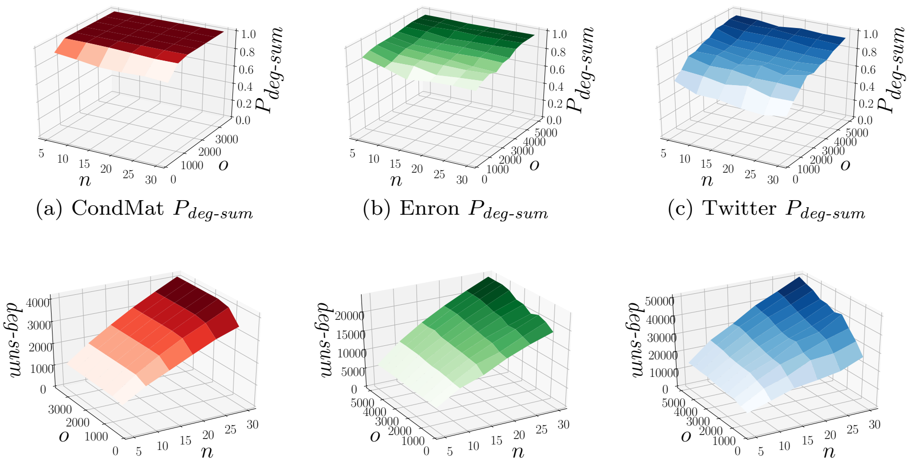

(a) CondMat $P_{deg\text{-}sum}$ (b) Enron $P_{deg\text{-}sum}$ (c) Twitter $P_{deg\text{-}sum}$

(d) CondMat *deg-sum* (e) Enron *deg-sum* (f) Twitter *deg-sum*

Figure 4.1: $P_{deg\text{-}sum}$ and *deg-sum* performance of PO-Place for varying n (# of locations) and o (# of observations) in the real-world networks.

Table 4.1 for ease. We varied both the number of nodes ($n = 5$ to $n = 30$) being requested as well as the number of observations provided ($o = 500$ to $o = 5000$ [$o = 3500$ for CondMat]). To establish an upper bound and allow comparison a full-observability degree placement was also performed for each of the networks with the same range of values. Each set of parameters was averaged over 30 runs.

4.3.2 PO-Place Performance

In this section we present the analysis of PO-Place and compare it to the upper bound from degree placement. We explore the effects of the various parameters on PO-Place at different levels of observation. We then use these findings as insight to compare the performance of PO-Place to degree for convention emergence when used to place FS agents into the chosen networks.

We begin by looking at the isolated algorithm output, comparing it to the output generated by a degree placement scheme. As the aim of PO-Place is to maximise *deg-sum* (Equation (4.1)) this is our primary metric by which to evaluate PO-Place. The highest *deg-sum* possible in each network is that of the set of highest degree nodes. Establishing this as an upper bound allows evaluating the performance of PO-Place by comparing the *deg-sum* of its output as a proportion of that of the pure degree network. We denote this as $P_{deg\text{-}sum}$ and note that a $P_{deg\text{-}sum}$ of 1 means that the *deg-sum* is the same as that found under full observability.

Whilst *deg-sum* describes the maximum reach of the nodes selected, another useful metric is the size of the 1-hop neighbourhood of those nodes. This can be defined as:

$$\text{1-Hop}(S, G) = \{v \in \{V \setminus S\} | \exists (u, v) \in E \wedge u \in S\} \tag{4.2}$$

where S is the set of nodes selected for placement and $G = (V, E)$ is the network. That is, the 1-Hop neighbourhood is the set of nodes that are connected to a member of S but are not in S themselves. The 1-Hop neighbourhood offers a slightly different measure of influence by discounting nodes that are connected to multiple members of S. Whilst normally tied closely to *deg-sum* a noticeable disparity indicates that the selected nodes are likely to be clustered close to one another, which is undesirable. As with *deg-sum* we concern ourselves with the proportionate behaviour of 1-Hop size, $P_{|\text{1-Hop}|}$. Again we note that a $P_{|\text{1-Hop}|}$ of 1 means that PO-Place has found a set of nodes with the same 1-Hop value as the nodes selected under full observability.

The final metric we use to evaluate the performance is based on the Jaccard Index which measures similarity between two sets [Jaccard, 1912]. The Jaccard Index is defined as $J(A, B) = |A \cap B|/|A \cup B|$. However, in our instance, one of the sets is static. We are trying to approximate that set with the other (i.e. a one-way similarity), whilst the Jaccard Index is looking at the two-way similarity between them. Instead we want to measure how close the selection of PO-Place is to the baseline, and so we define a distance measure, D_{Base}, thus:

$$D_{Base}(S, Base) = |S \cap Base|/|Base| \tag{4.3}$$

That is, the fraction that elements of S (the set of nodes selected for placement) make up of the baseline set, *Base*. If PO-Place selects the same nodes as placement under full observability then D_{Base} will have a value of 1. This metric enables evaluation of how close the actual node selection of PO-Place is to that of degree placement, whilst the previous two measure the selection's features. However, it should be noted that a poor D_{Base} score is not necessarily indicative of poor performance. In graphs with many nodes with the same degree we can expect that, even if PO-Place has found nodes of equal or close *deg-sum* that may not be the same sets that have been selected by degree placement. Indeed, multiple degree placements may select different sets of nodes depending on implementation. Because of this, a high D_{Base} score is indicative of good PO-Place performance but a bad D_{Base} score is not indicative of poor performance.

These metrics offer insight into the influence and reach of the nodes selected by PO-Place as well as allowing a direct comparison to degree-based placement with full observability. Thus they should be good predictors of the performance of PO-Place in the convention emergence setting.

Varying Observations

We begin by considering the base case of the algorithm where $s = p = f = 1$. This allows us to study the effect of varying the number of observations and provides a lower bound on the expected performance of PO-Place. With these settings, PO-Place closely resembles the algorithms presented by [Borgs et al., 2012b; Mihara et al., 2015] in how it expands the search area.

We examine the effects of varying both the number of observations available (o) as well as the number of locations requested (n) in all three networks. For all networks, n was varied between 5 and 30 in increments of 5 and o was varied from 500 observations up to 3500 (for CondMat) or 5000 (for Enron and Twitter). The results are presented in Figure 4.1.

As can be seen in Figure 4.1, all networks respond well, even with minimal numbers of observations. Even at $o = 500$, the degree sum of the nodes selected by PO-Place is often a substantial proportion of the optimal, full observability one. The performance varies across the three networks, with placement in CondMat doing best where it varies from 90% (±5%) at $n = 5$ to 83% (±5%) at $n = 30$. The algorithm similarly performs well in Enron, though to a lesser extent. The performance in Twitter is noticeably worse, varying from 61% to 48% with larger standard deviations for both. This is to be expected, as 500 observations represents a substantially smaller proportion of the population in Twitter than it does in CondMat or Enron (0.61%, 2.34% and 1.48% respectively). Even with this, the percentage achieved in Twitter with such limitations substantially outperforms the naïve solution of using all observations at random locations (16% (±6%) for $n = 5$, $o = 500$, averaged over 100 runs).

Performance rapidly increases with the number of observations. For $n = 30$, the worst performing value of n, in both CondMat and Twitter $P_{deg\text{-}sum}$ exceeds 90% at around 5% network observation ($o = 1000$ for CondMat and $o = 5000$ for Twitter) and Enron exceeds 90% at around 10% observation ($o = 3500$). Figure 4.1 also shows that the relationship between $P_{deg\text{-}sum}$ and increasing o is one of diminishing returns, with improvements in $P_{deg\text{-}sum}$ most noticeable at lower values of o. This is to be expected, the relative increase in o is smaller at higher values, but dictates that increasing the effectiveness of PO-Place at low values of o will have the most benefit. Additionally, in each network, the difference in performance across the values of n becomes less noticeable at higher o. Thus, any increased performance from PO-Place will be most noticeable early on.

The other metrics we use to evaluate PO-Place show similar behaviour to $P_{deg\text{-}sum}$, increasing rapidly with the number of observations. Figure 4.2 shows a representative example of the three metrics' variation with o for the Twitter network when requesting 20 locations. The shaded regions represent the standard deviations. As can be seen, both the *deg-sum* and 1-Hop proportions increase rapidly up until $o = 2000$ and then any further gains occur over larger spans of increases in the number of observations. The standard deviations for each of these decrease as well, from approximately 15% at $o = 500$ down to around 5% at $o = 5000$. This indicates that, not only is PO-PLACE finding sets of nodes with higher degree, it is doing so consistently at higher numbers of observations, a finding that is repeated across all networks and values of n. $P_{|\text{1-HOP}|}$ is consistently at the same level, if not better than, $P_{deg\text{-}sum}$. Whilst it was expected that the two should be well-correlated, this shows that PO-PLACE is not simply choosing nodes close to one another and, indeed, is often choosing nodes that have a better neighbourhood size than the *deg-sum* would indicate.

The performance of PO-PLACE when evaluated by D_{Base} is noticeably different than the other two metrics and offers an interesting insight. The same pattern of diminishing returns is not present and D_{Base} continues to increase with additional observations in the range investigated. Note that, although both the degree sum and neighbourhood size are comparable to that of pure degree placement, the low values of D_{Base} indicate that the nodes selected are not the same as the actual highest degree nodes. The results in this chapter evaluate whether this difference has a noticeable effect on convention emergence or if the reach and influence indicated by high *deg-sum* and 1-HOP scores is the best indicator of success as hypothesised.

Overall, these initial results show that PO-PLACE performs nearly as well as the full observability case whilst observing only 5-10% of the network. In comparison to other applicable approaches, it also performs well. The approach of Borgs et al. [2012b] is difficult to quantify as they do not include the necessary observations of nodes within their fringe in their calculation. Thus, the $O(\log^4 n)$ nodes in their output set must be increased to account for all other nodes observed but not included in the output set. Assuming average degree for all nodes in the output set the fringe could be substantially larger. Whilst there would likely be some overlap in neighbours, this gives a rough approximation. Using this, the approach of Borgs et al. would need to observe between 14-24% of the examined networks and would only be guaranteed of finding a

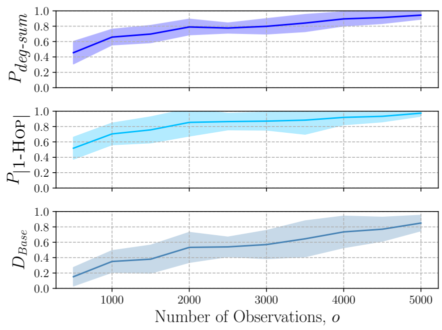

Figure 4.2: Metric performances of PO-Place for the Twitter network, $n = 20$.

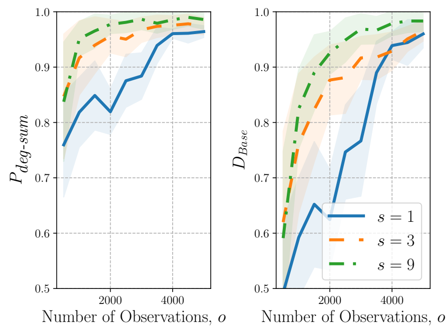

Figure 4.3: Effect of varying s on $P_{deg\text{-}sum}$ and D_{Base}. Enron network, $n = 30$.

set of nodes which provides an $O(0.01)$ approximation of $P_{deg\text{-}sum}$ [Borgs et al., 2012b]. The approach of Brautbar & Kearns [2010] is substantially worse. In order to only require the same number of observations, the multiplicative approximation for high degree nodes would be exceptionally poor. The approach of Mihara et al. [2015], whilst different in desired outcome, is closest here but PO-Place outperforms even this with all networks achieving 90+% by 10% observation compared to 60-90% for 1-10% observation for Mihara et al. Whether these results can be applied to the benefit of convention emergence is explored later in this chapter.

Varying Concurrent Searches

Having established a baseline for PO-Place and explored the effects of limited observations we now explore the variants of the algorithm. As noted in the previous section, at low values of o the $P_{deg\text{-}sum}$ performance of PO-Place is consistently lower, with performance in the Twitter network as low as 48%. With very limited observations, making the best use of them is paramount. At the beginning of this section we hypothesised that splitting the available observations between multiple locations in the network and exploring them in parallel may offer improvements over crawling from a singular location.

To test this hypothesis, we varied s from 1 to 9 to determine the effect that these concurrent searches would have. Figure 4.3 shows a typical case in the Enron network for $n = 30$. Shaded areas represent the errors of each plot. The left-hand graph shows the effect on $P_{deg\text{-}sum}$ of varying the number of concurrent searches, splitting the observations between them. As can be seen, adding concurrent starting points has an immediate and noticeable effect, especially at low numbers of observations. At $o = 500$ the proportion achieved by $P_{deg\text{-}sum}$ is 10% higher when additional starting locations are introduced and this difference becomes even more noticeable as o increases. Indeed, for most values of o, adding additional starting locations had significant benefits in both the Enron and Twitter networks, with the benefits becoming less marked at high o where $P_{deg\text{-}sum}$ approaches 1.0 unaided. Whilst there is a noticeable drop-off in effectiveness after initial parallelisation ($s = 5$ and $s = 7$, not included in the results to aid readability, offer little improvement over $s = 3$ for example) the effect at low values of s is substantial as can be seen. Concurrent starting points enable saturation of the algorithm's effectiveness at much lower values of o and not only increase $P_{deg\text{-}sum}$ and $P_{|\text{1-Hop}|}$ (which exhibits a nearly identical pattern of increase as $P_{deg\text{-}sum}$) but, as shown in the right-hand figure of Figure 4.3, cause marked improvement in D_{Base} as well, indicating that this change facilitates much better approximation of the degree placement.

However, it should be noted that this pattern is not present in all networks. In the CondMat network, increasing s had little effect and in a few settings was actually detrimental. This indicates that there is perhaps an underlying feature of the CondMat topology that benefits from localised crawling and that splitting the observations between multiple areas reduces the number of observations dedicated to this effort. The results of CondMat in Figure 4.1a lend additional weight to this hypothesis, with behaviour that is substantially different than the other two topologies despite being of comparable size to Enron (see Table 4.1). Overall though, increasing s by even a small amount is likely to benefit the performance of PO-Place.

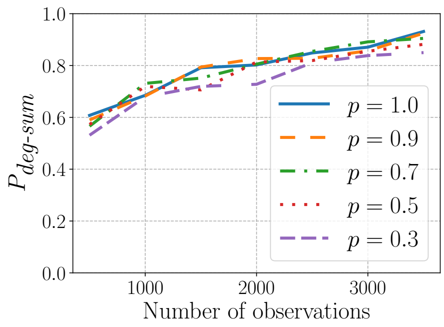

Figure 4.4: Effect of varying p on $P_{deg\text{-}sum}$. Twitter network, $n = 5$.

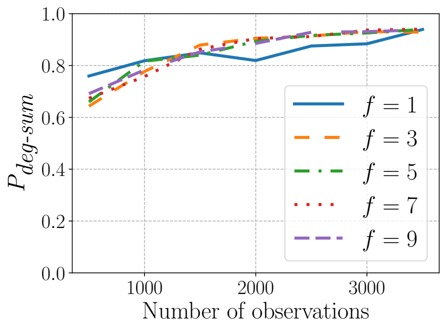

Figure 4.5: Effect of varying f on $P_{deg\text{-}sum}$. Enron network, $n = 30$.

Partial Neighbour Lists

In many settings, retrieving the whole of a node's neighbour list may also be impossible. Whether this is due to a technical limitation (only being able to retrieve a certain percentage of information) or because such information is not publicly available and is instead reserved for 'premium' or 'subscribed' users of such a network, ensuring that PO-Place is robust to such issues is a necessity to make it widely viable.

To simulate these restrictions, and measure their effect on the performance of PO-Place, the parameter, p, controls the proportion of a node's neighbours that may be explored. Neighbours are chosen uniformly at random from the full list to produce the restricted list that is presented to the algorithm. Results until this point have assumed that the full neighbour list for any agent is available upon request (i.e. $p = 1.0$). p is varied between 0.3 and 0.9 to determine the impact of this limitation. Representative results are shown in Figure 4.4 for the Twitter network and $n = 5$ but are applicable across all networks and values of o and n.

The results in Figure 4.4 show that different values of p have minimal effect on the performance of PO-Place. For all values of p, $P_{deg\text{-}sum}$ is comparable. Performing a 95% confidence interval Welch's t-test against the $p = 1.0$ results at each point, only $p = 0.3$ ($o = 1500, 2000, 3500$) and $p = 0.5$ ($o = 1500, 3500$)

are significantly worse. This pattern of minimal difference is repeated in all networks, with none seemingly more susceptible or affected by partial neighbour lists. We conclude that PO-Place is robust to receiving only partial information of this nature and is primarily unaffected by such limitations.

Breadth-First vs Depth-First Expansion

Finally, we turn our attention to the concept of breadth-first vs depth-first expansion in PO-Place. That is, when crawling the local area, should additional current area expansion be performed before considering new additions (breadth-first) or purely iteratively (depth-first). Where there is locally a clearly defined degree gradient we expect the latter to perform better. However, depth-first expansion also risks expending all the observations whilst exploring a suboptimal, locally maximal path.

Parameter f allows study of this by controlling how many of the current highest degree nodes that PO-Place is aware of are expanded concurrently. Experiments up until now have had $f = 1$ (depth-first). We now vary f from 1 to 9. Figure 4.5 presents these findings in the Enron network for $n = 30$. As with the previous results, it is our finding that the patterns here are replicated throughout the different topologies and values of n.

Similar to the findings when varying p, varying f has little absolute impact on the capabilities of PO-Place. However, using a 95% confidence interval Welch's t-test, all but $f = 9$ are statistically significantly worse at $o = 500$. This is likely due to the limited observations being focused too locally. All are significantly better between $o = 2000$ and $o = 3000$ but there is little gain in selecting values of f beyond 3 as the performance of PO-Place is almost identical. Overall, PO-Place seems to gain little from considering the local area more thoroughly before further expansion. Whether this is intrinsic in the design or a facet of the topologies being explored is outside the scope of this thesis but should be considered for future expansions.

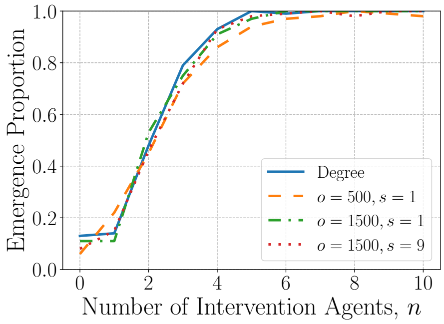

(a) CondMat

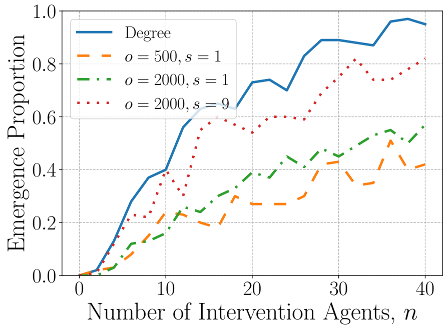

(b) Enron

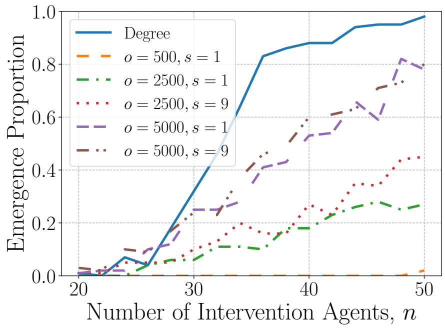

(c) Twitter

Figure 4.6: Comparison of PO-Place and degree IA placement for convention emergence in real-world topologies. The y-axis indicates the proportion of runs where the desired strategy emerged as the convention.

4.4 Initial Intervention with Limited Observations

In this section we use the findings of the performance of PO-Place with various parameters and apply it to the problem of encouraging convention emergence in the real-world topologies specified. We then expand this investigation into synthetic graphs and examine the performance of PO-Place within these.

4.4.1 Real-World Networks

Having explored the effectiveness of PO-Place in finding high-degree locations under different real-world topologies, ensuring the algorithm is robust to different aspects of partial observability, we now examine how PO-Place compares to degree placement for IAs in convention emergence in static networks. Having established ranges of parameters that offer the best performance improvements for each topology, these will be utilised to compare the algorithm to degree placement. Additionally, basic settings (small numbers of observations, no concurrent placements) provide a baseline comparison of PO-Place.

For the convention emergence experiments, a population of agents is situated in the 3 real-world topologies described previously. Each timestep, each agent chooses one of its neighbours uniformly at random to play the 10-action coordination game [Sen & Airiau, 2007] receiving positive or negative payoffs depending on whether their choices match. Agents use a simplified Q-Learning algorithm to learn the most beneficial choice. We utilise the 10-action game as used by [Marchant et al., 2015b] to avoid the issues of small convention spaces raised in Section 4.2 and to allow comparison to previous work. They have a chance to randomly choose their action ($p_{explore} = 0.25$) or else choose the most beneficial one. IAs replace the agents at the chosen locations and always choose their predetermined action.

A convention has emerged when the population has converged to have one action as the dominant choice of agents in the network. Most work considers this to be the case when the convention reaches 90% dominance [Griffiths & Anand, 2012; Marchant et al., 2015a; Sen & Airiau, 2007]. However, much of this work utilises synthetic networks rather than real-world topologies and populations that are substantially smaller than those we consider. Preliminary experiments show that the topologies are relatively resistant to convention emergence, requiring both high numbers of IAs well as substantial time. As we are concerned with a comparison of the performance of PO-Place against pure degree placement we wish to find settings that are guaranteed to repeatably experience convention emergence. As such, we consider a convention to have emerged when the 80% Kittock criteria is met, $K_{80\%}$ [Kittock, 1995]. That is, a convention has emerged when 80% of the population, when not exploring, would choose the same action. This indicates a high level of dominance of the desired action and allows more robust comparisons. We find that such a threshold is reliably reached, if it is likely to be reached at all, within 10000 iterations for the CondMat and Twitter networks and within 15000 iterations for the Enron network. As such, we measure the proportion of runs that have converged to the desired strategy within these time-frames across all networks.

The results are presented in Figure 4.6. We begin by varying the number of IAs, n, and finding a range where degree placement exhibits noticeable changes in convention emergence rates. We then utilise PO-Place across this same range with the parameters indicated. The values of o chosen within each topology are such that the number of observations is, at most, approximately 5% of the agent population. All runs are performed 100 times and the proportion of runs that produce the desired convention (strategy chosen uniformly at random at time $t = 0$ and assigned to all IAs) is measured.

Figure 4.6a shows the results for the CondMat topology. As was expected, due to the behaviour of CondMat in the PO-Place experimentation, all of the chosen parameters produce comparable results to the pure degree placement. Even at the worst performing parameters ($o = 500, s = 1$) there is no discernible difference between the performance of degree placement and PO-Place, whilst at higher number of observations (where PO-Place was entirely approximating the highest-degree nodes as seen in Figure 4.1a) the performance is as expected. Of note is the fact that, whilst it resulted in worse output of PO-Place in the prior section, increasing s does not noticeably affect the performance here.

Within the other networks the difference in performance is more noticeable but still indicates that PO-Place is generating close approximation of the degree placement. In both Enron and Twitter (Figures 4.6b and 4.6c) the minimal observation situation performs substantially worse than degree placement, particularly in the Twitter network. However, when given observations of only around 5% of the network ($o = 2000$ for Enron, $o = 5000$ for Twitter), the performance of PO-Place increases significantly. Whilst it still falls behind the performance of degree placement in both networks the difference is substantially smaller with PO-Place performing around 50-65% as effectively on average as degree placement in both networks (0.50 ± 0.08 in Enron, 0.64 ± 0.13 in Twitter). However, when we increase s, as was found in Section 4.3.2, it improves this substantially to 0.82 ± 0.11 average effectiveness compared to degree placement in Enron and, less substantially, to 0.69 ± 0.12 in the Twitter network. We quantify these values by comparing the emergence proportions of PO-Place and degree at each value of n and calculating the ratio between them which we then average. We discount values where either placement is achieving less than a 0.1 emergence proportion to avoid noisy results influencing the measure. As 0.1 is the expected emergence proportion of our desired strategy in a convention emergence we do not influence, we believe discounting values below this allows a more accurate comparison between the two algorithms. In the Twitter network, we also consider $o = 2500$ as the effect of increased s was more pronounced for this value during Section 4.3.2. Whilst there is a noticeable improvement at higher n the average compared effectiveness has a smaller difference: 0.23 ± 0.05 for $s = 1$ and 0.32 ± 0.09 for $s = 9$.

Additionally, we use a one-tailed Z-Test with a 95% confidence interval to compare the emergence proportions achieved between the $s = 1$ and $s = 9$ cases (where appropriate). We find that in the Twitter network ($o = 2500$) the difference is only significant at $n = 24, 34, 48, 50$. By comparison the results in the Enron network ($o = 2000$) are significant at $n = 4 - 10, 14 - 40$ (note that only even values of n are included in Figure 4.6b) showing that the improvement exhibited here is consistent and significant at nearly all tested values of n. We conclude that concurrent searches are less effective in the Twitter network, with o being the dominant factor, but that it makes a substantial improvement in the Enron network.

Overall, we have shown that even when only observing a small portion of the underlying topology, and strategically using these observations to maximise their effect, it is possible to achieve comparable performance to degree placement with full network visibility using PO-Place.

4.4.2 Synthetic Networks and Limited Observations

Having established the efficacy of PO-Place in encouraging convention emergence in real-world topologies we now turn our attention to using PO-Place in synthetic graphs. Whilst it is important to test the algorithm in graphs constructed from real-world social data this has a number of shortcomings. Firstly, that we are unable to vary any of the aspects of the network topology to explore the effect that these have on the performance and robustness of PO-Place. Unless we sample the network we cannot change the size, edge count, density or any other topological value. Secondly, there are noticeable distinctions and differences between synthetic and real-world networks [Leskovec et al., 2008; Newman, 2003] that often mean they are poor approximations of one another. Finally, that there are real-world *constructed* networks, in contrast to real-world social networks, that exhibit features similar to synthetic topologies, particularly in computing and telecommunication networks such as the Internet itself [Dorogovtsev & Mendes, 2003; Lewis, 2006; Lian-Ming et al., 2011]. Because of these differences and the capability to vary the network topology it is important to establish the effectiveness of PO-Place in synthetic networks as well as those already seen and to evaluate any differences that arise.

We use the Java Universal Network/Graph (JUNG) framework[1] to generate both scale-free [Barabási & Albert, 1999] and small-world [Kleinberg, 2000b] networks upon which to test PO-Place. Both of these types of topologies are frequently used in the literature due to having features often found in real-world topologies [Delgado et al., 2003; Sen & Airiau, 2007]. Namely, in scale-free networks there is a power-law distribution in degree (scale-free networks) which results in a few high-degree "hub" nodes and a much larger number of low degree nodes. In small-world networks the topology is characterised by a small diameter and a well-connected local area around each node which results in the

[1] JUNG version 2.1.1 (`http://jrtom.github.io/jung/`)

well-known Six Degrees of Separation phenomena [Travers & Milgram, 1969]. More details on these topologies and their generation can be found in Chapter 2.

Unlike previous usage of synthetic graphs in Chapter 3 initially in this section we pre-generate synthetic graphs to allow a baseline comparison to occur. If we were to generate a new graph for each run of PO-Place then we would be unable to compare D_{Base} meaningfully between runs as each would generate a different set of high-degree nodes. Using pre-generated graphs allows us to compare the average performance of PO-Place upon a given synthetic graph. Additionally, this approach is the same as that taken in Section 4.4.1 where the networks were unchanging real-world topologies. When PO-Place is used to encourage convention emergence later in this Section this condition will be changed as at that point we wish to establish the general performance of PO-Place in synthetic graphs of a given type.

Experimental Setup

We begin as we did in Section 4.3.2 by establishing the performance of PO-Place independent of its use in encouraging convention emergence. To this end we create two synthetic networks:

Scale-free As in the previous chapter the scale-free network is generated using the Barabási-Albert model [Barabási & Albert, 1999]. We use an initial number of vertices, $m_0 = 4$, with $m = 3$ edges being attached per iteration of the generator. We create a graph consisting of 50000 vertices and 149988 edges.

Small-world A small-world topology is generated using Kleinberg's model [Kleinberg, 2000a]. As with the scale-free network it has 50000 vertices with a clustering exponent of 2 and 1 additional long-range connection per vertex. This produced a graph with 149932 edges.

The sizes of the graphs has been chosen to allow direct comparison to the real-world topologies used before with 50000 vertices being between the Enron

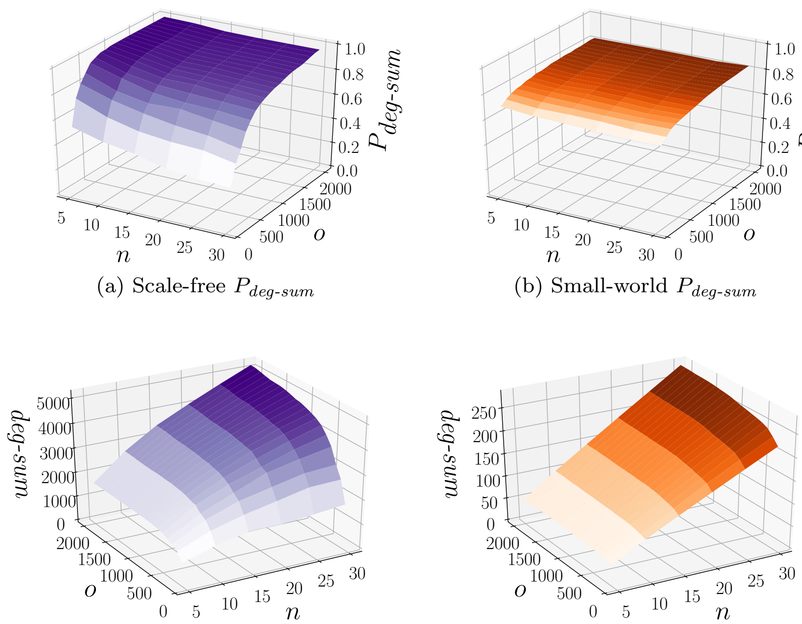

(a) Scale-free $P_{deg\text{-}sum}$ (b) Small-world $P_{deg\text{-}sum}$

(c) Scale-free *deg-sum* (d) Small-world *deg-sum*

Figure 4.7: $P_{deg\text{-}sum}$ and *deg-sum* performance of PO-Place for varying n (number of locations) and o (number of observations) in synthetic networks with 50000 nodes.

and Twitter network sizes. This will minimise the likelihood that any differences are due to the size of the topologies and rather is likely to be from other differences between them.

These graphs are unchanging and are used to examine the general performance of PO-Place. Whilst there is a risk of individual aspects of these graphs being unique to them the generic nature of the graph generation algorithms and the size of these topologies makes this unlikely. Additionally, as when examining the performance of PO-Place on the real-world topologies, the placement algorithms are averaged from multiple independent runs. In this instance each setting of the algorithm parameters is run 100 times and the results averaged.

PO-Place Performance

We begin as we did before by investigating the performance of the algorithm with $s = p = f = 1$ and focusing on the difference in performance as we vary n and o. We vary the number of locations to find, n, between 0 and 30 in steps of 5. The number of observations, o, is varied between 100 and 2000 in steps of 100. Note that this maximum is less than even the smallest maximum used when examining the real-world topologies (CondMat with $o = 3500$) and the reason for this is presented in Figure 4.7. Figures 4.7a and 4.7b show how $P_{deg\text{-}sum}$ changes as we vary these parameters whilst Figures 4.7c and 4.7d show the same but for the full *deg-sum*. As can be seen in the top row of the figure, the performance of PO-Place even at these levels of observation (which represents between 0.2% and 4% of each of the networks) is substantial and, in terms of $P_{deg\text{-}sum}$, close to the pure degree placement. For comparison, at $n = 10, o = 1000$ in the Enron network $P_{deg\text{-}sum}$ was 0.82 ± 0.10 and in Twitter it was 0.66 ± 0.12. In the scale-free network (which it should be noted is larger than Enron) at these values $P_{deg\text{-}sum}$ performance was 0.96 ± 0.02. This level of increased performance even at lower percentages of network observation is seen throughout Figure 4.7a.

As can be seen in Figure 4.7b the performance in the small-world network is noticeably different from the previous plots of both the real-world topologies and the scale-free network with little difference in performance from 0.2% network observation to 4% network observation. The increase in $P_{deg\text{-}sum}$ between these number of observations is minimal with the performance plateauing around 0.8. This is in contrast to each of the other examinations thus far where the performance has rapidly approached 1.0. The underlying features of small-world networks, in comparison to the others, mean that there are less likely to be the "hub" nodes which occur frequently in other topologies. As PO-Place moves along gradients of increasing degree the absence of these will likely hinder the algorithm. However, even with this the algorithm easily finds nodes with approximately 80% of the degrees of the pure degree placement. Indeed, at very low levels of observation ($o \lessapprox 500$) PO-Place in small-world outperforms

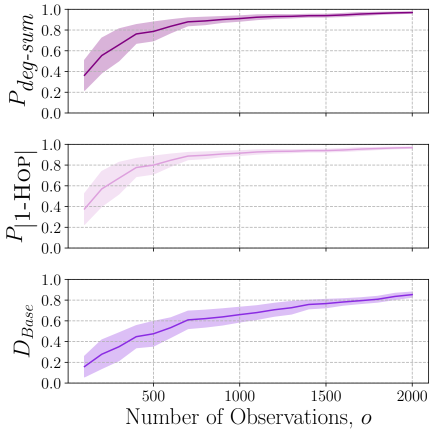

(a) Scalefree

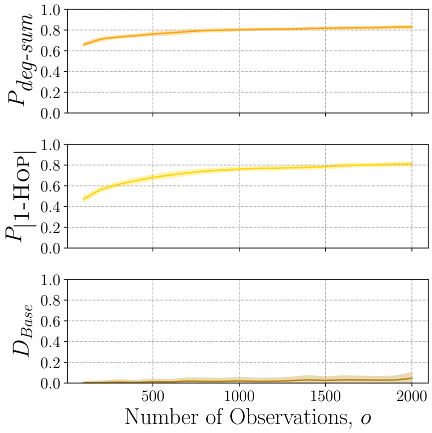

(b) Smallworld

Figure 4.8: Performance of PO-Place as measured by $P_{deg\text{-}sum}$, $P_{|1\text{-Hop}|}$ and D_{Base} for varying numbers of observations in scale-free and small-world networks for $n = 20$.

PO-Place in scale-free topologies.

A more in depth analysis of the PO-Place performance metrics in synthetic networks is presented in Figure 4.8 which shows the change in $P_{deg\text{-}sum}$, $P_{|1\text{-Hop}|}$ and D_{Base} in both networks as we vary the number of observations. For both networks we look at the case when $n = 20$ but similar patterns are observed throughout. Each of the metrics is plotted against the number of observations, o, and the shaded areas in each plot represent the standard deviations amongst the runs. As can be seen in Figure 4.8a the metrics in the scale-free network follow a similar pattern of diminishing returns as was seen in Figure 4.2. However, there are a number of differences that we highlight here. Primarily, as was seen in Figure 4.7, the performance in all metrics is substantially better than was seen in Figure 4.2. Noting that Figure 4.2 starts at $o = 500$ whilst this starts at $o = 100$ we see that at all points performance here is noticeably improved. Additionally, the standard deviations seen decrease to negligible as the observations increase, a trend not seen in Figure 4.2. This means that PO-Place is choosing *consistently* better in the scale-free topologies than in the real-world ones and this can be explained due to the gradient-climbing nature

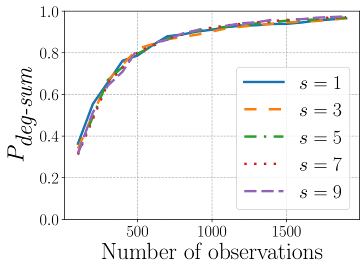

(a) Scalefree

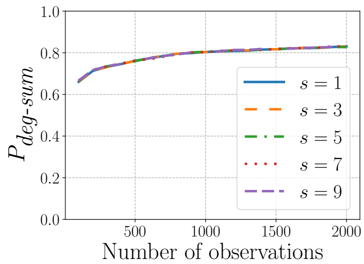

(b) Smallworld

Figure 4.9: $P_{deg\text{-}sum}$ performance of PO-Place for varying s in scale-free and small-world networks for $n = 20$.

of the algorithm once again as the scale-free nature of the real-world topologies is likely noisier than those of the generated topologies.

In comparison, Figure 4.8b shows the performance of the metrics in the small-world topology at $n = 20$. Whilst both $P_{deg\text{-}sum}$ and $P_{|1\text{-Hop}|}$ stay around 0.8 after initial small amounts of growth at lower levels of observations, the performance of D_{Base} is consistently low, barely above 0. This supports the hypothesis that PO-Place is performing poorly in small-world networks when it comes to traversing and locating the highest degree nodes. Whilst it is finding *some* nodes of high degree it is not finding the same ones as pure degree placement does and the lack of the degree gradient that PO-Place makes use of is likely the underlying cause. However, as was mentioned in Section 4.3.2, a poor D_{Base} score is not necessarily indicative of poor performance when using PO-Place for IAs.

As with the evaluation in real-world networks we also seek to establish the effect that varying other parameters of PO-Place has on performance. These are the number of concurrent starting locations (s), limiting the proportion of neighbours that are available for expansion (p) and the fringe size of vertices expanded before choosing again (f). Figure 4.9 shows the effect of varying s between 1 and 9 in both of the synthetic networks for $n = 20$. In both of them there is no significant difference in performance for any value of s. Indeed there is no difference in performance for any value of n examined. This is in

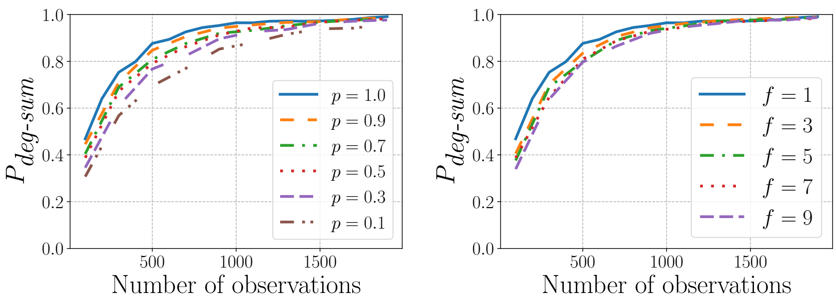

(a) Varying p

(b) Varying f

Figure 4.10: Effect on PO-Place in a scale-free network of varying p and f. $n = 10$.

	Number of Observations, o									
p	100	300	500	700	900	1100	1300	1500	1700	1900
1.0	0.469	0.753	0.877	0.926	0.953	0.965	0.972	0.973	0.979	0.991
0.9	0.446	**0.709**	0.847	**0.906**	0.942	0.956	0.966	0.97	0.981	**0.987**
0.7	**0.404**	**0.69**	**0.808**	**0.876**	**0.92**	**0.932**	**0.952**	0.967	0.977	**0.98**
0.5	**0.388**	**0.673**	**0.791**	**0.863**	**0.913**	**0.943**	**0.946**	0.968	0.975	**0.983**
0.3	**0.345**	**0.609**	**0.767**	**0.824**	**0.896**	**0.927**	**0.936**	0.963	0.972	**0.976**
0.1	**0.307**	**0.568**	**0.695**	**0.77**	**0.852**	**0.88**	**0.916**	**0.939**	**0.942**	**0.96**

Table 4.2: $P_{deg\text{-}sum}$ scores for o and p in the 50000 node scale-free network. Values which have statistically significant difference from the $p = 1.0$ performance are shown in bold.

contrast to the real-world topologies where, although it was not consistent or for all values, there were substantial and significant increases in efficacy when increasing s.

We also vary p and f to see the effect this has on the performance of PO-Place. The results of this for the scale-free network are shown in Figure 4.10. We use the results from $n = 10$ but similar results are found for all values of n. We vary p from 1.0 down to 0.1 and f from 1 to 9. As seen in the figures, unlike in the real-world topologies, there is a noticeable effect on the performance of PO-Place when both of these are varied. Decreasing the proportion of neighbours that is provided to the algorithm, as in Figure 4.10a decreases the effectiveness of PO-Place as it was originally hypothesised that it would. However, even

	Number of Observations, o									
f	100	300	500	700	900	1100	1300	1500	1700	1900
1	0.469	0.753	0.877	0.926	0.953	0.965	0.972	0.973	0.979	0.991
3	**0.405**	**0.703**	**0.835**	0.906	0.943	0.961	0.973	**0.981**	**0.986**	0.989
5	**0.376**	**0.69**	**0.799**	**0.889**	0.931	0.951	0.972	**0.977**	0.981	0.989
7	**0.387**	**0.647**	**0.809**	**0.888**	**0.928**	**0.946**	0.967	0.971	0.978	**0.983**
9	**0.339**	**0.639**	**0.797**	**0.865**	**0.919**	**0.947**	**0.961**	**0.977**	0.976	0.987

Table 4.3: $P_{deg\text{-}sum}$ scores for o and f in the 50000 node scale-free network. Values which have statistically significant difference from the $f = 1$ performance are shown in bold.

$p = 0.1$ only reduces the effectiveness of the algorithm slightly although it means that only 10% of neighbours are available to the algorithm. Similarly, increasing the fringe size as shown in Figure 4.10b has a similar detrimental effect although less pronounced. Increasing f means that more observations are spend exploring the immediate area rather than traversing the degree gradient so this effect is to be expected.

To highlight these differences more the results are also presented in Tables 4.2 and 4.3. The $P_{deg\text{-}sum}$ values are shown in the table and those that are statistically significant in their difference from the baseline are highlighted in bold. These were found using a 95% confidence interval Mann-Whitney U test [Fay & Proschan, 2010; Mann & Whitney, 1947] and highlight that the differences are significant for nearly all values of p and f at each number of observations. The effect of varying p is particularly negative with most scores being significantly different even if the difference is not large in an absolute sense. However, these small absolute differences highlight that PO-Place is robust to these variations as well as reassuring us that a simple depth-first growth approach is the best one in all topologies.

Figure 4.11 shows the effect of varying p and f on the performance of PO-Place in the small-world topology. We use $n = 30$ as the value where differences should be most noticeable due to the algorithm needing to find a larger number of high-degree locations. As has been the case throughout, these results are

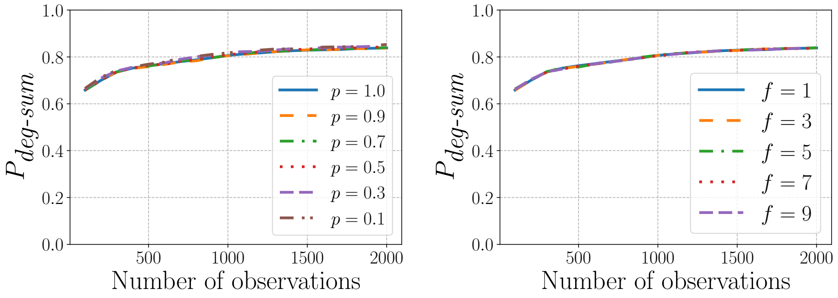

(a) Varying p (b) Varying f

Figure 4.11: Effect on PO-Place in a small-world network of varying p and f. $n = 30$.

	$\|V\|$	$\|E\|$	$\overline{deg}$	$\|E\|/\|V\|$
CondMat	21,363	91,286	8.546	4.273
Enron	33,696	180,811	10.732	5.366
Twitter	81,306	1,342,296	33.018	16.509
Scale-free-50k	50,000	149,988	6.000	3.000
Small-world-50k	50,000	149,932	5.997	2.999

Table 4.4: Average degrees for graphs used. Real-world networks are the largest WCC within them.

substantially distinct from those in the scale-free topology with the changes in parameters having no discernible effect on the algorithm. This again supports the hypothesis that the algorithm in small-world graphs is able to find nodes of approximately high enough degree in its local area, regardless of whether they are the nodes that would be selected through pure degree placement. Whether this will be enough to be of use in convention emergence tasks will be explored later.

Varying Topology Features

Thus far we have varied the parameters of the PO-Place algorithm and seen the effect of this on its performance in the synthetic graphs. We now consider the effect of the underlying topology and the variations that can be generated in the synthetic graphs themselves.

The differences in the performance of PO-Place can be explained by the different features of these underlying topologies with the lack of high degree "hubs" being a prominent facet in the performance in small-world topologies. The noisier nature of the real-world topologies in comparison to the scale-free networks generated by the Barabási-Albert model similarly explains aspects of this difference in performance.

One noticeable difference in the performance of PO-Place in the real-world and synthetic networks is that varying s has no effect in the synthetic networks. Whilst its effect was not found throughout the real-world networks (being most prominent within the Enron topology) its lack of any effect, positive or negative, highlights that something is distinct in the real-world networks compared to the synthetic ones.

As noted and explored by Franks et al. [2014] the global network metrics generated by the synthetic graphs often differ from those exhibited by real-world networks. They highlight a number of metrics such as graph diameter and clustering coefficient which are substantially different in the generated graphs vs the real-world and real-world sampled graphs within their work.

One other metric they highlight is that of average degree which also differs between real and synthetic graphs. Average degree is defined as:

$$\overline{deg} = \frac{\sum_{v \in V} deg(v)}{|V|} = \frac{2|E|}{|V|} \tag{4.4}$$

The average degrees for the graphs used in this chapter are shown in Table 4.4 as well as the number of vertices and edges in each graph. The real-world networks are again reduced to their largest weakly-connected component. As can be seen, the synthetic graphs have lower average degrees than even the CondMat network (which is much smaller) and substantially lower average degrees when compared to the Enron and Twitter networks. Whilst the number of nodes within the synthetic networks is comparable to those in the real-world networks other aspects such as average degree are very different.

We now investigate the effect this has on the performance of PO-Place by producing a number of topologies of both scale-free and small-world types and varying the average degree within them. From this point onwards, for ease of notation, we refer to the edge:vertex ratio, $|E|/|V|$, instead of average degree. These are the same barring a constant factor (as can be seen above) and as such we lose no finesse by doing this. These values are similarly presented in Table 4.4. Note that we thus have already been using synthetic topologies for $|E|/|V| = 3$.

We generate new graphs with $|E|/|V| = 6, 9, 12, 15$ to see how PO-Place is effected. For each synthetic graph this is done by changing the graph generation parameters associated with the topologies as follows:

Scale-free The generated graph has $|E| = m(|V| - m_0)$ and so $|E|/|V| \approxeq m$. $m_0 = m = |E|/|V|$ for $m = 6, 9, 12, 15$.

Small-world In the Kleinberg model each node is connected to each of 4 neighbours in a lattice with l additional long-range connections. Thus $|E| = |V|(2 + l)$ (due to avoiding double counting). This means that $|E|/|V| = l + 2$ and thus we use $l = 4, 7, 10, 13$.

Having already established how PO-Place in synthetic networks is affected by varying p and f we now focus on the effect of increasing s. As an increased $s = 9$ was beneficial in the real-world networks with higher $|E|/|V|$ we use this to evaluate our algorithm in denser synthetic networks.

For each of the networks generated above we vary the number of locations requested and number of observations as before. Additionally we use $s = 1$ and $s = 9$ for all these values and compare the $P_{deg\text{-}sum}$ of the results. Each combination of network, n, o and s are run 100 times and the averaged $P_{deg\text{-}sum}$ used. As we are interested in whether the increased s setting improves the performance of PO-Place or not, we then compare the results using the Mann-Whitney U test at the 95% confidence interval ($p < 0.05$) to establish if the difference is significant.

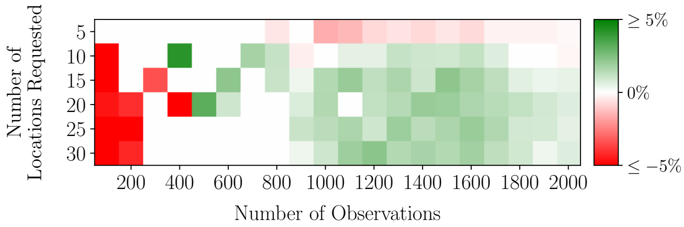

(a) $|E|/|V| = 3$

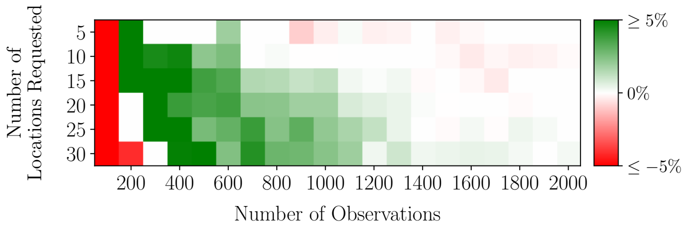

(b) $|E|/|V| = 6$

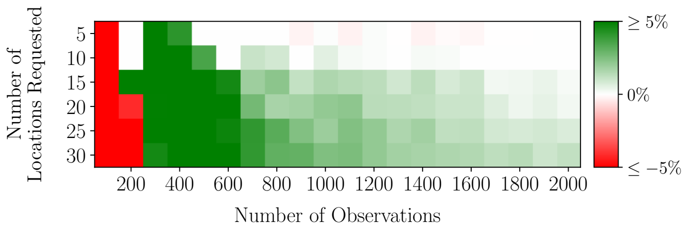

(c) $|E|/|V| = 9$

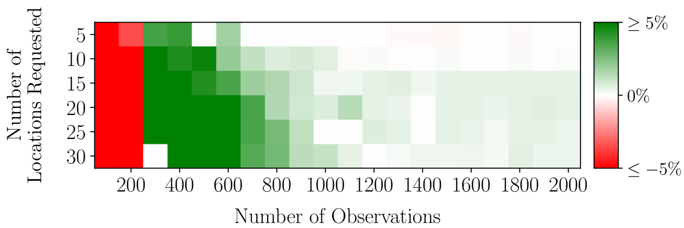

(d) $|E|/|V| = 12$

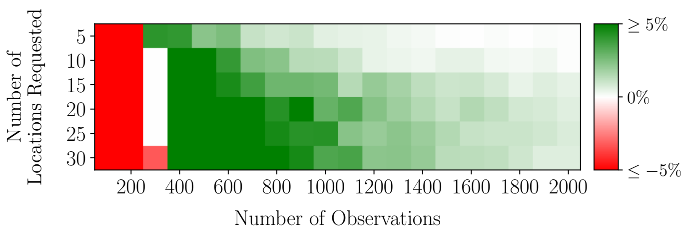

(e) $|E|/|V| = 15$

Figure 4.12: Significant differences ($p \leq 0.05$) when increasing number of starting points for PO-Place in scale-free graphs of 50000 vertices at different edge:vertex ratios.

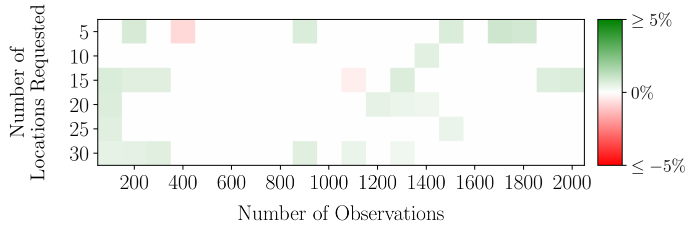

(a) $|E|/|V| = 3$

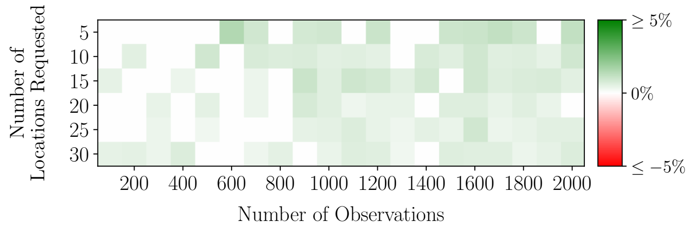

(b) $|E|/|V| = 6$

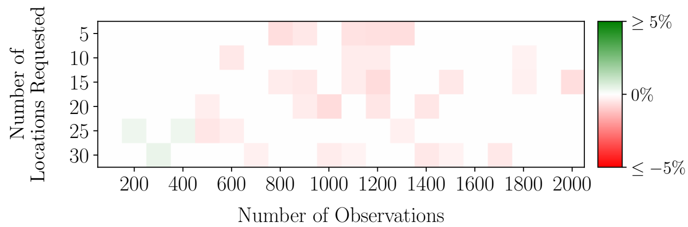

(c) $|E|/|V| = 9$

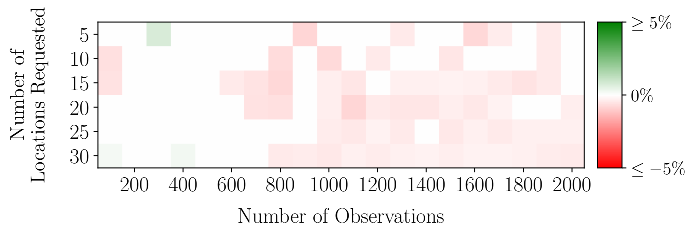

(d) $|E|/|V| = 12$

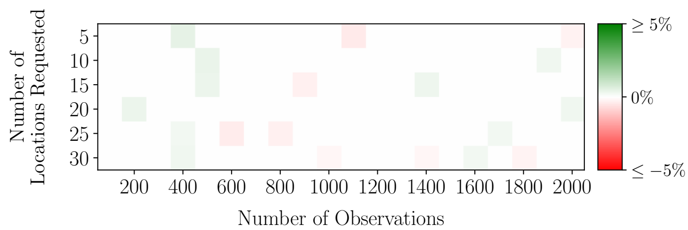

(e) $|E|/|V| = 15$

Figure 4.13: Significant differences ($p \leq 0.05$) when increasing number of starting points for PO-Place in small-world graphs 50000 vertices at different edge:vertex ratios.

The results of these evaluations are presented in Figure 4.12 for the scale-free networks. Each combination of n and o is assigned a colour based on the difference and significance of the results. White indicates that there was no statistically significant difference at the 95% confidence interval between the $P_{deg\text{-}sum}$ performance of PO-PLACE with $s = 1$ and $s = 9$ for these values. Red indicates that $s = 9$ was statistically significantly worse than $s = 1$ whilst green indicates that it was better to a statistically significant level. The saturation of this colour is then determined by the differences in $P_{deg\text{-}sum}$ as indicated in the colourbar, up to a maximum of $+5\%/+0.05$ for solid green and $-5\%/-0.05$ for solid red. For instance, if $s = 9$ had $P_{deg\text{-}sum} = 0.85$ and $s = 1$ achieved $P_{deg\text{-}sum} = 0.77$ the location would be coloured solid green. Alternatively if $s = 9$ produced $P_{deg\text{-}sum} = 0.66$ and $s = 1$ had $P_{deg\text{-}sum} = 0.67$ it would be coloured pale red.

Figure 4.12 shows that the average degree (and edge:vertex ratio) has a substantial effect on the performance of PO-PLACE. At low average degree, as has been the case thus far, increasing s has a mostly neutral or barely beneficial effect, explaining the behaviour seen in Figure 4.9a. At very low numbers of observations it causes decreases in performance of greater than 5% consistently. This is to be expected as the algorithm is splitting an already highly limited number of observations between multiple locations and is something that continues regardless of the density. For many other combinations of n and o increasing s has a non-significant effect although at higher numbers of observations there are improvements but all of these are less than 5%.

As we increase the density the performance improvement caused by $s = 9$ increases rapidly. At $|E|/|V| = 6$, comparable to the Enron network, we see that a number of locations are now undergoing substantial $P_{deg\text{-}sum}$ increases, particularly at the lower end of the number of observations, although the poor performance at the lowest number of observations persists. Additionally a larger contingent of increased performance is present in the lower to middle ranges of observations indicating that the increased density is being exploited by PO-

Place to better find the high degree nodes.

This pattern of increasing performance continues for higher levels of $|E|/|V|$ with $|E|/|V| = 9$ and $|E|/|V| = 12$ increasing the number of locations that exhibit greater than 5% improvement as well as continuing to expand the number of locations that have statistically significant improvements to lesser amounts as well. It is worth noting however that the number of locations where performance decreases substantially also increases for the lowest values of o. This is to be expected as the increased average degree means that the algorithm is unable to explore as far along the degree gradient before using up its available observations due to encountered nodes generally having more neighbours that must be explored. As the levels of observation at which this occurs are less than 0.4% of the graph this is a small fraction of the situations where PO-Place may be applied. However, the switch between this negative effect and the positive increases is a sharp delineation and indicates that there is a phase shift in the effect of having multiple starting locations from which to search. Increasing the number of observations from 0.4% of the network size to even 1% dramatically increases the effectiveness of PO-Place in this way and overall gives credence to the manner in which PO-Place operates.

When we reach an $|E|/|V|$ value that is close to that exhibited by the Twitter network, PO-Place is consistently and nearly universally benefiting from an increased s except at very low numbers of observations. Although the effect diminishes at lower numbers of locations requested or at higher numbers of observations it is not detrimental in this situations, being at worst not statistically significant. At all other combinations of n and o the increase in starting locations has a significant and often substantial effect. The areas where increasing s does not benefit PO-Place are areas not where increasing s makes the algorithm worse but rather areas where the $s = 1$ approach already performs well. For lower numbers of locations requested it only has to choose the top few locations well rather than choosing well consistently. At higher numbers of observations provided, ensuring they are used effectively becomes less important.

However, the effect of increased average degree in small-world networks exhibits very different dynamics. As shown in Figure 4.13, after a brief but shallow increase in $P_{deg\text{-}sum}$ performance for $|E|/|V| = 6$ (where any statistically significant increases are for only a few percentage points), increasing $|E|/|V|$ to 9 and 12 has the opposite effect, with small but statistically significant decreases in effectiveness. Unlike in the scale-free topologies there is no consistent pattern or range of combinations where this differences may occur and given the low level of difference that is present it is likely to be facets of the individual topologies rather than something intrinsic. This is supported by the performance differences shown for $|E|/|V| = 15$ which follow no pattern that could be inferred from the previous levels of edge:vertex ratio. This all lends further evidence to the theory that PO-Place in small-world networks is unable to make use of many of the assumptions built into the algorithm and, whilst it does not perform poorly by the metrics employed, it does not benefit in the same way that the real-world and scale-free topologies do.

We have shown that PO-Place performs well in synthetic topologies as well as in real-world networks. In scale-free graphs it performs even better than with comparable parameters in the real-world topologies and benefits from the additional functionality of PO-Place, achieving greater than 90% performance with only 1-2% graph observation. In small-world networks PO-Place performs well, achieving approximately 80% of the pure degree performance with 1-2% network observation. However increasing beyond this is difficult due to aspects of the underlying small-world topology.

Convention Emergence

Having established that PO-Place is applicable and effective in synthetic networks we now use it to select locations for IAs and examine the effectiveness of its selection in comparison to that of pure degree placement with full network visibility.

To evaluate the effectiveness of PO-Place in synthetic networks we perform the same types of simulation as in Chapter 3 as well as in Section 4.4.1. We use two different networks, scale-free and small-world, at two different scales, 20000 and 50000 nodes, to produce four different combinations. Each of the graphs has $|E|/|V| = 9$ due to its being between the values of CondMat/Enron and Twitter. The parameters used to generate graphs with this value of $|E|/|V|$ are the same as in the previous section. Each individual run generates a new graph in contrast to the approach performed for PO-Place without convention emergence. This is to ensure that the results give a general picture of the behaviour of convention emergence in each graph type and scale.

We create a population of agents and situate them within the topology. As before, in each timestep each agent chooses one of its neighbours with whom to play the 10-action coordination game and every agent uses Q-Learning [Griffiths & Anand, 2012; Sen & Airiau, 2007] to update its knowledge based on the payoff received. Each agent also explores a random action choice with $p_{explore} = 0.25$. For the purposes of these simulations we consider a convention to have emerged when the 90% Kittock Criteria, $K_{90\%}$ is achieved rather than the $K_{80\%}$ threshold used in the real-world network evaluations. This is because the synthetic networks have been shown to consistently achieve this level of convention even without IAs in the system. We use both degree placement and PO-Place with various settings derived from well-performing parameters from the previous section to insert n IAs into the population at time $t = 0$. The simulations run for 3000 timesteps which was found to be long enough for conventions to robustly emerge even without outside aid. We measure the proportion of 100 runs that emerge the desired convention (that assigned to the IAs at $t = 0$, selected uniformly at random from those actions available) within that time frame.

Figure 4.14a shows convention emergence proportions in scale-free networks with 20000 nodes. As was found in Chapter 3 and in previous work by Griffiths & Anand [2012] and Sen & Airiau [2007], only a few IAs are needed to affect a population substantially larger than themselves with less than 10 being able to consistently cause a robust convention to emerge whilst only representing

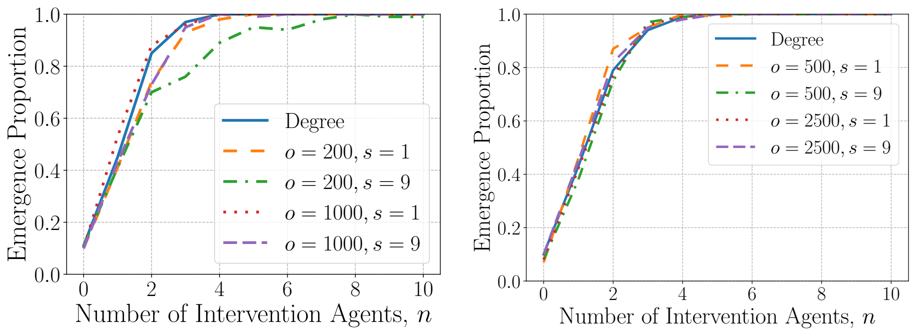

(a) 20000 nodes

(b) 50000 nodes

Figure 4.14: Convention emergence using PO-Place in scale-free networks.

0.05% of the population. PO-Place performs well, keeping pace with degree placement and being able to effect convention emergence to the same scale in the same range of IAs with $o = 200, s = 1$ achieving a proportion of 0.97 ± 0.05 of degrees performance. Of particular note is the fact that, due to the low numbers of observations as well as such a small amount of locations being needed, $o = 200, s = 9$ falls into the area in which PO-Place performs worse, being significantly ($p < 0.05$) different at $n = 3 - 6$ and only achieving a proportion of 0.93 ± 0.07. With $o = 1000$ PO-Place performs equally as well as the pure degree placement, despite only observing 5% of the network and achieving performance of above 95% with only 1% of the network observed. Figure 4.14b shows similarly high-levels of performance in the 50000 node scale-free network with 1% observation ($o = 500$) achieving an averaged proportion of 0.99 ± 0.01 of degree placements performance. Overall, PO-Place is exceptionally effective in encouraging initial convention emergence in the scale-free network achieving near perfect performance with a tiny fraction of the network observed.

Figures 4.15a and 4.15b show the performance of PO-Place for small-world networks. As was hypothesised in the previous section, although PO-Place struggled when it came to the performance evaluation when used to find locations for IAs it still performs very well in small-world topologies. In both 20000 and 50000 node topologies even just 1% network observation allows

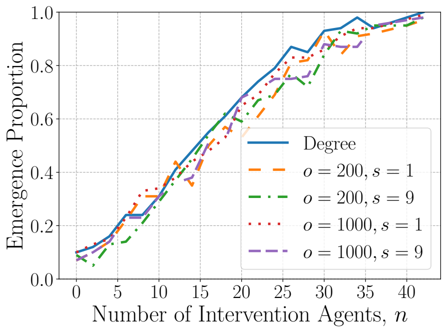

(a) 20000 nodes

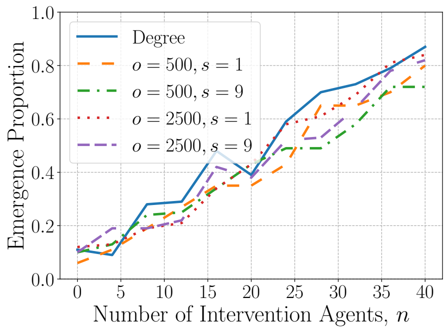

(b) 50000 nodes

Figure 4.15: Convention emergence using PO-PLACE in small-world networks.

averaged proportional performance compared to degree of 0.92 ± 0.12 in the 20000 node graph and 0.85 ± 0.17 in the 50000 node graph, performing better than the $P_{deg\text{-}sum}$ would indicate. Whilst larger numbers of IAs are needed than in scale-free topologies this is to be expected and is consistent with the general performance of convention emergence in small-world networks. As shown before, varying s has no effect on the efficacy of PO-PLACE when used for convention emergence in this domain making it neither better nor worse.

Overall we have shown that PO-PLACE is a powerful tool to be deployed in synthetic networks encouraging convention emergence comparably with pure degree placement whilst only observing small fractions of the underlying topology. It performs better in the synthetic networks than in the real-world networks and the various parameters available to the algorithm can be used to increase its performance dependent on network density.

4.5 Limited Observations in Dynamic Networks

The issue of partial observability in dynamic networks adds a number of new dimensions to the problem of finding highly influential nodes under these constraints. With the constantly changing nature of the dynamic topologies, additional consideration of how to best utilise available observations over different timesteps is paramount.

There are two likely use cases that we believe cover most situations with limited observations in dynamic topologies:

1. The user is supplied a total, finite number of observations and must determine how best to allocate these across the time available. For instance, situations where the controlling/hosting party of the network allows precisely limited access to the system, perhaps in exchange for remuneration (e.g. 5000 observations for £X).

2. The user is supplied a finite number of observations at regular intervals throughout the lifespan of the system. This represents the typical access limitations of many social network APIs: you may make a limited amount of API requests per hour/day/week.

We may identify 1 as *Bulk Supply* and 2 as *Staggered Supply*. Note that this does not cover the case where the number of observations supplied may vary at different timesteps. However, most API limitations observed (as the primary motivator for this research) are entirely static in nature and do not vary over time. An exception to this is where an API may have multiple levels of rate limiting such that, for example, you may make n observations per hour or m observations per day, whichever occurs first. In this instance we can treat the larger limit, m observations per day, as the actual limit of the supply of observations without loss of general applicability.

4.5.1 Formalisation

Numerous related but unique formalisations of dynamic graphs exist in the literature with no solid consensus on how best to describe them [Kim & Anderson, 2012; Kostakos, 2009; Michail, 2015; Nicosia et al., 2013; Pan & Saramäki, 2011; Tang et al., 2010]. Numerous approaches treat the edges as labelled with the timesteps the edge exists in. Others devolve the temporal graph into a multi-layered graph with connections between the vertex and itself in subsequent timesteps. Some treat the graphs in the individual timesteps as an ordered set of individual static graphs. However, in each of these approaches, the set of nodes is treated as static and unchanging. In our case, using the networks of González et al. [2006a] and Ichinose et al. [2013] as described in Chapter 2 results in a node set that is constantly changing. With this in mind, we propose and will use the following formalisation of dynamic graphs.

A dynamic graph, $G_T = (V_T, E_T)$, consists of a set of temporal vertices, V, and a set of temporal edges, E_T, where a temporal vertex $v_{i,j} \in V_T$ exists in the graph between timesteps i and j and a temporal edge $(u, v)_{k,l} \in E_T$ exists between vertices u and v between timesteps k and l. As such, it can alternatively be viewed as a sequence of individual static graphs at various timesteps with vertices and edges in those graphs that exist at those particular timesteps. That is, $G_T = (G_{t_1} = (V_{t_1}, E_{t_1}), \ldots, G_\tau = (V_\tau, E_\tau))$ where τ is the maximum timestep that the graph exists in.

Additionally, for the purposes of partial observability in dynamic networks we must formalise the nature of observations and their supply in dynamic topologies. We break the timesteps of the graph up into contiguous non-overlapping groups which we refer to as *blocks*. Given the number of observations supplied in each block, $O = (o_1, ..., o_i, ...)$, and the timesteps where these are supplied, $T_{supply} = (t_{o_1}, ..., t_{o_i}, ...)$ which denotes the start points of the blocks, the problem definition becomes how should these observations be best used to find highly influential nodes amongst the population.

This formalisation is general but slightly cumbersome for the way observation supplies work in many application domains. In particular, often the individual blocks are defined by the number of observations supplied and the frequency with which they are so. For instance, the rate-limited Twitter API [Twitter Developers, 2017a] allows 15 calls every 15 minutes. As such, we can instead denote situations such as this by using 3 parameters: o_{supply}, the number of observations received in each supply; f_{supply}, the number of timesteps between supplies; and t_{start}, the timestep at which the first supply is received.

Using these parameters, $O = (o_{supply}, o_{supply}, \dots)$ and $T_{supply} = (t_{start}, t_{start} + f_{supply}, t_{start} + 2f supply, \dots)$ allowing easier specification of common setups of observation supplies. Note that in this formalisation a *Bulk Supply* becomes a special case of *Staggered Supply* where $f_{supply} = \infty$ and staggered supplies are readily defined using the available parameters.

Thus, we formalise the Dynamic Partial Observability Problem as:

Given the following:

- A dynamic graph, $G_T = (V_T, E_T)$.
- A finite number of timesteps, τ.
- A number of observations per supply, o_{supply}.
- The number of timesteps between each supply, f_{supply}.
- The timestep of the first supply, t_{start}, where $0 \leq t_{start} \leq \tau$.

select a set of nodes of size n each timestep that maximises the *deg-sum* of the selected nodes.

We also assume the following restrictions on the observations:

- the observations cannot be "stored" between supplies; when a new supply comes in, any unused observations from the previous supply are discarded.
- a single observation will retrieve the list of neighbours (and hence degree) of a given node. This functionality is frequently available in APIs as a single command and hence is a good descriptor of what an "observation" is.

Using this formalisation we can now approach the problem of using the limited number of observations to produce a dynamic variant of PO-Place that attempts to maximise the *deg-sum* of the found nodes.

4.5.2 Design Considerations

With the addition of a dynamic component to the topologies, a number of new considerations arise that are not present in the static version of PO-Place.

Frequency of Updates

One of the factors explored in the previous dynamic topology work is that of how frequently the node lists should be updated. With limited observations this becomes even more problematic. Updating the lists multiple times between observation supplies requires spreading out the limited observations between these updates, potentially making each individual search much weaker. Only updating when a new supply is provided however runs the risk of the node lists being substantially out-of-date, particularly if the observation supplies are infrequent. The effects of these two opposing constraints must be balanced. Observing the performance as the frequency of updates increases will indicate the priorities that need to be addressed.

This approach will thus give rise to two new parameters: f_{update} and o_{update} which dictate how frequently and with how many observations from the supply respectively DynaPO updates should be performed.

Node Removal

Another aspect to consider is whether to update upon node removal and to what extent. In the previous dynamic work, if one of the selected FS nodes was removed an entire recalculation would occur to allow selecting the replacement FS node. In partial observation however, this would mean holding some observations in reserve to facilitate this and runs into many of the same problems as above. Three options are available to deal with this event:

No Update - do not update in this situation, instead using the (potentially outdated) node list from the previous update, selecting the next highest node on this list that is still present and is not already selected. This means no additional observations need to be used but, depending on how long ago the last update was, the selection may be substantially suboptimal.

Partial Update - perform an update but with fewer observations than a full update would be allowed. Specifically target likely candidate nodes for this update. For instance, target the highest non-selected nodes on the list from the previous update or nodes that are likely candidates for high influence. This helps to reduce the issue of having an outdated node list but also is conservative with using the limited supplies available.

Full Update - perform a full update, using the full number of observations that such an update would be allowed. Whilst this will ensure that up-to-date information about nodes is available, it runs the risk of running down the supply long before the next one is due, particularly if node removal happens frequently as it does in the models of González et al. [2006a] and Ichinose et al. [2013].

These different approaches may thus introduce/change the o_{update} parameter from above into $o_{fullUpdate}$ and $o_{partialUpdate}$ which represent the number of observations given to full updates and partial updates respectively. These may need to be variable, functions of the state of the simulation or graph.

Exploration Targets

As observations are highly limited, trying to target nodes we believe to be good candidates is important. Whilst certain approaches from the static algorithms can be used (exploring high degree neighbours, growing along the gradient, etc.) the dynamic nature can also be used to inform this process. In particular, re-exploration of known nodes, particularly those that are likely to have shifted or are actively being used will allow updating of the relevant rankings of nodes with little use of observations. This in turn will allow the selection of DynaPO to remain relevant without having to perform full updates.

Algorithm 3 DynaPO

procedure DynaPO(G_T, n, o_{sup}, f_{sup}, o_{upd}, f_{upd}, *updStyle*, p_{rex})
    create empty list, *selected*
    create empty max-priority degree queue, Q
    create empty map, N
    for all timesteps **do**
        if first timestep or f_{sup} timesteps since last supply **then**
            $obs \leftarrow o_{sup}$
        if first timestep or f_{upd} timesteps since last full update **then**
            perform full update: update($\min(o_{upd}, obs)$)
        else if selected node has been removed **then**
            Remove node from N, Q and *Selected*
            if *updStyle* is "full" and PerformFullUpdate? **then**
                perform full update: update($\min(o_{upd}, obs)$)
            else if *updStyle* is "partial" **then**
                perform partial update: update($\min(\text{PartialObs}(), obs)$)
        Update *obs* by removing the number of observations used
        selected $\leftarrow$ top n highest degree nodes as currently known
        if size of selected $< n$ **then**
            add nodes selected u.a.r. from rest of graph until correct number

function PerformFullUpdate?
    if time since last full update $< f_{upd}/2$ **then return** FALSE
    else return TRUE with prob. proportional to time since last full update's distance between $f_{upd}/2$ and f_{upd}

function PartialObs
    return $(o_{upd}/f_{upd}) \times$ (time since last partial or full update)

4.5.3 Dynamic Partial Observability Algorithm: DynaPO

Taking these concerns and considerations into account, in this section we present and describe the algorithm for influential node location detection in dynamic partially observable networks: DynaPO.

The core of the algorithm is presented in Algorithm 3 and has many distinct functions related to monitoring when updates should be performed, when the supply should be updated and dealing with node removal. The algorithm takes the following as arguments: the dynamic graph, G_T, which is needed to request neighbour lists; the number of locations to be found, n; the number of observation per supply, o_{sup}; the frequency of supplies, f_{sup} (we assume that t_start is when DynaPO will first be called and hence it can be omitted); the number of observations to use per update, o_{upd}; the frequency of full updates, f_{upd} which

Algorithm 4 DynaPO: Update

```
procedure UPDATE(o)
    reExplorationObs ← p_rex o
    newExplorationObs ← (1 − p_rex)o

    while known but not re-explored node exists and reExplorationObs > 0 do
        take next not re-explored node with highest degree, v
        NodeExplored(v)
        reExplorationObs ← reExplorationObs − 1

    newExplorationObs ← newExplorationObs + reExplorationObs
    if |Q| = 0 and |selected| < n and newExplorationObs > 0 then
        NodeExplored(v selected u.a.r. from unexplored part of graph)
        newExplorationObs ← newExplorationObs − 1
    while |Q| > 0 and newExplorationObs > 0 do
        v ← Q.pop()
        for all w in N[v] do
            if w not in N then
                NodeExplored(w)
                newExplorationObs ← newExplorationObs − 1
                if newExplorationObs = 0 then
                    add v to Q
                    return 0
    return newExplorationObs

function NodeExplored(v)
    N[v] ← G_T.neighbours(v)
    add v to selected if not already
    add v to Q
```

represents the maximum number of timesteps between full updates; the node removal update style, *updStyle*, which is one of "Full", "Partial" or "None"; the re-exploration proportion, p_{rex}, which is the fraction of observations per update that will be used for re-exploring already known nodes.

The main flow is as follows:

1. Create the list to store nodes in descending degree order, the queue of nodes that are next to be explored and the map that stores neighbour lists of explored nodes to avoid having to use observations each time.

2. Each timestep, check how long it has been since the last supply and update the number of observations with the new supply if necessary (remember that this removes the old, unused observations).

3. Check how long it has been since the last full update and if this is f_{upd}

perform a full update with o_{upd} observations (or all observations remaining in the supply if this is less).

4. If a full update is not performed but a node chosen as one of the n locations has been removed then perform a node removal update as dictated by $updStyle$.

5. After any update, remove the observations used from the remaining observations and resort the known nodes in descending order of degree to get the new top n locations. If there are not enough nodes known, which is most likely to occur due to no observations being available and nodes being removed from the graph, then supplement the list up to n locations with random nodes from the graph. The degree of these nodes remains unknown and they will be replaced any time the list updates. This is an undesirable situation, will negatively impact performance and should be avoided at all costs.

The node removal update strategy dictates which of 3 different approaches, each taking into consideration different aspects of the concerns raised in Section 4.5.2, should be used. These are detailed below.

None DynaPO will not perform any additional exploration upon node removal, instead simply removing the node from the list and shifting up those behind it. This approach is the most conservative with observations and relies on the full updates dictated by o_{sup} and f_{sup} to find new nodes and update the node list. This may mean that its information is outdated but it should also avoid using all the observations and being forced to rely on random placement.

Partial This approach will perform an update with fewer observations than would be expected in a full update. Rather than specify ahead of time how many observations should be used, DynaPO will scale the number based on how long ago the last update was performed and the size and frequency of the full updates of the system. Each timestep since the last update will accrue o_{upd}/f_{upd} observations. This approach should avoid the situation where multiple updates are performed in close temporal proximity to one another and observations are thus wasted in re-exploring recently updated nodes and exploring new nodes when the knowledge is fairly up-to-date. It should also avoid all observations being used up due to frequent node removal and DynaPO then having to rely on random placement.

Full In this approach a full update is performed using the normal number of observations, o_{upd}. However, similar to the Partial update approach, DynaPO attempts to reduce the likelihood of wasted or redundant updates due to close temporal proximity. It does this by refusing to perform another full update if we are within $f_{upd}/2$ timesteps of the last one. After $f_{upd}/2$ timesteps it will only perform the full update with probability directly proportional to where the current timestep places between $f_{upd}/2$ timesteps since the last update and f_{upd} timesteps since the last update. That is, it approaches probability 1 as we get closer and closer to when the full update would be performed regardless. This introduces an intelligent and dynamic aspect to when DynaPO performs full updates but also helps to minimise the risk of performing too many updates due to node removal.

The update functionality of DynaPO is shown in Algorithm 4 and shares many similarities with PO-Place. When exploring new nodes it builds a growing fringe of nodes to explore and attempts to explore along the gradient of higher degree nodes similar to other approaches such as [Brautbar & Kearns, 2010; Chen et al., 2009]. However the primary difference is that of node re-exploration which is necessitated due to the changing nature of dynamic networks. As nodes leave and rejoin, the degrees of already explored nodes are likely to change. In the approach of PO-Place these nodes are already marked as explored and so would not be considered again. Indeed, allowing them to be fully explored again would potentially lead to an infinite loop of only looking at the same few nodes. Instead we assign a proportion of the observations given to the update, p_{rex}, for the purpose of re-exploring known nodes. As discussed in Section 4.5.2 simply selecting nodes at random to re-explore would be wasteful as many nodes are unlikely to become highly influential locations. Instead we wish to focus on those nodes that we think are highly influential (in case they have become not so) and those we think might become highly influential (in case they have). To this end, DynaPO re-explores nodes in descending order of their last known degree. This allows for both desirable aspects to be achieved and, as only the nodes themselves are being re-explored and not their neighbours, a small number of re-exploration observations can traverse a large number of the high degree known nodes.

DynaPO Performance Evaluation

Having described the functionality and design of DynaPO we now seek to investigate its efficacy at finding high degree locations. Examining which type of node removal update style has the greatest effect and how the combinations of o_{upd} and f_{upd} affect the performance will allow us to focus on these when using DynaPO for convention emergence.

To facilitate this study we focus on examining the performance of DynaPO in the González model [González et al., 2006a]. Due to DynaPO's similarity to PO-Place and the Ichinose model's similarity to scale-free topologies we look to investigate the performance of DynaPO on the more complex González model and then transfer this knowledge to Ichinose for convention emergence. We use the settings most used in the González model work in Chapter 3 and described in Section 2.6: an arena size of 42, $v_0 = \bar{v} = 0.3$, $TTL_{max} = 500$, particle radius of 0.1 and a simulation rate of 100. We produce graphs of 1000 vertices and, as before, our González model automatically adjusts this up in order to have a largest WCC in the underlying graph that is approximately 1000 vertices. The graph has a burn-in of 1000 timesteps.

To evaluate the performance we use 4 different settings looking to investigate each of the different settings and their effect on DynaPO. Unless stated otherwise all 4 have $f_{upd} = 200$, $f_{sup} = 1000$ and $o_{sup} = 1000$. All have $p_{rex} = 0.33$:

None $o_{upd} = 200$, $updStyle$ = "$None$" - allows us to study the conservative approach, evenly spaced updates between supplies and no additional observations used between updates.

Partial $o_{upd} = 100$, $updStyle$ = "$Partial$" - examines the effect of allowing partial updates. Early simulations showed that higher o_{upd} resulted in the observations of o_{sup} being used up too quickly. As we are concerned with the actual performance of partial updates we set it lower.

Full, Small $o_{upd} = 20$, $f_{upd} = 20$, $updStyle$ = "$Full$" - small and consistent full updates with little space between them for additional full updates to use all the observations. This approach is in many ways similar to partial updates and will be compared to it.

Full, Large $o_{upd} = 200$, $updStyle$ = "$Full$" - standard setup, fewer big updates which will allow lots of full updates in between if necessary. Seeing the effect this has is the primary purpose of this setting.

For each setting we request $n = 10$ locations and run each 30 times, combining the results. *deg-sum* and 1-Hop data was collected for both degree and DynaPO placement for every timestep from 1 to 2000. For better baseline purposes, degree placement was run alongside DynaPO simultaneously on the same graph so that they could be directly compared. Due to the constantly changing nature of the topologies, D_{Base} as a metric is not as useful and is not monitored.

The results of this evaluation are shown in Figure 4.16 as a series of violin plots [Hintze & Nelson, 1998], one for each setting. Violin plots are an extended form of box plot which contain the same information (values of the median/mean, quartiles and overall range) but additionally combine it with a

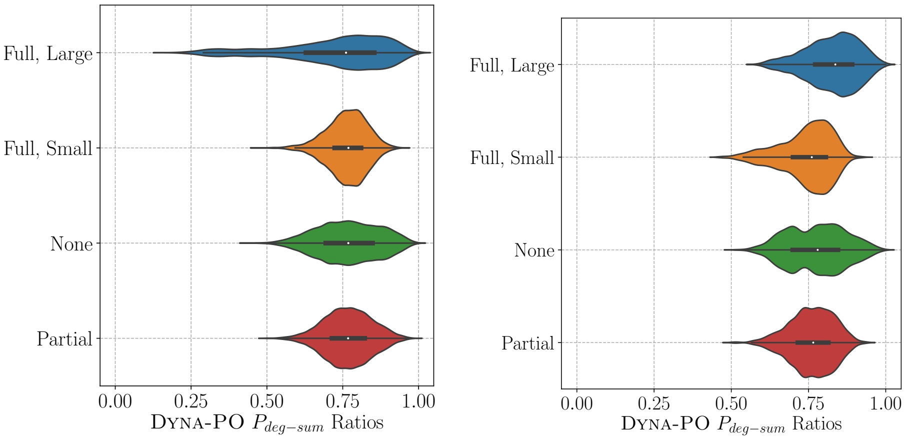

(a) $0 \leq t \leq 2000$

(b) $0 \leq t \leq 200$

Figure 4.16: Violin Plots of $P_{deg\text{-}sum}$ ratios for DYNAPO in González networks

kernel density estimate which highlights the distribution of the dataset. The $P_{deg\text{-}sum}$ scores for every timestep in each run is combined to give a large-scale overview of the performance and typical ranges that each approach scored in Figure 4.16a.

As can be seen, the means for each setting are approximately the same at just under 0.8. It is thus standard deviations and the shape of the violin plots that indicate the differences in the performance of the settings. Full, Large is by far the worst performing setting despite having the highest maximum $P_{deg\text{-}sum}$ scores of the group. This setting's worst performances look to occur not infrequently given the shape of the kernel density plot and are substantially worse than any other setting. This was expected as the full update style and high value of o_{upd} mean that the algorithm is likely to use all its observations fairly quickly and then have to rely on random placement which will perform very poorly in comparison. Full, Small, performs much better with a narrow range of values and most of them clustered around the mean. Using the constant, small, iterative approach looks to be beneficial, not detrimental and this is supported further by the performance of Partial being comparable to Full, Small. Both have quite pronounced clustering around their means indicating that they

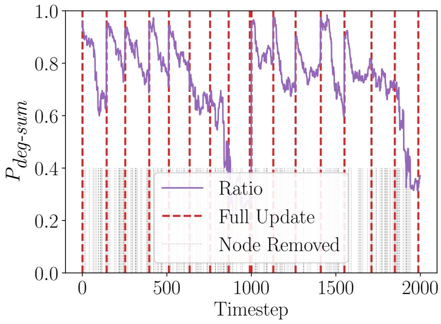

(a) Full

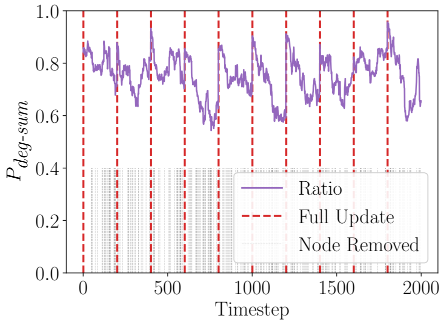

(b) Partial

Figure 4.17: $P_{deg\text{-}sum}$ ratios over time for DynaPO in González networks showing updates and node removals

reliably have $P_{deg\text{-}sum}$ scored of around 0.8. None has a similar range to Partial and Full, Small but less clustering. Instead it is more spread out to the upper and lower ends of its effective range, indicating less reliability but better performance at certain points.

Figure 4.16b shows the same plots but this time limiting $P_{deg\text{-}sum}$ scores to only those in the first 200 timesteps. This allows us to see the performance of DynaPO at the early stages of the simulation where influencing convention emergence is important. Whilst most of the settings performances are broadly comparable, Full, Large performs substantially better when viewed over this period. This is likely due to it frequently updating its data which allows reliable location selection but causes the issues shown in Figure 4.16a once the observations have been used up because of this. Similar behaviours and relationships are seen for $P_{|1\text{-Hop}|}$.

Given this disparity with performance at the beginning of the simulation and over the whole simulation, Figure 4.17 shows the performance of DynaPO in typical runs for both Full and Partial update methods. As can be seen, the peaks for Full are much higher than for Partial and the full update lines indicate how frequently they occur in this update methodology. However in the time span just before the new supply arrives, the performance of Full update drops precipitously as all observations have been used. Whether this early performance boost is important in convention emergence will be examined in the next section. Figures 4.17a and 4.17b both also highlight the importance of ensuring that node removal updates do not use too many observations by showing the frequency with which they occur.

Convention Emergence using DynaPO

Having established that the performance of DynaPO can reach good levels of approximation for pure degree placement we now look to use this for placing IAs and encouraging convention emergence.

To evaluate DynaPO in this domain we use the González settings from the previous section and additionally utilise Ichinose dynamic networks [Ichinose et al., 2013] with $m_0 = 4$, $m = 3$ to test whether the good performance in the González model is also applicable to the different dynamic nature of Ichinose graphs.

The experimental setup for the convention emergence runs in dynamic topologies is the same as that in the real-world and synthetic topologies explored earlier: the 10-action coordination game with $p_{explore} = 0.25$ and agents using Q-Learning. 100 runs were performed for each of the values of n for each placement setting and the proportion of these results that emerge the desired convention is calculated. As a baseline we also include degree placement which is updated in a static manner (see Section 3.5).

Establishing what proportion of the graph is being observed is much harder than in the static topologies. Simply comparing the number of observations provided to the algorithm is problematic as this is reliant on the length of the simulation. Instead we must find a relative measure, one that can expressed as a per X timesteps value. One approach is to consider the maximum *possible* number of observations that DynaPO could be allowed to make use of. This would be the use case where the algorithm was observing every node, every timestep. Given the approximate size of the dynamic graphs that are being examined this would mean that, for every 1000 timesteps, the maximum number of possible observations would be 1000 timesteps $\times$ 1000 vertices = 1000000 possible observations. As the settings we are using only supply DynaPO with 1000 observations per 1000 timesteps this would mean that DynaPO is only observing 0.1% of the network.

However, this analysis is very liberal in its assertion of equivalent observations as the degree placement metric is not "updating" in nature (as examined in Chapter 3) and so is not observing every node every timestep. A fairer measure would be to compare the number of "observations" that the degree metric uses to the number of observations that DynaPO uses. In the Static approach we are discussing, degree only updates when a node it has selected is removed. This means that the number of node removals that should be counted as triggering an update is dependent on n. Whilst having a dependency like this is not desirable, due to not being invariant across runs, it provides the best measure for a comparison on how many observations degree placement makes compared to DynaPO.

To allow comparison we define the *degree equivalent observations*, o_{deg}, for a network as:

$$o_{deg} = (\text{number of node removal events}) \times |V| \tag{4.5}$$

As each time a selected node is removed the degree placement metric is recalculated across the entire graph, o_{deg} is the equivalent number of observations that would be needed to enact the same effect. Having calculated the degree equivalent observation for a given graph and value of n over some arbitrary time window we can then compare the number of observations that are available to DynaPO over that same time window and establish what percentage of the network DynaPO is observing compared to pure degree placement.

Table 4.5 shows the average number of node removal events for the various dynamic graphs when $n = 10$. As this is the largest number of IAs used in the

	Ichinose				Gonzalez
	RP	RR	TP	TR	
Average NRE	9.97	10.13	843.23	657.33	96.43
StdDev NRE	2.48	3.29	34.80	41.13	8.88
o_{deg}	9,966.67	10,133.33	843,233.33	657,333.33	96,433.33
DynaPO % obs	10.03	9.87	0.12	0.15	1.04

Table 4.5: Node removal events and degree equivalent observations for dynamic topologies when $n = 10$.

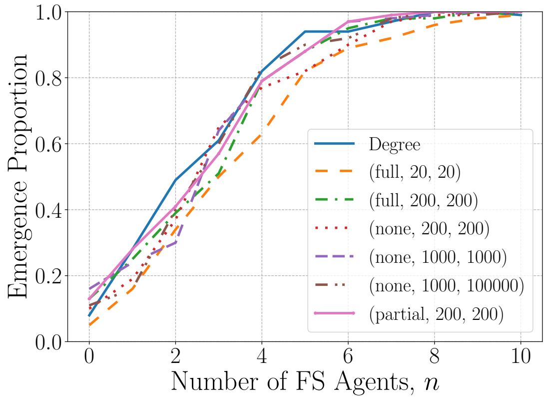

Figure 4.18: Convention Emergence using DynaPO in the González Model

convention emergence runs this represents the worst-case scenario in terms of observation percentage as the number of node removal events will be highest as there are more nodes whose removal would trigger such an event. The number of node removal events over 1000 timesteps was monitored for 30 runs in graphs of $|V| = 1000$ and the resulting average is shown in the table. We can then calculate o_{deg} and hence, the equivalent percentage of the network that DynaPO observes. The low number of node removal events in Ichinose-RP and Ichinose-RR topologies results in the observation percentage being an order of magnitude larger for these topologies but still no higher than approximately 10%.

With a way to quantify the percentage of the network being observed we now turn to the results of the convention emergence. Figure 4.18 shows the effectiveness of DynaPO in encouraging convention emergence in González networks.

The different types of DynaPO being used are similar to those explored in Section 4.5.3 but with a few additional entries. $(None, 1000, 100000)$ represents a bulk supply as discussed earlier in this chapter; a single supply and immediate utilisation of available observations with no others appearing later. In this scenario the algorithm uses all the observations at once which is equivalent to PO-Place with basic settings being used at the beginning of the simulation and then not updated at all. $(None, 1000, 1000)$ is a similar approach, using all observations at the beginning of the supply and then not updating until the next supply. This is included to see if there are differences in efficacy between it and $(None, 1000, 100000)$. It appears however that finding the perfect pure degree placement at the beginning is sufficient to counteract the lack of later supplies.

All the settings chosen have the same number of observations per 1000 timesteps, 1000, and hence represent the same observation percentage, approximately 1%. How they use those observations in that time frame is the variable to monitor.

All settings exhibit close approximations of the pure degree placement despite only observing a small fraction of the temporal network. Indeed all of them bar one have average proportion performance, compared to the value of degree at the same n, of higher than 90%. The only one that performs worse and is statistically significantly worse (95% CI one-tailed z-test) than degree at a majority of values of n is $(Full, 20, 20)$. This is likely due to the low amounts of observations available to it at the important early stages of convention emergence.

The same settings and their effectiveness in the various Ichinose networks is presented in Figure 4.19. Similar levels of effectiveness are found in these topologies as well with $(Full, 20, 20)$ again being the only setting of DynaPO that is significantly worse than the pure degree placement. With the targeted node removal nature of Ichinose-TP and Ichinose-TR [Ichinose et al., 2013] the risk of DynaPO being unable to accommodate with the rapidly changing set of

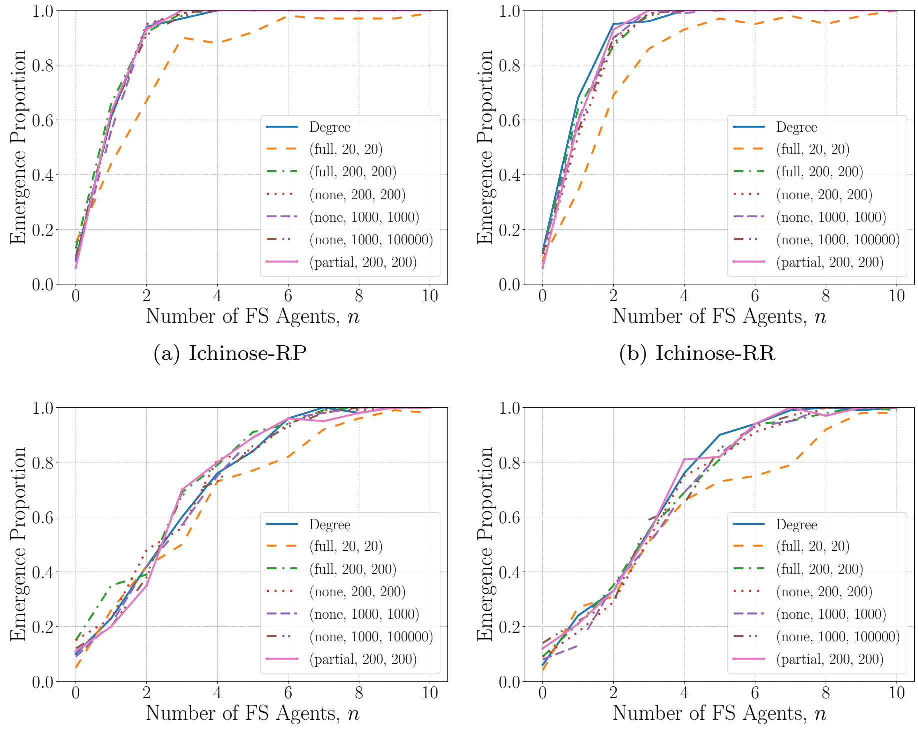

(a) Ichinose-RP (b) Ichinose-RR

(c) Ichinose-TP (d) Ichinose-TR

Figure 4.19: Convention Emergence using DynaPO in the Ichinose Models

high degree nodes was higher than in the other dynamic topologies. However, it is able to easily keep pace with the pure degree placements whilst observing less than 0.15% of the number of nodes that the pure degree placement strategy must do.

We have now evaluated the performance of the potential variants of DynaPO whilst purely concerned with their performance in finding locations as well as when using DynaPO to encourage convention emergence. We have shown that, with a very small percentage of the observations that fully-observable degree placement requires, DynaPO exhibits the same level of performance. This is the case for both González and Ichinose dynamic topologies and indicates that DynaPO should be generally applicable for encouraging convention emergence in partially observable dynamic networks.

4.6 Late Interventions with Limited Observations

We have shown that it is possible to effect convention emergence under the constraint of partial observability in both static and dynamic networks. We have introduced two algorithms, PO-Place and DynaPO, which are able to encourage convention emergence to a desired action choice to similar levels as using placement by degree whilst observing much smaller fractions of the networks to do so.

In this and the following section we now investigate the performance of these two algorithms in causing *destabilisation* of existing conventions as we did in Chapter 3. Being able to efficiently replace established conventions with a different one of the designer's choice whilst only having to have limited knowledge of the network allows external actors to make use of these results. The work in Chapter 3 required and assumed full network knowledge and hence is more applicable to those who already have high levels of access to the system.

We begin by using the settings that performed well in Section 4.4.1 to elicit destabilisation of conventions in the real-world networks. Given the positive effect of increasing the number of starting locations, s, in the Enron and Twitter networks we again use both $s = 1$ and $s = 9$ at various levels of observation to establish whether this continues to have a positive effect on the performance of PO-Place.

As discussed before, the real-world networks are resilient to conventions emerging compared to synthetic networks of similar size and density. They require both larger numbers of IAs and longer periods of time to effect the same level of convention emergence. Due to this, and our desire to simply establish the difference in performance between the different settings of PO-Place and degree placement, we again utilise the 80% Kittock Criterion, $K_{80\%}$, as the threshold at which we consider a convention to have emerged. In this instance we consider the dominant strategy to have been destabilised and replaced when another strategy achieves this level of choice amongst the population of agents.

Conventions infrequently emerge unaided within these networks, even at the $K_{80\%}$ level (see the proportion emerging with $n = 0$ in Figure 4.6). When they do so, this also takes substantial time. Given this, and as we wish to focus on destabilisation, we force a convention to rapidly emerge by saturating the networks with initial IAs in the same way as occurs in initial intervention. In each network we place 500 IAs at timestep 0 and leave them in the network long enough to be assured that a convention has emerged to the $K_{80\%}$ level. Initial simulations find that this occurs by timestep 500 for the CondMat network and timestep 1500 for the Enron and Twitter networks with probability 1 (based on 100 runs). Whilst this artificial saturation differs from our approaches in other network types (where we let the network emerge the initial convention unaided) it allows rapid and robust creation of the convention to be destabilised and differs in no noticeable way from the naturally emerged conventions.

We evaluate the performance of PO-Place in late interventions using the same agent simulation as used in all previous sections of this chapter. In each simulation we select the 500 highest degree nodes at timestep 0 and use them for the saturation IAs discussed above. The simulation is run for the length discussed previously to allow the convention to emerge to the desired level. The saturation IAs are then removed and a new intervention of the desired type, degree or PO-Place, is used to select the IAs for destabilisation. These new IAs are then assigned a strategy, chosen uniformly at random, from those strategies that are not the current dominant strategy. The simulation then continues until timestep 15000 with the levels of adherence to each strategy monitored. If destabilisation is likely to occur it predominantly does so within this window. For those simulations where destabilisation would have occurred given a longer time frame, they will be marked as failures due to this cut-off. However, as this cut-off is applied equally to all runs within a given graph, and as we are primarily concerned with the *relative* performance of different approaches this does not affect the conclusions drawn regarding algorithm performance within

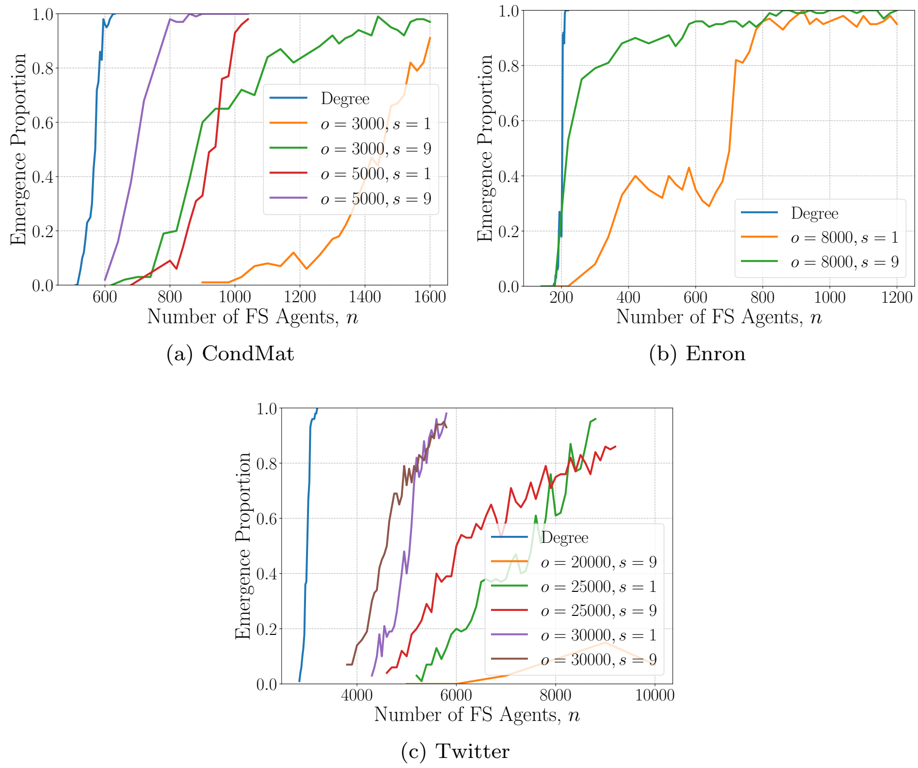

Figure 4.20: The proportion of runs switching away from the established convention during late intervention in the real-world networks.

the range of values chosen and within this stated time frame. For each set of parameters, the simulation was performed 100 times and the proportion of runs that successfully caused the dominant strategy to be replaced by the selected one to the $K_{80\%}$ level was recorded. Figure 4.20 shows select performances of PO-Place compared to pure degree placement in this regard.

As found in Chapter 3, unlike in initial intervention the real-world networks need varying and distinct numbers of IAs to guarantee destabilisation even when using degree placement: CondMat destabilises with ~2.81% IAs, Enron with ~0.59% and Twitter with ~3.94%.

The number of observations PO-Place requires to enact consistent destabilisation is substantially higher than those required to have close to baseline performance during initial intervention. This is primarily due to the large num-

Graph	n or o and equivalent percentage					
CondMat	600	800	1000	1600	3000	5000
	2.81%	3.74%	4.68%	7.49%	14.04%	23.4%
Enron	200	400	600	800	8000	
	0.59%	1.19%	1.78%	2.37%	23.74%	
Twitter	3200	5800	9000	9200	25000	30000
	3.94%	7.13%	11.07%	11.32%	30.75%	36.9%

Table 4.6: Differing numbers of nodes or observations and their percentage equivalent for the real-world networks

bers of high-degree nodes required. Not only must PO-Place find a significant number of nodes with limited observations in order to reach comparable performance to its use in initial intervention, it must find a much larger number of high degree nodes within these found nodes. For instance, with $o = 5000$ in the CondMat network PO-Place requires 12% of its observations just to add the 600 highest-degree nodes required to its knowledge base, let alone actually find them. In contrast, in initial intervention, even at $o = 500$ only 2% of its observations were required to actively add the highest-degree nodes if it located them. As even pure degree placement requires high numbers of IAs in all 3 networks, similar patterns of much larger proportions of the available observations being needed just to be able to select these nodes, even after finding them, persist. This contributes to the difference in performance exhibited.

However, even with this limitation, PO-Place performs well in all 3 networks. Whilst requiring more IAs to cause the same level of destabilisation it does so whilst requiring observation of only a fraction of the network. The percentages of the network that the various numbers of nodes (or observations) are equivalent to is shown in Table 4.6 for easier parsing.

For instance, in the CondMat network, when observing 14.04% of the network PO-Place requires 1600 IAs compared to degree placement's 600 to effectively guarantee destabilisation and replacement of the dominant convention. This represents an additional 4.68% of the CondMat network, a small percentage increase given the fraction of the network that has to be observed. Increasing the percentage observed to 23.4% substantially decreases this difference, with PO-Place requiring 1000 IAs to perform as degree placement does – an increase of only 1.87%, a relative drop of 60%, an absolute one of 2.81%.

This trade-off, increasing the percentage of the network observed to reduce the increase in number of IAs required, varies in effectiveness between each of the networks. For instance, in the Twitter network an increase in observations from 25000 to 30000 (an increase of 6.15% of the network size) allows a reduction in number of IAs from ~9200 to ~5800 for the same level of performance – an absolute reduction of 4.19%, 37% relative. This contrasts with CondMat where the number of observations had to increase by 9.36% of network size to cause an absolute reduction of only 2.81%. We can thus see that underlying features of networks heavily influence the amount gained when increasing the number of observations available. Analytically understanding this is beyond the scope of this thesis but it is sufficient to say that this is a factor that must be considered and expected when deciding what number of observations to supply PO-Place for destabilisation.

In each of the topologies however, increasing the value of s has a marked and highly beneficial effect. With $s = 9$ we see increases in the effectiveness of PO-Place throughout. In particular, increased s dramatically increases the performance of PO-Place at mid-levels of destabilisation. For instance, in the CondMat network with $o = 3000$ increasing s can reduce the number of IAs needed by up to 600 to elicit the same level of performance, 2.81% of the network size. Similarly, in the Enron network, reductions of almost 500 IAs is shown, 1.48% of the network size. Indeed, for Enron, this makes a substantial difference in performance with $o = 8000, s = 9$ performing as well as pure degree placement for much of the destabilisation proportions compared to a markedly worse performance for $o = 8000, s = 1$. Whilst the relative differences it causes in the Twitter network are less marked, it still exhibits improved performance with the higher number of concurrent searches. Overall this indicates that this

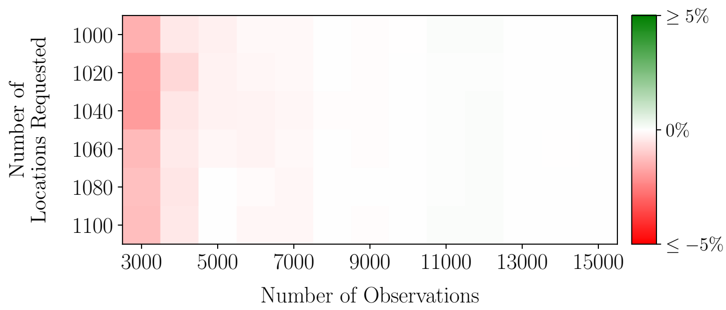

Figure 4.21: Statistically significant $P_{deg\text{-}sum}$ differences between $s = 1$ and $s = 9$ for late intervention placement in the CondMat network.

disjoint search is beneficial for causing destabilisation in the real-world networks. This is likely caused by two factors: (1) splitting the observations increases the likelihood of finding high-degree nodes, or (2) splitting the areas where the nodes are selected from into distinct sections of the graph increases the effectiveness of the IAs in causing destabilisation.

To test both of these hypotheses we explore the pure placement performance of PO-Place at the ranges of n and o that are causing destabilisation. Figure 4.21 shows this for the CondMat network for values of n where the greatest difference in performance is exhibited in Figure 4.20a and for a wider range of o. Each square is coloured based on the difference in $P_{deg\text{-}sum}$ when $s = 1$ and $s = 9$ for the given n and o, with white meaning that there was no statistically significant difference (95% Mann Whitney U significance test based off of the average $P_{deg\text{-}sum}$ of 30 runs at each point). As can be seen, for most values of o there is little if any difference in performance. For the values of o that are represented in Figure 4.20a however, the $P_{deg\text{-}sum}$ performance is consistently worse with increased s. This means that option (1) from above is thus very unlikely and instead means that the disparate nature of the nodes selected by PO-Place is most likely the largest contributing factor.

Considering the best performing options presented in Figure 4.20 gives us the following levels of performance for each network. In CondMat, by observing 23.4% of the network PO-Place can achieve destabilisation whilst requiring

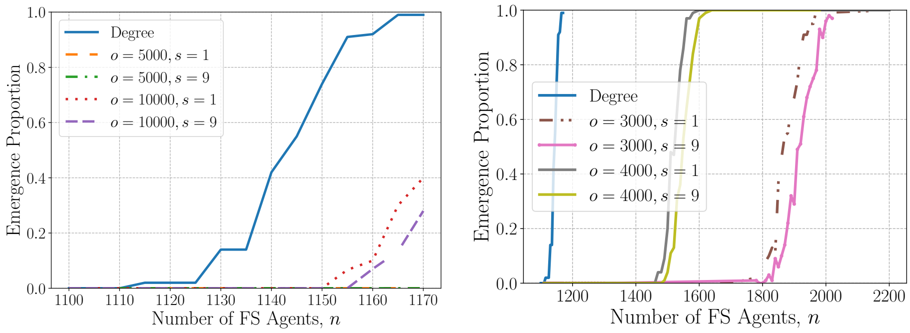

(a) (b)

Figure 4.22: The proportion of runs switching away from the established convention during late intervention in scale-free networks of 20000 nodes. (a) focuses on the area where degree placement exhibits its change whilst (b) does the same for PO-Place.

only 0.93% more of the network as IAs than degree placement does. In Enron, for a destabilisation probability of 95%, PO-Place requires an additional 0.6% of the network as IAs when observing 23.74% of it. In Twitter, for a destabilisation probability of 95%, if able to observe 36.9% of the network, PO-Place can elicit this probability whilst requiring an additional 3.19% of the network as IAs.

Overall, we can state that PO-Place performs well when used to find nodes for destabilisation in real-world topologies. Whilst these parameters are likely not optimal, nor necessarily the best trade-offs available, we have shown that PO-Place can provide the same level of performance as pure degree placement by using less than 1-5% more of the network as IAs and whilst only observing 25-35% of said network.

4.6.1 Synthetic Networks

Having explored the effectiveness of PO-Place for late destabilisation in the real-world topologies we now, as before, similarly explore the performance in synthetic networks. We focus on generated scale-free and small-world topologies as before and both of these are generated with $|V| = 20000$ and $|E|/|V| = 9$.

The settings used to generate these topologies are otherwise identical to those used in Section 4.4.2.

Unlike in real-world topologies, conventions emerge unaided with ease in the synthetic networks and consistently reach the 90% Kittock Criterion of a convention, $K_{90\%}$. We thus, instead of artificially saturating the network as before, allow the conventions to emerge naturally and use this higher Kittock threshold when establishing that conventions have emerged and when they have been replaced. The simulation setup is otherwise identical to that of previous sections. We find that initial conventions have become established in the network with near certainty by timestep 1000 and so we introduce our IAs into the system at this point with the strategy chosen uniformly at random from those that are not the current established convention. When monitoring the destabilisation and replacement of these conventions we find that this happens, if it is likely to happen at all, by timestep 10000. As with the real-world simulations, any destabilisation that would occur after this time is marked as a failure but as we are concerned with the relative performance of our approaches, this does not affect the conclusions drawn in a substantive manner. All parameter settings were run for 100 simulations and the proportion causing destabilisation and replacement to the desired convention is noted.

We display the results of this process for scale-free networks in Figure 4.22. The first thing to note is the high percentage of nodes required to cause destabilisation even when placing by degree. As in the real-world networks, where destabilisation occurred with at most 4% of the network as IAs, scale-free networks require around 6% of the network to be IAs to facilitate destabilisation, despite being closest in size to the CondMat network which only required 2.81%. The range over which this change occurs, from destabilisation occurring with probability 0 to occurring with probability 1, is only 50 nodes which is 0.25% of the network size. This narrow range during which the transition occurs indicates that there is a critical "tipping point", more noticeable here than in many of the real-world networks.

Figure 4.22a focuses on the range of IAs during which degree placement undergoes this transition. In particular, and different from many of the real-world networks particularly Enron, this figure shows that even relatively high numbers of observations (in this case $o = 10000$, 50% of the network) do not allow PO-Place to closely approximate degree placement. Whilst we start to see change with this number of observations it is still only minor and indicates that the scale-free network requires much larger proportions of the network to be observed compared to the real-world networks, even those of comparable size. This highlights the underlying differences between real and synthetic networks as described by Franks et al. [2013].

Figure 4.22b looks at a wider range of n and shows the performance of PO-Place using observation proportions closer to those needed in the real-world networks. For $o = 3000$, representing observing 15% of the network, destabilisation is guaranteed at approximately $n = 2000$, 800 more than pure degree placement which represents 4% of the network size. $o = 4000$ (20% of the network) causes destabilisation with $n = 1600$, an increase of 2% of the network size compared to degree placement. This is comparable to the performance differences exhibited in the real-world networks but the effect of increasing observations is better than the increases seen in many of the real-world networks. Given the performance of $o = 10000$, these increases are likely ones of diminishing returns but the improvement from such a small increase in observation percentage, whilst still being a small percentage of the network observed overall, shows the effectiveness of PO-Place.

The major difference between the performance of PO-Place in real-world and scale-free networks is the effect of increasing s. Whereas in the real-world topologies increasing s from 1 to 9 always produced improvement, with the change in both CondMat and Enron being substantial, in the scale-free network it always causes a, however minor, *decrease* in performance. This occurs at all levels of observations shown in Figure 4.22 and so is a wide-ranging effect of PO-Place in this network.

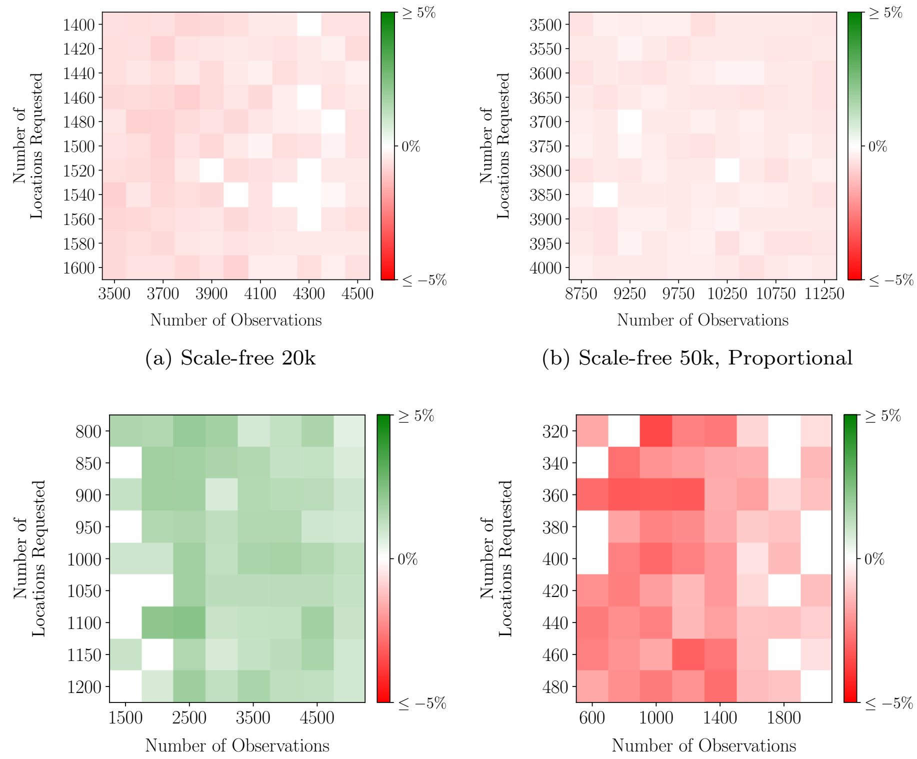

(a) Scale-free 20k

(b) Scale-free 50k, Proportional

(c) Scale-free 50k

(d) Scale-free 20k, Proportional

Figure 4.23: The statistically significant differences in $P_{deg\text{-}sum}$ performances for the scale-free networks at scales appropriate for late intervention.

Given the notable increases that s caused in Figure 4.12 this result is counter-intuitive. To establish how PO-Place is performing, independently of the convention emergence, we evaluate the statistical differences between its $P_{deg\text{-}sum}$ performance when $s = 1$ and $s = 9$. Figure 4.23 shows these differences and represents the same approach as Figure 4.12 with the average of 30 runs representing each point. Figure 4.23a shows the $P_{deg\text{-}sum}$ performance changes at values of n and o where PO-Place exhibits change in Figure 4.22. As can be seen, the increase in s causes worse $P_{deg\text{-}sum}$ performance nearly universally, more so than for CondMat as seen in Figure 4.21. Unlike in CondMat and the other real-world networks this performance deficit is not accompanied by a better performance in destabilisation and this indicates that in scale-free networks having disjoint searches is less important than finding the higher degree nodes. This is likely due to the very nature of scale-free networks with their emphasis on preferential attachment and the power-law distribution which creates the "hub" like structure [Barabási & Albert, 1999].

Given the positive increases found in Figure 4.12, we seek to establish at what values of n and o this effect diminishes. Given that Figure 4.12 is the performance difference in scale-free networks with 50000 nodes we wish to check whether the equivalent area in the parameter space of that size of network exhibits better performance. This is shown in Figure 4.23b which is the proportionate area of the 50k scale-free network parameter space. That is, the ranges of n and o used here are the same percentage of the network size as those in Figure 4.23a. The $P_{deg\text{-}sum}$ performance difference here is very similar however, with nearly universal detriment from increasing s, indicating that this is not a facet of simply the size of the network but rather the scales of n and o that are being used at this point.

However, knowing that there is a point in both 20000 and 50000 node scale-free networks where increasing s has a substantial and significant benefit, we seek to find where such a benefit still exists. This is shown in Figure 4.23c which represents an area of the 50000 node scale-free network's parameter space closer to that shown in Figure 4.12. As can be seen, at this range of n and o there are substantial benefits to increasing s from 1 to 9 with almost universal improvement in the $P_{deg\text{-}sum}$ scores achieved by PO-Place. Translating this proportionately to the 20000 node scale-free topology is shown in Figure 4.23d. In this figure the performance is in direct contrast, with the $P_{deg\text{-}sum}$ consistently affected negatively by the increase in s. Indeed, the level of the negative effect is far more than in Figures 4.23a and 4.23b. This indicates that graph size is indeed a partial factor for what ranges of n and o increasing s is beneficial but Figures 4.23a and 4.23b show that the scales needed to cause destabilisation are beyond where this transition takes place for both sizes of scale-free network. Figure 4.23d further supports this theory with the values of n and o being far

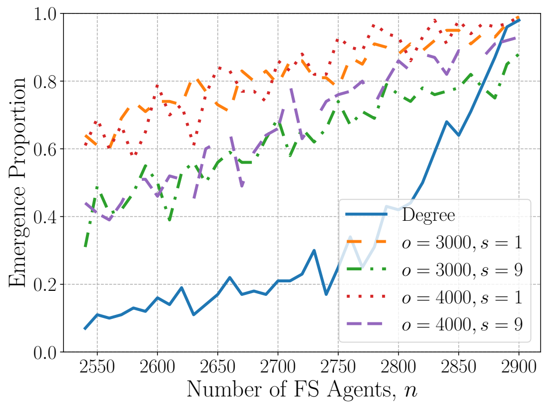

Figure 4.24: The proportion of runs switching away from the established convention during late intervention in small-world networks of 20000 nodes.

below those that are shown to effect destabilisation in the 20000 node scale-free network. Overall, these findings suggest that increases to s are not a panacea that can be applied universally to all types of network, although the detriment of doing so is only small in Figure 4.22. It hints that there are underlying facets of the topologies which make the application of multiple concurrent starting locations a boon or not. Finding these features is an area that future work should address to allow determination of how to best apply PO-Place but is beyond the scope of this thesis.

We also utilise PO-Place for late intervention in the 20000 node small-world network, the results of which are shown in Figure 4.24. The behaviour here is substantially different from all other types of topology investigated in this chapter and suggests even further benefits that might be gleaned from PO-Place when applying it to convention emergence and destabilisation. As can be seen, small-world networks require a much larger percentage of the graph to be IAs in order for destabilisation to occur (only being guaranteed at $n = 2900$ which is 14.5% of the network) and the change occurs over a much larger range of values than the sharp transitions of the real-world and scale-free topologies. The most striking feature however is that PO-Place does not just match the performance of degree placement, it substantially outperforms it over this range of n. As degree is well-understood to be a good measure of influence and as this effect is not present in initial intervention or in destabilisation in other topologies it is unlikely that PO-Place has found objectively "better" locations in this sense. Instead, we hypothesise that the localised nature of the choices available to PO-Place act beneficially within small-world topologies. Rather than placing the IAs at high degree locations that are spread out and distant from one another, the growing fringe approach of PO-Place here means that the locations will necessarily be closer to one another, forming a localised cluster of agents of which a high proportion have been selected. The fact that n is very close to o further necessitates that this is the case and indicates that nearly all agents in this cluster will be IAs. This localised cluster of agents are all forced into the same strategy because of this and hence form a self-reinforcing region as discussed by [Villatoro et al., 2011a]: an artificial meta-stable subconvention. This region is able to spread its influence without risk of becoming converted itself and hence is better at causing destabilisation than the isolated regions of IAs created in degree placement. If this were true we would expect increasing s to have a negative effect due to splitting this self-reinforcing region and this is what we see for both values of o in Figure 4.24, lending further support to this theory. Overall however, whatever the underlying reason, it is the case that, for small-world topologies, PO-Place performs markedly better than degree placement when attempting to destabilise established conventions.

We have now shown that PO-Place continues to be effective at facilitating destabilisation in both types synthetic networks. In scale-free networks, when $s = 1$ PO-Place performance is comparable in its achievements to its use in real-world topologies. It is able to, with around 20% network observation, consistently cause destabilisation of established conventions with only 2% more of the network as IAs. However increases in s are detrimental to the performance of PO-Place in scale-free topologies, in contrast to its effect in the real-world networks. In small-world topologies, PO-Place is able to outperform degree placement for the purposes of destabilisation. The underlying mechanism is unknown but is likely due to the features of small-world topologies, the larger number of IAs needed to elicit destabilisation and the nature of PO-Place.

4.7 Late Interventions in Partially Observable in Dynamic Networks

Finally, we turn our attention to destabilisation in dynamic networks under the constraint of partial observability. Having shown that PO-Place is able to perform well in this domain, we now seek to do the same for DynaPO.

We use the same general simulation settings as all other sections: the 10-action coordination game, with Q-Learning. All settings are as before. As with the synthetic networks, conventions are able to emerge without encouragement within the dynamic networks rapidly and robustly, reaching the 90% Kittock Criterion unaided. The time required for this to occur in the dynamic networks examined was found to be within 1000 timesteps consistently and so this timestep is used for the start of the late intervention. Similarly, the length required for destabilisation to occur was found to fall consistently within 5000 timesteps for both González and all Ichinose models and so this cut-off is used as the threshold of time for destabilisation. The topology models used are the same as in Section 4.5.3 to allow direct comparison and as an established set of networks which are known to facilitate convention emergence. Each simulation was run 100 times and the proportion which cause the desired destabilisation and replacement calculated. Static Degree is also included and is used as a baseline to compare against.

The results for the González runs are shown in Figure 4.25. As with the majority of other networks, this requires ~5% of the network to be IAs in order to consistently cause destabilisation. The DynaPO runs with the same settings as used previously are shown as well. Overall, the exhibit similar performance to one another with destabilisation guaranteed between 72 and 80 IAs regardless of settings. However, there is a clear ranking between them with

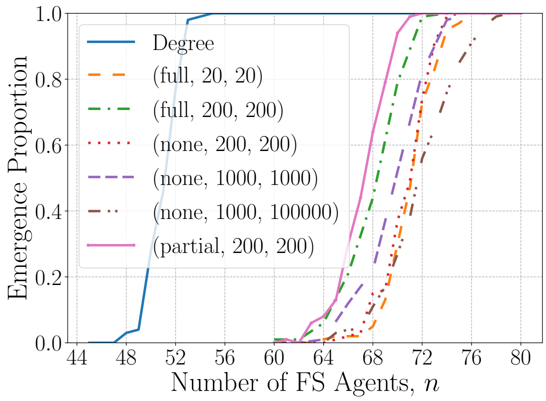

Figure 4.25: Late intervention and destabilisation in González networks.

$(Full, 20, 20)$, $(None, 200, 200)$ and $(None, 1000, 100000)$ all performing worse than the others. This is to be expected, as we have previously shown that having up-to-date information is important for destabilisation in dynamic networks and these settings lack this criteria due to either not updating enough at a time $[(Full, 20, 20)]$, not having up-to-date information as node removals occur $[(None, 200, 200)]$ or not having up-to-date information beyond the initial search $[(None, 1000, 100000)]$. Each of these represent the extremes of these positions but it is interesting to note $(None, 1000, 1000)$, which could be expected to suffer from similar problems, outperforms the other 3 settings. A balance between finding very good locations at the beginning of the intervention (due to being able to effectively view the entire network at that timestep) and ensuring that information is up-to-date seems to be reached in this setting. The settings which encourage up-to-date information in one of two ways, $(Full, 200, 200)$ and $(Partial, 200, 200)$ both outperform all of these with $(Partial, 200, 200)$ doing best overall likely due to the consistent need for accurate information failing in $(Full, 200, 200)$ towards the end of each supply due to lack of observations as detailed in Section 4.5.3.

Overall there is little difference in the performance levels of these approaches which is likely a factor of graph scale. However, the clear and consistent ranking between them highlights that the approaches which optimise for continuous

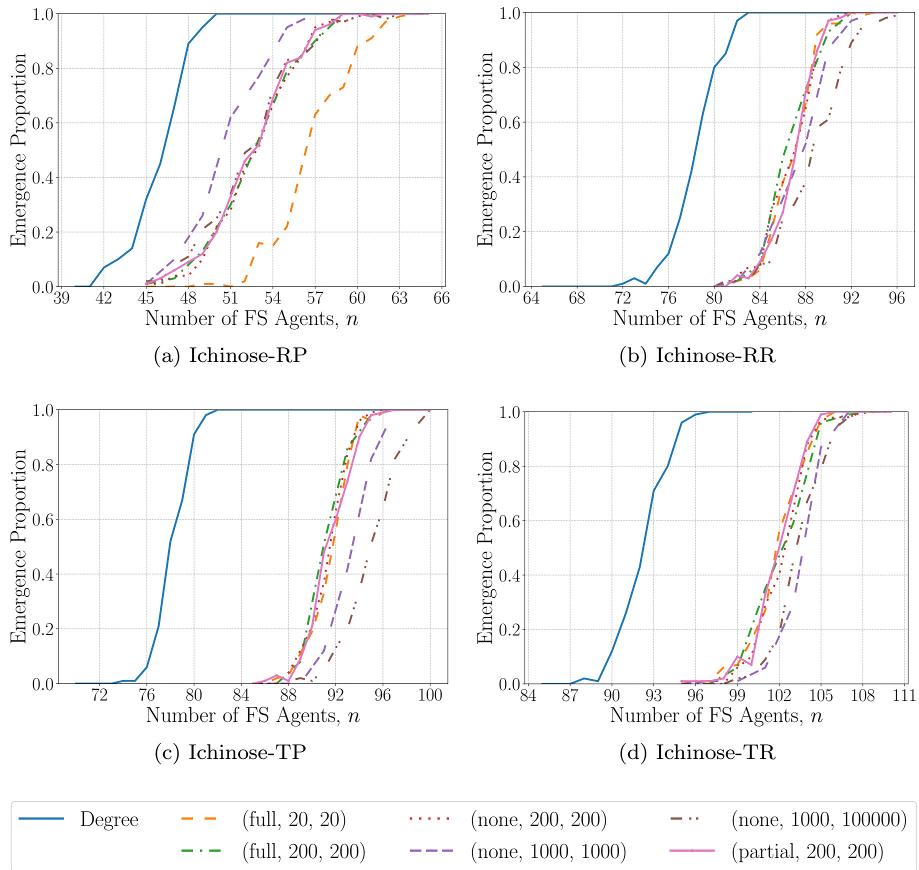

Figure 4.26: Late intervention and destabilisation in the Ichinose models.

and intelligent information gathering are the most beneficial. Each of these approaches still only observes a tiny fraction of the González graph compared to what Static Degree placement must do as shown in Table 4.7. With this in mind, DynaPO performs exceptionally well in this domain with the best performance, $(Partial, 200, 200)$, only requiring 17 more IAs than the Static Degree approach or 1.7% of the network. Whilst none currently approach the pure degree placement performance, this is not the case in either synthetic or real-world networks either which require much larger percentages of available observations.

Figure 4.26 shows the performance of DynaPO for late intervention in each of the Ichinose model variants. As is to be expected, given the underlying dynamic of each of them, the approaches which perform best varies between each model. In many of the models, many of the variants perform similarly but there are a few things to highlight. In Figure 4.26a, which represents the Ichinose-RP model, $(Full, 20, 20)$ performs substantially worse than all other DynaPO variants. This is likely due to the slow update approach taken by this variant. In other models this setting benefits from the fact that it is constantly updating, as high-degree nodes that other variants are reliant on are more likely to be removed and $(Full, 20, 20)$ is more resilient to this due to having constantly updated information. However in Ichinose-RP this fact cannot be relied upon as nodes are selected at random for removal, with no priority for higher degree. We would thus also expect to see a similar effect in Ichinose-RR but this is likely masked due to the fact that all approaches are hindered by the lack of preferential attachment and hence degree gradient. By the same fact, $(None, 1000, 1000)$ does better than the other approaches in this model due to having a full picture of the network at supply time and the high-degree nodes it has selected not being explicitly targeted for removal. In the Ichinose-RR model all variants perform at the same level due to these issues. The lack of a degree gradient due to no preferential attachment hinders all approaches and although $(None, 1000, 100000)$ performs slightly worse at higher levels the distinction is small and explained by the few high-degree nodes it is aware of being removed and not being able to update.

Ichinose-TP highlights similar shortcomings in both "one-and-done" approaches: $(None, 1000, 1000)$ and $(None, 1000, 100000)$. Due to the targeted and consistent removal of high-degree nodes these two approaches will perform poorly as their information rapidly becomes outdated and they must select what they consider low-degree nodes as the high-degree ones are removed. The other approaches have enough up-to-date information, or are able to collect it, that they are not influenced by these issues. Similar issues arise in Ichinose-TR where high-degree nodes are once again targeted and removed in order. However, due

	Ichinose				Gonzalez
	RP	RR	TP	TR	
Average NRE	81.93	79.43	1,000.00	1,000.00	652.10
StdDev NRE	7.98	7.05	0.00	0.00	20.39
o_{deg}	81,933.33	79,433.33	1,000,000.00	1,000,000.00	652,100.00
DynaPO % obs	12.21E-03	12.59E-03	1.00E-03	1.00E-03	1.53E-03

Table 4.7: Node removal events and degree equivalent observations for dynamic topologies when $n = 80$.

to the lack of preferential attachment the other approaches suffer similarly and the difference is thus less pronounced.

Each of the Ichinose models requires different numbers of IAs, even from degree placement, to guarantee destabilisation. This is due to the underlying aspects of the models and makes comparisons between them difficult. However, in the model in which degree and DynaPO have the greatest difference in IA number (Ichinose-TP) this is still only 15 IAs which accounts for 1.5% of the network size. Given the percentage of observation that DynaPO makes of the network, this difference and that found when using DynaPO in the González model are very efficient uses of limited observations.

Overall there is no singular approach that outperforms each of the others within the Ichinose models. Each of $(Full, 200, 200)$, $(None, 200, 200)$ and $(Partial, 200, 200)$ perform similarly to each other in each of the 4 models but none of them outperform the other two in any model. However, they are consistent with none of them ever performing poorly compared to the others. As such, and given their performance in the González model as well, these approaches offer consistent, robust and balanced destabilisation. Given the better performance of $(Partial, 200, 200)$ in the González model, this should be considered the best option to place within an unknown dynamic network.

Thus we have shown that DynaPO can be used to efficiently effect destabilisation of existing conventions. Whilst performing less than 0.1% of the number of observations that degree placement must make, DynaPO is able to achieve

Result	Section
PO-Place, with at most 10% of the network observed, is able to find a set of high-degree nodes 90+% as well as under full observability. Increasing the number of concurrent searches increase this effectiveness.	4.3.2
When using PO-Place to place IAs, when observing only 5% of the population PO-Place produces convention emergence rates 70-90% as effective as full observability.	4.4, 4.4.2
When using DynaPO for initial intervention, when only performing at most 1% of the observations that full observability does, conventions can be emerged with 90+% effectiveness in all topologies.	4.5
Using PO-Place for destabilisation efforts requires observing only 25-35% of the underlying topology and recruiting 1-5% more of the network as IAs in order to guarantee the same level of destabilisation as full observability. In small-world networks it outperforms full observability degree placement.	4.6
DynaPO can be effectively used for destabilisation in dynamic topologies requiring observation of only 0.1% of the possible network over time and 1-2% more of the population as IAs than full observability.	4.7

Table 4.8: A summary of the major results and contributions from this chapter.

comparable performance whilst only requiring 1-2% more of the network population to be made into IAs. With additional observations it is highly likely that performance equivalent to degree could be achieved. Overall, using a partial updating methodology as espoused by DynaPO's design is the best use of the limited observations available and consistently performs as well if not better than other approaches.

4.8 Conclusions

Finding influential positions within a network topology to maximise the effectiveness of fixed strategy Intervention Agents (IAs) is an ongoing area of research in convention emergence. The problem has many facets and variations that make it difficult to find an optimal yet general approach. In many cases, placing the fixed strategy agents at high degree nodes provides effective convention emergence with little computational overhead. Finding high-degree nodes in a network is trivial when the network is fully observable. In many domains, this may not always be possible. Technical limitations such as memory constraints or incomplete information and usage limitations such as finite API calls mean that often a network topology may only be *partially observable*. Finding effective placement for Intervention Agents with these restrictions adds another level of complexity.

In this chapter we presented two placement algorithms, PO-Place and DynaPO, that are designed for use in partially observable static and dynamic topologies respectively. They use finite observations to find sets of high-degree nodes and approximate the set of nodes that would be selected given full observability. Table 4.8 presents a summary of the major results from this chapter and we examine these in more detail below.

With small proportions of the network being observable, PO-Place can locate nodes with similar reach and influence as degree placement. We evaluate the performance in three real-world topologies and show that the addition of concurrent searches and splitting of observations improves the performance of the algorithm across all metrics. With 1-10% observation the algorithm is able to find sets of nodes with >90% of the reach and influence of degree placement. We similarly applied PO-Place to synthetic scale-free and small-world graphs evaluating the performance of the algorithm when applied to topologies with different average degrees and sizes. We showed that PO-Place could achieve similar levels of performance as in real-world topologies but whilst only observing 1-2% of synthetic graphs. This level of performance is higher than other applicable approaches such as that of Brautbar & Kearns [2010] and Borgs et al. [2012b] who can only offer coarse approximations of the highest degree nodes or need to examine 14-24% of the network, respectively. The closest performance comes from Mihara et al. [2015] but PO-Place outperforms their approach by guaranteeing 90+% performance at similar levels of observation.

We then showed that PO-Place performs comparably to degree placement when used to facilitate convention emergence using Intervention Agents whilst only observing 5% of a network topology. We found that the additional aspects of PO-Place benefit the placement mechanism and demonstrated that convention emergence is easily facilitated in partially observable static networks.

We designed and analysed the dynamic placement algorithm DynaPO, evaluating its characteristics, resilience and performance. Formulating what it means for the partial observability problem to exist in dynamic graphs we showed that the variations of DynaPO could achieve high levels of node selection with highly limited observations compared to dynamic degree placement. We then showed that DynaPO could facilitate convention emergence in dynamic networks using less than 1% of the observations of degree placement.

Finally we turned our attention to using both algorithms to cause convention *destabilisation*; removing an existing convention within a system. We showed that both algorithms were potent in doing so and by observing between 20-35% of static networks and 0.1% of dynamic networks were able to achieve similar levels of performance as placement by degree with less than 5% additional agents. In particular, performance in small-world and real-world networks was shown to highly benefit from the approaches taken by PO-Place.

Overall we have shown that it is still possible to direct and encourage convention emergence and destabilisation in a range of partially observable network types and that our algorithms offer quick and robust methods to do so.

CHAPTER 5
Temporary and Budgeted Interventions

Throughout this thesis thus far we have shown the effect that a small number of Intervention Agents (IAs) can have when trying to elicit convention emergence or destabilisation. In Chapter 3 we showed that the primary drivers for their effectiveness were (i) the number of IAs and (ii) where they were placed in the system. In the previous chapter we utilised this knowledge but under the additional constraint that the network might not be fully visible. In this chapter we consider an alternative constraint: *time*. We explore how long IAs must be located in the system to cause permanent change and explore the notion of the *cost* this might incur when trying to destabilise existing conventions. We then use these findings to consider the placement of IAs in budgeted scenarios.

5.1 Introduction

Coordination is fundamental to multi-agent systems (MAS) and self-organisation as it increases the efficiency of systems. Coordination is required as incompatible actions cause conflicts or incur costs. However, it is often impossible to constrain agents beforehand to ensure coordination. This can be due to lacking knowledge of clashing actions or the inability or unwillingness to dictate behaviour. This is of particular importance in systems without centralised control or where the range of possible actions makes pre-determination infeasible.

Hence, as seen previously, many MAS rely on the emergence of conventions, in the form of expected behaviour adopted by agents, with minimal prior involvement by system designers. As such, conventions allow coordinated actions to emerge through self-organisation. In particular, conventions have been shown to emerge given only agent rationality and the ability to learn from previous interactions.

Fixed strategy agents, that always choose the same action regardless of others' choices, have been shown to facilitate rapid convention emergence and to influence the adopted action. A small number of such agents, placed suitably, are able to influence a much larger population [Franks et al., 2014; Griffiths & Anand, 2012; Sen & Airiau, 2007]. However, in realistic domains there is likely to be a cost associated with inserting a fixed strategy agent, or persuading an agent to act in a particular way, and it is desirable to minimise this cost.

Given the self-reinforcing nature of conventions [Boyer & Orléan, 1992; Lewis, 1969], once convention convergence has begun it is likely to continue unless an outside force acts on it. Given this, the permanent inclusion of fixed strategy agents, what we call IAs, is unlikely a requirement to guarantee convention emergence or destabilisation as desired. Giving a "nudge" and allowing the self-reinforcing nature and force of precedence to facilitate the rest of the change would allow the presence of IAs to be temporary within multi-agent systems (MAS) whilst still effecting the same level of change. Additionally, in scenarios where the recruitment of agents as IAs has such an associated cost, finding the minimum amount of intervention needed will reduce these costs.

This chapter considers what the minimum levels of intervention are to effect permanent change within a system. We consider the inclusion of IAs in a temporary manner and study the reduction in their efficacy in order to establish the minimum times that IAs must be present to cause such a change. We show that in initial interventions a small proportion of IAs placed at targeted locations in the population for a short length of time can guarantee convention emergence, showing that the time required can be as small as a couple of hundred timesteps to be as effective as permanent inclusion. We then consider how *when* an initial intervention starts is important and show that very early stages of interactions are formative, requiring direction early to be effective. When considering using similarly placed IAs to destabilise an established convention, replacing it with another of our choosing, we examine how temporary application can be effective when left for sufficient finite time and show that there exists an exponential relationship between the number of IAs and how short this application can be. We also establish how the cost of these interventions varies inversely with the number of IAs used and that this effect is replicated across different pricing mechanisms and all topologies examined. We can then use these findings to consider the best application of limited budgets in effecting change.

The rest of this chapter is organised as follow. In Section 5.2 we introduce relevant parts of the literature that inform the decisions and design of this chapter. In Section 5.3 we begin by exploring the nature of minimum temporary interventions when used to encourage convention emergence at the start of a simulation. We then expand on this Section 5.4 to explore how the start time of an intervention can affect the ability of IAs to enact change. In Section 5.5 we consider similar notions of minimum intervention when it comes to efforts to destabilise existing conventions and examine the notion of the *minimum intervention* required to cause destabilisation. In Section 5.6 we expand on findings in the previous section to consider the notion of placing IAs by specific costs and how this impacts the cost of minimum intervention. Then, in Section 5.7, we use the findings thus far to denote the concept of a *budgeted intervention* and introduce a placement heuristic, BudgetedPlacement that can maximise the performance from a given budget to spend on IAs. Finally, in Section 5.8, we present our conclusions and final thoughts.

5.2 Background

Little work exists in the literature concerned with the usage of fixed strategy agents to effect convention emergence and none does so when considering the agents as temporal and finite.

Both Griffiths & Anand [2012] and Sen & Airiau [2007] explore the use of fixed strategy agents when placed at the beginning of a simulation to direct and encourage convention emergence. They show that the presence of these agents has a marked effect on the time taken to reach convention emergence with increasing numbers of fixed strategy agents reducing the amount of time taken. This indicates that there is a temporal nature to convention emergence that can be affected by the inclusion of fixed strategy agents and we expand on this by considering how the inclusion of such agents affects the time of convention emergence even when their presence is removed.

Both Boyer & Orléan [1992] and Lewis [1969] discuss the self-reinforcing nature of conventions, that the force of precedence and the ease of utilisation compared to choosing clashing actions means that conventions, through positive reinforcement are self-sustaining. Tied into this is the belief that once one of the equilibria represented by the possible conventions has been chosen the force of precedence within the system will further add pressure to select this equilibria to all agents. We thus believe that the initial direction given to the convention emergence is the most important part with the force of precedence allowing the system to continue to emerge the convention without additional external force.

Franks [2013] briefly considers the notion of interfering with conventions later in their life cycles after they have emerged or whilst they are in the process of doing so. The model used is quite distinct from our own however and allows modification of the individual agents payoff matrices to allow additional rewards to be granted to them. In our model we assume that agent architecture cannot be modified in this manner. Additionally, they consider interventions at quite distant timesteps after the start of the simulation whilst we believe the force of precedence will be too high by that point and would transform the intervention into destabilisation.

Previous work often assumes no restrictions when placing fixed strategy agents into the network. We follow this assumption, but add that such an insertion has an associated *cost*. In real-world domains, inserting fixed strategy agents likely has such a cost, and understanding how to minimise this is crucial. In this chapter, we investigate the effect of the cost of insertion and its relation to the duration and efficacy of intervention. Delre et al. [2010], in the context of marketing budgets, raises the question of how best to choose which consumers to target for viral marketing and it is in a similar vein that we consider the costs of recruiting agents as IAs.

5.3 Temporary Initial Interventions

We begin with a consideration of temporary initial interventions, using IAs at the beginning of a simulation to elicit and direct convention emergence. We previously have shown this is to be effective but studied the case where the IAs were placed at time $t = 0$ and left within the system for the duration. We now wish to find the minimum amount of time that the IAs must be left in the population in order to cause a change that is permanent enough to direct convention emergence even without their presence. We believe that the self-reinforcing properties of conventions and the precedence that the IAs instil to a particular action choice will enable them to be removed quite early and still have the desired effect. Being able to do so, and understanding what time-frames such agents must be left in the system to guarantee it, would allow much easier initial interventions by not requiring the permanent change of agents into IAs.

Having already shown the minimum number of IAs needed with each placement metric to guarantee convention emergence in Section 3.4.1 we investigate the minimum amount of time they must be present by placing that number of IAs into the system at time $t = 0$. We then remove them some time, $t_{removal}$, later and allow the simulation to continue. We increase $t_{removal}$ in steps starting at 10 (as a minimum) until the presence of the IAs is once again causing 100% of simulations to converge to the desired convention. These steps were 5 for the scale-free graphs and 10 for the small-world graphs as we found that 10 gave insufficient granularity of results for the scale-free graphs. We otherwise use the same interaction model as in the previous chapters: 5000 node scale-free and small-world graphs, constructed with $m_0 = m = 3$ for the scale-free graphs and

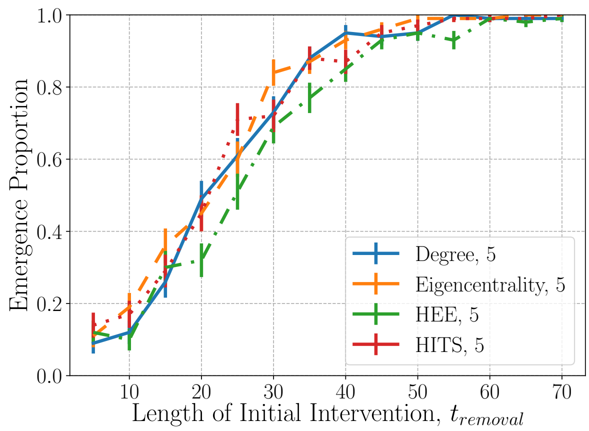

Figure 5.1: Minimum length of initial intervention in the 5000 node scale-free network.

$ce = 1$, $l = 1$ for the small-world graphs. Each timestep each agent chooses one of its neighbours to play the 10-action coordination game $(+4, -1)$ with and both update their internal knowledge using Q-Learning. We measure the proportion, over 100 runs, that each setting emerges the IAs strategy as a global convention, defining a convention to have emerged when 90% of agents, when not exploring, would choose it.

Figure 5.1 shows the results of this for the scale-free network for each placement metric and the number of IAs indicated in the legend. For each metric, the relationship between the length of initial intervention and the effectiveness of the intervention is fairly linear with longer interventions producing corresponding increases in performance. What is most interesting however is that every metric reaches peak performance, that achieved when the IAs were placed within the system permanently, with a $t_{removal}$ of only ~60. This indicates that it is the very early period in the emergence of conventions where most of the results are decided with a small number of IAs able to enact permanent change whilst only being present for a short time period. Beyond this initial window of influence, the IAs can be removed from the system without the direction of convention emergence changing; those agents they have already converted are enough for

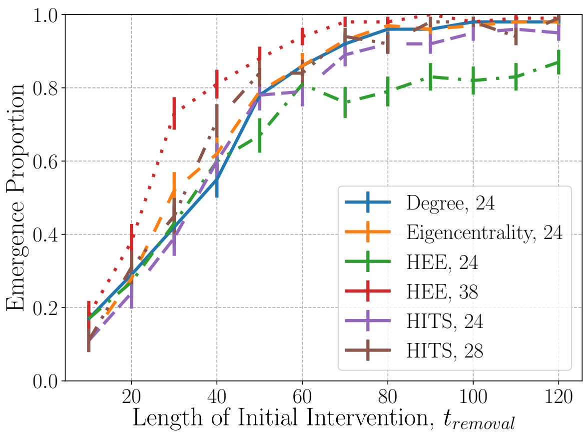

Figure 5.2: Minimum length of initial intervention in the 5000 node small-world network.

the convention to be self-sustaining. Interestingly, the worse performance of highest edge embeddedness (HEE) in eliciting convention emergence is present here as well where, despite having enough IAs that it can guarantee convention emergence normally, it performs statistically significantly worse than degree at a number of points (two-tailed proportion test, $p < 0.05$). This indicates that not only is HEE placement worse at causing conventions to emerge, they are also less stable than those created when placing by other metrics, being affected more by the removal of the IAs.

Figure 5.2 shows similar results for the 5000 node small-world network. In this instance, given the disparity between the number of IAs needed to guarantee convention emergence by degree/eigencentrality and HEE or hyperlink-induced topic search (HITS) we investigate two values for each of the latter: the same number of IAs as needed by degree/eigencentrality and the actual number needed for those metrics to guarantee convention emergence as shown in Section 3.4.1. The behaviour here is slightly different than in the scale-free topology with increases in $t_{removal}$ causing performance to approach its peak in a noticeably asymptotic manner; whilst initial increases in length have a marked effect on performance, the extra benefit gained from increasing $t_{removal}$ between

60 and 120 have less and less impact on performance. This indicates again that it is the initial period that is most important and that, in small-world networks, increases beyond this are limited in how much they affect performance. Again, we find that, within a relatively short time frame, the presence of IAs is no longer needed to ensure the desired convention emerges, though this is notably longer than in scale-free networks and matches with the slower emergence of conventions in general for small-world networks. The locally clustered nature of the links in small-world topologies means they are more resilient to change [Franks et al., 2013] and so we must include IAs longer to ensure that change is permanent.

Of particular note is the performance of HEE with 38 IAs. Despite this being the same criteria as for the others, namely the minimum number of IAs that was required to guarantee convention emergence when included permanently, its performance in this scenario is markedly better than degree/eigencentrality with the same constraints. We can conclude from this that larger numbers of IAs benefit temporary initial interventions, making shorter interventions more likely to succeed with less concern for the placement metric being used. Whilst the performance of HEE when IAs are included permanently is worse than the other placement metrics, it is not substantially so and thus the additional IAs still provide benefits over fewer IAs slightly better placed.

We similarly must consider the nature of minimum initial intervention in the dynamic networks previously explored. These have been shown to have major differences in behaviour due to the changing nature of both nodes and edges and thus minimal initial convention emergence within them is likely to differ as well. We consider 1000 node graphs of the González network and all 4 modes of the Ichinose model with settings otherwise the same as these graphs in previous chapters.

The results for minimal initial intervention in the González network are shown in Figure 5.3 and show broadly the same effect as in the small-world and scale-free networks, indicating that these findings are general rather than

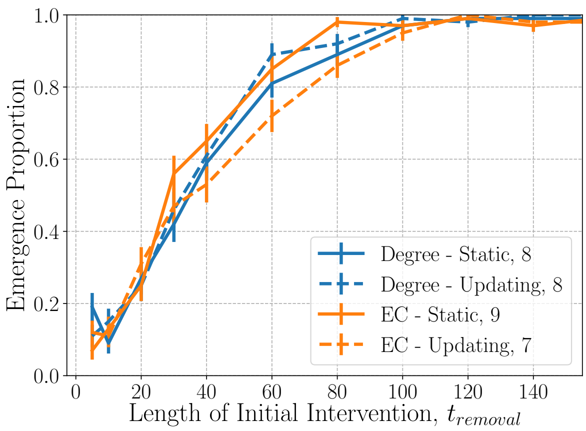

Figure 5.3: Minimum length of initial intervention in the 1000 node González network.

specific to static networks. Given the more distinctive differences in performance for the placement metrics in dynamic networks we only include degree and eigencentrality (both static and updating) although we observe similar effects for all metrics. As in the static networks we find that initial increases in $t_{removal}$ have commensurate effects on increasing the performance of the IAs and that this tapers off at longer interventions as the performance approaches that of permanent inclusion. Most markedly there is no significant difference between the static and updating approaches indicating that up-to-date information is not a primary concern in maximising the effect of short interventions. The length of intervention required in the González network is comparable to those found previously, requiring less than 150 timesteps to reach the same levels as permanent inclusion. Given that this network is 1/5 the size of the static networks this indicates that the dynamic networks require the IAs to be present longer comparatively in order to generate a permanent effect and is likely due to the changing topological nature as converting the local area is no longer as much a guarantee that it won't easily revert as new connections are made.

We find that the behaviours of the 4 Ichinose models are nearly identical, leading us to conclude that the temporary initial interventions are primarily

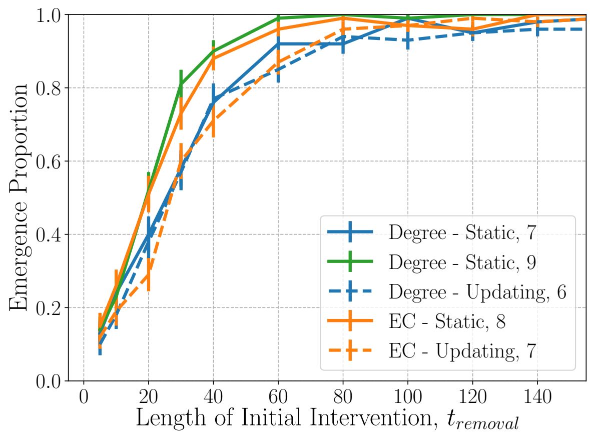

Figure 5.4: Minimum length of initial intervention in the 1000 node Ichinose TP network.

unaffected by the different node removal and edge attachment rules and that the length of the intervention is the predominant factor. Figure 5.4 shows the results for the Ichinose TP model as a representative case. As can be seen, the length of time that the IAs must be included to maximise the performance is comparable to those required in the González network with the asymptotic nature highlighting again the importance of the initial length increases in their effect on the performance of IAs. Of particular note is the performance of static eigencentrality which performs markedly better than the other placement metrics despite still having the minimum number of IAs. We attribute this, as with the performance of HEE earlier, to the effect that additional IAs can have in helping a convention become established in the small timeframe made available to them. To test this hypothesis we additionally include a larger number of IAs for another placement metric, beyond the minimum number required, with 9 IAs and the static degree placement heuristic. This results in marked improvements in performance, particularly at shorter lengths, than is the case when static degree only has 7 IAs available and supports our notion that small increases in the numbers of IAs, despite not changing the proportion when included permanently, can facilitate faster initial intervention and would need to be included for

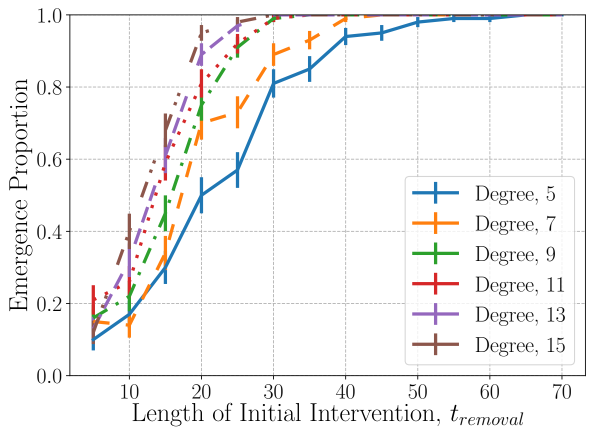

Figure 5.5: Minimum length of initial intervention in the 5000 node scale-free network with varying numbers of IAs

less time to enact the same level of permanent change. Whilst previous work by Griffiths & Anand [2012] and Sen & Airiau [2007] has shown that the inclusion of more IAs reduces the overall time for convention emergence this is the first indication that the effects occur so early in the convention emergence and can benefit from such a small increase in IAs.

To better quantify the effect that the increased number of IAs has on the minimum time we consider a representative case in the scale-free network, as shown in Figure 5.5. As can be seen, just small increases in the number of IAs beyond the minimum needed increase the effectiveness at shorter lengths of intervention to a statistically significant level compared to 5 IAs (two-tailed proportion test, $p < 0.05$). This supports our notion that even a few more IAs can make a meaningful impact due to their ability to influence the emerging convention more readily. However, this increase in impact is one of diminishing returns with the difference between higher and higher numbers of IAs becoming less meaningful. Indeed, at the highest levels shown, they are not statistically significantly different from one another at any point although the plot still shows marginal increases in efficacy. Additionally, increasing the number of IAs decreases the minimum $t_{removal}$ needed for full effectiveness from around ~65 with

5 IAs down to ~30 with 15 IAs. This decrease, and the non-linear nature of it, is something we revisit the importance of later in Section 5.5.

Overall, we find that there is indeed a minimum length of temporary initial intervention rather than the requirement to include IAs until the convention has emerged fully. We find that IAs need only be included in the system for a short period of time (<150 timesteps in all the simulations above) before their effectiveness is practically indistinguishable from the case where they are left within the simulation. This has important implications for the nature of convention emergence as it allows those who wish to facilitate convention emergence in MAS to not have to find a way to permanently change agents into IAs but instead shows that such a change is only necessary for a short time.

5.4 Staggered Temporary Interventions

We have shown that there is a minimum length of time that IAs must remain in a system to facilitate the level of desired convention emergence as caused by the inclusion of the same number of IAs permanently. We have seen that increasing the number of IAs reduces this minimum length of time and increases the effectiveness of the IAs even when included for only short periods at the beginning of a simulation. Additionally, in the previous chapters, we have seen that the number of IAs needed for destabilisation is substantially higher than the number needed for initial intervention. This indicates that the influence power of the IAs is tied to *when* they are introduced to the simulation and that the number of IAs sufficient to direct an undecided population towards one option lacks the influence to go against the precedence of an already established convention. Finding the speed with which this transition occurs, when the number of IAs that can enact initial intervention become unable to change the population consistently, will highlight the point at which intervention switches from initial to late, when the IAs will be trying to counteract an already established (or rapidly becoming established) convention rather than simply directing population choice. This information is important for convention emergence in open MAS as it is unlikely that all users will be present at the very start of interaction in such systems but may well join slightly later, after interactions have been ongoing.

In this setting we seek to measure this effect by varying the time at which IAs are initially placed within the system compared to the case where they are inserted at $t = 0$. We call these *staggered interventions* and will denote the start time of the intervention as t_{start}. As we are concerned with the rate at which initial intervention ceases to be effective as a means of ensuring the desired convention emerges, we utilise the same settings as used in the previous section. Namely, we use the minimum number of IAs that were required to guarantee convention emergence for a given topology and placement metric that were found in Section 3.4.1. Whilst this will result in an unequal number of IAs for each placement metric, even within the same topology, it unifies them in the fact that this was what consistently allowed conventions to emerge when used at $t = 0$. As we are concerned in this chapter with temporary interventions we similarly limit the time that the IAs are placed within the system to that shown in the previous section to allow maximum performance: 100 timesteps in the scale-free network, 150 in the small-world, Ichinose RR and Ichinose RP networks and 200 in the González, Ichinose TR and Ichinose TP networks. We perform 100 simulations with each setting and find the proportion of results that still emerge the desired convention under these constraints.

We begin by examining the effect of staggered intervention in the 5000 node scale-free network as shown in Figure 5.6. As well as the minimum number of IAs we also include a larger than minimum number for each placement metric to study the effect that this increase has under staggered interventions. The placement metric and the number of IAs is shown in the legend. As can be seen, within the same number of IAs there is little difference in the performance of each individual placement metric, with most not statistically significantly different consistently. An exception to this is HITS with 5 IAs which is worse than

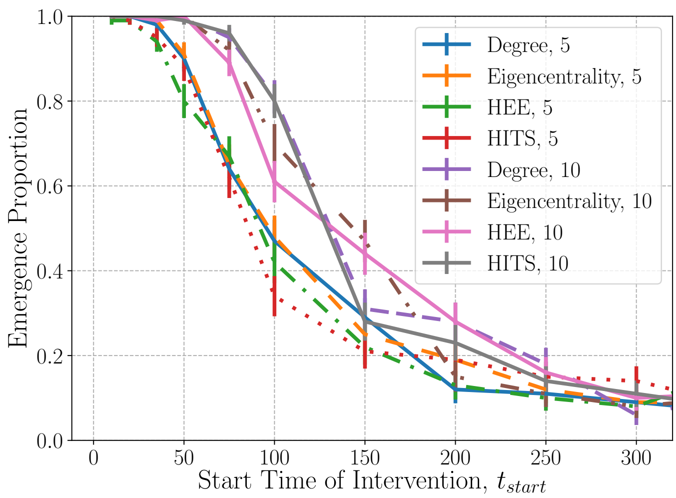

Figure 5.6: Effect of intervention start time for staggered interventions in the 5000 node scale-free network.

the runs for the other metrics with 5 IAs over an extended range of start times. The primary observation from these results is that the start time, t_{start}, of a staggered intervention has an almost immediate and drastic effect on the efficacy of the IAs. Starting the intervention as little as 100 timesteps after the beginning of the simulation reduces the effectiveness of the intervention by almost 55-60% across all metrics with 5 IAs. Whilst an increased number of IAs performs marginally better it too is reduced by the same level by $t_{start} = 150$. This highlights again the importance of early intervention with the system already resilient to outside attempts to influence it very early on. The speed with which this change occurs is rapid, reducing all intervention attempts to no better than random chance by $t_{start} = 300$. The relationship between efficacy and t_{start} is again asymptotic in nature with the later start times reducing the performance by reduced amounts each time. This again highlights that it is the very start of the simulations and interactions that determines the direction of likely convention emergence and that the force of precedence and positive-reinforcement of an emerging convention rapidly become too potent for the IAs to overcome, requiring destabilisation efforts rather than just initial intervention.

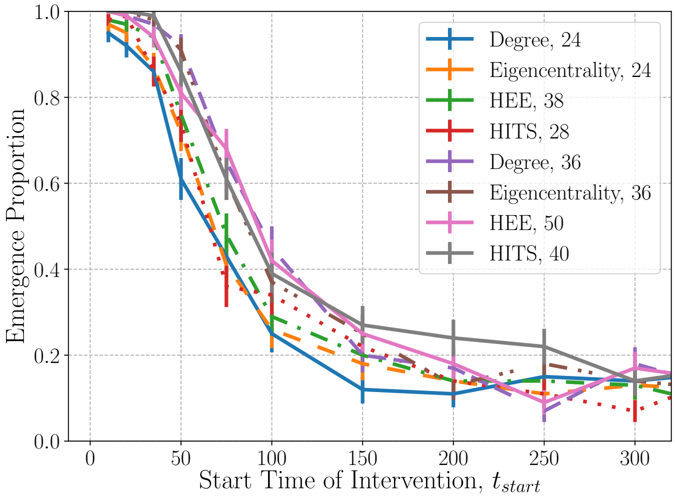

Figure 5.7: Effect of intervention start time for staggered interventions in the 5000 node small-world network.

Figure 5.7 shows similar results in the 5000 node small-world network. The detrimental effects of starting intervention attempts at later times are even more pronounced in this topology, highlighting again the resilient and self-reinforcing nature of the local clusters present in the small-world topologies. We again include results when using more than the minimum number of IAs (shown in the legend) but the inclusion of these additional agents has little effect on the rate at which efficacy decreases. The initial decrease from beginning the intervention attempts just slightly after the beginning of the simulation are more marked here, with the minimum IA placement interventions dropping to 40% efficacy with $t_{start} \approx 75$ and reduction to behaviour no better than chance again occurring by $t_{start} = 300$. We again find little difference between the performance of the placement metrics with none consistently performing better or worse despite the differing numbers of minimum IAs assigned to them. We can conclude that, whilst beneficial in temporary initial interventions, the presence, at these levels, of these extra IAs is insufficient to cause change in the system once another convention has begun emerging.

We similarly explore the effect in dynamic topologies, focusing on the per-

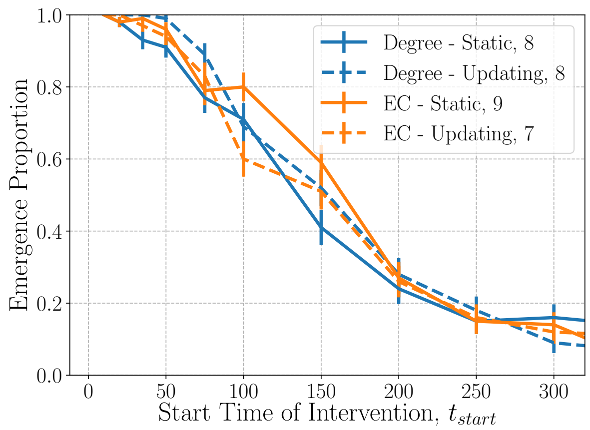

Figure 5.8: Effect of intervention start time for staggered interventions in the 1000 node González network.

formance of the best performing metrics from previous simulations: updating and static eigencentrality and degree. Figure 5.8 shows the results in the 1000 node González network. As can be seen, staggered interventions in the González network are slightly more resilient to the detrimental effects of increasing t_{start} with all metrics taking until a start time ~150-175 before efficacy is reduced to 40%, better than both the scale-free and small-world networks. The decrease is also more linear, without the sharp decline present for small shifts in intervention start time found in the static networks. This indicates that the temporary initial interventions in the González network are less affected by the initial gap, able to more readily overcome the emergent conventions that are becoming established and direct convergence to the desired strategy. The topological churn, the amount that the edges and nodes are changing, is greater in the González network than in the Ichinose model (which only removes/reattaches a single node and its edges each timestep) and we believe this to be to the benefit of staggered interventions as the emergent conventions are constantly exposed to each other, undermining the growth of each and allowing the influence of IAs to be relevant for longer as shown by its more linear decline. However, this slight benefit does not last forever with the performance of all metrics dropping to no

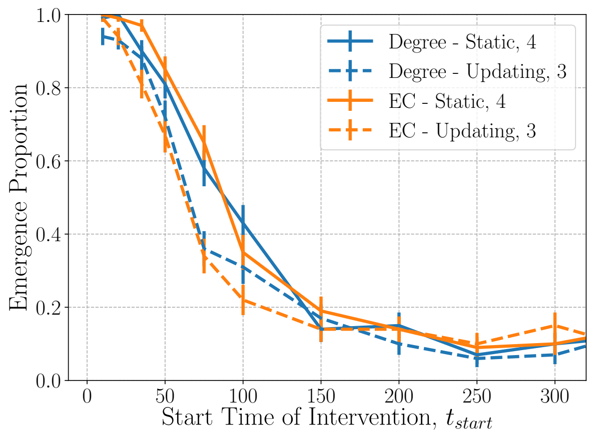

Figure 5.9: Effect of intervention start time for staggered interventions in the 1000 node Ichinose RR network.

better than chance by $t_{start} = 300$, similar to the other topologies.

We again find that the effects in the 4 Ichinose models are almost indistinguishable with only slight variations in performance at any given t_{start}. As such we include Figure 5.9 as a representative example of the broad features that appear during staggered interventions in the Ichinose networks. In each of the models, the effect of increasing t_{start} is closer to that of the synthetic networks than it is to that of González with the decrease to 40% efficacy happening between $t_{start} = 75$ and $t_{start} = 100$ in each model, a steep decrease that highlights the Ichinose models' sensitivity to the initial period of convention emergence. Also of note is the effect that the slightly higher number of IAs has on the performance of the static metrics, increasing their efficacy slightly unlike in the small-world network seen earlier. This is likely because these increases represent a much larger relative increase and, as seen in Figure 5.5, this can have a marked effect and this notion is reinforced by the fact that no such distinction in performance is present in the other Ichinose models where the number of IAs assigned to static and updating is the same. However, even with this boost in performance, the static metrics still rapidly decline in efficacy, highlighting again the importance of initial interventions being utilised as early as possible.

5.5 Temporary Interventions for Destabilisation

We have thus far shown that temporary interventions can be utilised in initial interventions to great effect, requiring the IAs to only be present for a short period of time to enact a self-reinforcing change within the system. We have shown that this allows target conventions to emerge with the same frequency as afforded by permanent inclusion of the IAs, despite the IAs only being present for a fraction of the time needed to reach convention emergence. We have also shown that the effectiveness of these temporary initial interventions rapidly diminishes when introduced later in the convention emergence life cycle, with the force of precedence amongst already emerging conventions too great to overcome. In this section, we explore the logical continuation of these two threads: using temporary interventions to cause the destabilisation of an already established convention.

5.5.1 Length of Intervention

We begin by considering the effect that temporary inclusion of IAs has on destabilisation by examining the convention membership sizes over time of both the IA strategy and the dominant one.

As a representative test-case we consider the averaged behaviour over 100 runs of temporary late intervention in the scale-free network used before. IAs are introduced at $t = 1500$, being placed at high-degree locations, and left in the system for a finite time, shown in the captions of Figure 5.10. As can be seen in Figure 5.10a, with 200 IAs the membership size of the dominant convention rapidly decreases after IAs are introduced, approaching parity with the rising membership of the IA convention. However, upon the removal of the IAs, the reduction in the dominant convention is mostly undone as it reclaims nodes that had shifted convention. The removal of the IAs happens too prematurely and the dominant convention, on average, remains established. This shows that there is indeed a minimum length that IAs must be present to effect

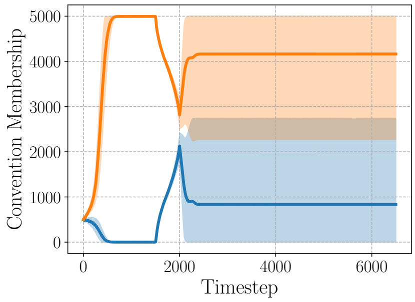

(a) 200 IAs, 500 Timesteps

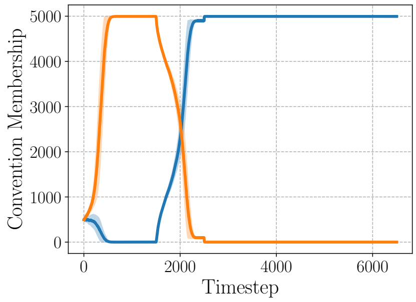

(b) 200 IAs, 1000 Timesteps

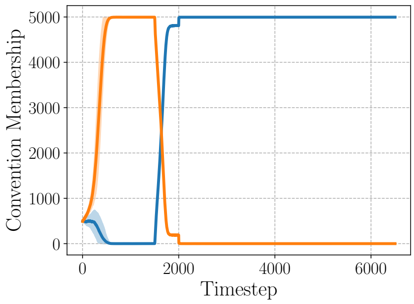

(c) 400 IAs, 500 Timesteps

Figure 5.10: The effect on scale-free graphs of different numbers of IAs when introduced for finite time. The IA strategy is shown in blue, the dominant strategy in orange. The shaded areas represent the standard deviation at each timestep over the runs.

the desired change, and their removal before this will allow the dominant strategy to rebound. Increasing the length of the intervention further, as is done in Figure 5.10b, shows this assertion to be true with the same number of IAs included slightly longer guaranteeing the destabilisation. However the length of the intervention is not the only thing we can vary and Figure 5.10c instead shows the effect of increasing the number of IAs whilst leaving the length of the intervention the same. In this scenario destabilisation is still guaranteed, happening rapidly and displacing the dominant convention. We can thus conclude that, as we have seen in the previous sections, there is both a minimum number of IAs and a minimum length that this number of IAs must be present to cause destabilisation. Increasing the number of IAs reduces the minimum length of time that they must be present in much the same way as it did for temporary initial interventions. Unlike in temporary initial interventions, dramatically increasing the number of IAs does not look to have an asymptotic effect.

We have previously seen that there is minimum number of IAs that must be present in order for destabilisation to occur. These initial findings indicate that there is also a minimum length of time that they must be present in order for them to induce destabilisation. That is, there is what we call a *minimum intervention* in each topology, a minimum number of IAs (which is topology specific) and their associated minimum length of intervention.

As indicated in Figure 5.10c, additional IAs decreases the minimum time for effectiveness. We wish to quantify this effect in order to establish how changing the number of IAs changes the minimum time required. To do so we vary the number of IAs from the minimum found in Chapter 3 in steps up to 2500 IAs, representing 50% of the 5000 node scale-free network. Whilst this upper limit is likely to be an infeasible target in many domains due to the open nature of many MAS, we include it for completion and to study the phenomena over as wide a range as possible. Having set the number of IAs we then look to find the minimum length of time these IAs must be included in order to cause destabilisation. We do this by varying the length of inclusion upwards from 0 in steps of 5 to 50, seeking to find the lowest value where destabilisation occurs in 100% of 30 runs. For those numbers of IAs where destabilisation does not occur by this point (which is likely for the lower numbers of IAs) we further increase the length in steps of 10 to 200. Similarly for any numbers of IAs which do not consistently cause destabilisation at this length we then increase the intervention time in steps of 50 to as high as is necessary for destabilisation to occur. This approach allows varying levels of granularity where it matters.

Figure 5.11 shows the results of this for the 5000 node scale-free network with degree placement of IAs. As can be seen, increasing the number of IAs even slightly has a dramatic effect on the minimum intervention length found, reducing it substantially. More importantly, this relationship is non-linear, with

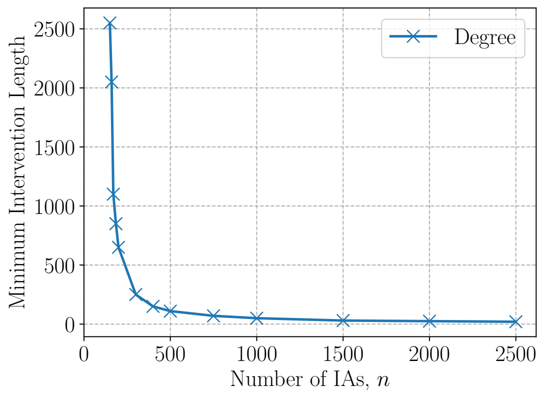

Figure 5.11: Number of IAs vs. the minimum length of intervention to cause destabilisation for scale-free topologies.

increases in the number of IAs causing much larger reductions in minimum length as we increase the number from the minimum required. Additional increases beyond ~300 agents cause much lesser reductions in minimum intervention time as the system asymptotically approaches a minimum intervention of zero; the relationship is one of rapidly diminishing returns with the most change happening at lower numbers of IAs.

To allow us to better quantify these notions of minimum interventions and better compare the effects at either end of the spectrum, we introduce the concept of the *cost* of minimum intervention. We assume that each IA has a cost associated with it that must be paid each timestep it is an IA in the system. This allows us to combine both the number of IAs and the minimum length they are required to allow comparisons between different minimum interventions in the system with different numbers of IAs and different lengths they must be included. We begin by considering a uniform cost, where each IA costs 1 unit per timestep. In this instance, the cost of the intervention is simply the number of IAs multiplied by the the length of time they are required.

In doing this we also expand the consideration of minimum intervention to the 5000 node small-world graph as well as considering eigencentrality placement in both topologies as well as degree. We find the minimum lengths of

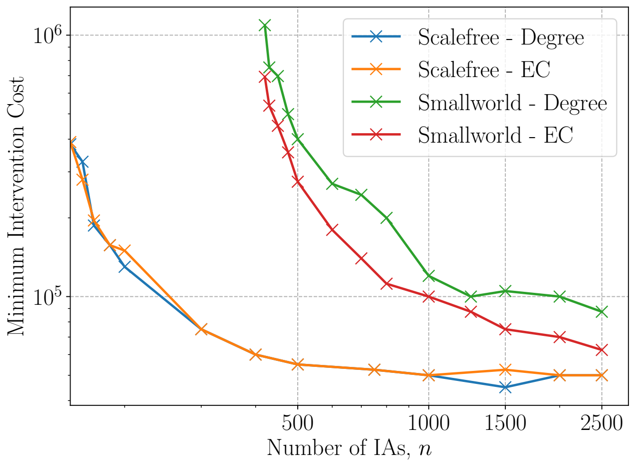

Figure 5.12: Number of IAs vs. the minimum intervention cost needed to cause destabilisation for static topologies.

intervention as before and use these and the number of IAs to calculate the costs of minimum interventions using each placement metric. Figure 5.12 shows the results of this. As can be seen, eigencentrality and degree are almost indistinguishable within the scale-free network, both requiring nearly the same level of minimum intervention regardless of the number of IAs. In both cases however, additional IAs act to continuously reduce the cost of intervention although the relationship is indeed one of diminishing returns as noted earlier with the costs of minimum interventions with the highest number of IAs being nearly identical. There is a marked difference in the two metrics when considering destabilisation in small-world topologies with eigencentrality allowing markedly lower costs of intervention than degree does for the same number of IAs. This indicates that eigencentrality placement is better at facilitating destabilisation than degree in small-world topologies allowing shorter temporary interventions to cause the same level of destabilisation as longer ones with degree placement. This difference is consistent throughout all levels of IAs, indicating this is a generalised trend. We also see that destabilisation in scale-free topologies is cheaper than in small-world ones at all points, highlighting again the robustness of small-world topologies to external influencers meaning that IAs must be

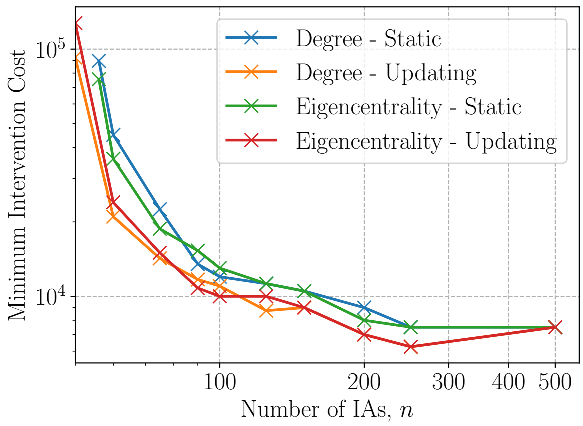

Figure 5.13: Number of IAs vs. the minimum intervention cost needed to cause destabilisation in the González network.

present longer to enact the same level of change in the small-world networks.

We apply the same methodology to the dynamic networks, varying the number of IAs up to 50% of the network again and utilising both static and updating degree and eigencentrality as the best performing metrics in each dynamic topology as found in Chapter 3. We present these findings for the González network in Figure 5.13. As the figure indicates, we find that the updating heuristics consistently produce lower cost minimum interventions in comparison to their static equivalents, across the entire range of IAs. This is to be expected as they consistently outperformed the static metrics in destabilisation efforts with permanent IAs inclusion but reveals that they are more efficient at nearly all levels of intervention and create a reduction in the length of time the IAs must be present in order to facilitate destabilisation. We find that there is little difference between degree and eigencentrality themselves however with the static and updating versions of both performing nearly identically. We observe the same pattern across each of the Ichinose models as well and so do not include them here.

Our primary finding here, and one that appears consistent across all placement metrics and network types, both static and dynamic, is that increasing the number of IAs used for destabilisation efforts nearly always produces a more than commensurate decrease in the minimum intervention time required meaning that the overall cost of intervention is lower. We can conclude that attempting to maximise the number of IAs that are being used for destabilisation is the best course of action across all topologies although this relationship is one of diminishing returns with small increases beyond the minimum number needed to facilitate destabilisation producing the most drastic decreases in cost.

5.5.2 Cost of Intervention

We have shown that we can minimise the cost of intervention by maximising the number of IAs. However, in the previous section we assumed that the cost of each agent to become an IAs was uniform, namely that it was a unit cost of 1 per timestep for each IA. We find this assumption to be unrealistic in many ways as agents with higher levels of influence are more likely to require higher levels of remuneration, either due to their ability to influence a larger number of people or due to their intrinsic importance in the domains being modelled. Consider for instance the scenario of trying to hire a brand ambassador on social media. Those with lower numbers of links will be worth less than those with much higher numbers and will require different levels of remuneration because of this. This is particularly true in any system where multiple individuals might be bidding for the same agent and hence price will be driven up for the most influential of them.

To model this we extend the previous notion of cost of minimum intervention to the case where cost is proportional (in this case equal) to node degree as a measure of node influence. We believe this captures the problems identified above and will enable us to establish the generality of the findings from the previous section.

Figure 5.14 shows the effect on minimum cost of intervention when pricing agents this way for the static networks. Whilst mostly the same as with uniform pricing, there are a number of important differences to highlight. Firstly,

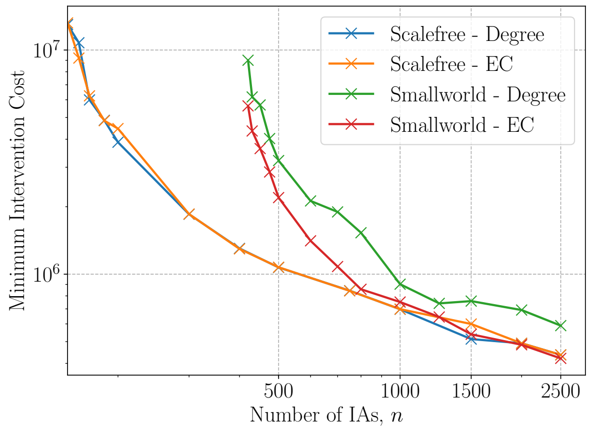

Figure 5.14: Number of IAs vs. the minimum intervention cost needed to cause destabilisation for static topologies. IAs are placed by the metrics indicated with vertex costs equal to degree.

that, under this pricing model, the costs of intervention in scale-free and small-world are much closer together, particularly at the larger numbers of IAs where eigencentrality placement in the small-world topologies is practically indistinguishable in terms of cost to the placement mechanisms in scale-free topologies. This disparity highlights the difference between the scale-free and small-world topologies, that the same number of IAs selected by degree from both will have higher total degree in the scale-free topology due to the skewed nature of the power-law degree distribution. When considering high-degree nodes to have a commensurate increase in cost, this means that influence is more expensive in the scale-free networks for a given number of IAs.

Additionally, when pricing by degree, we now see a continued drop in cost of intervention even at the highest numbers of IAs rather than the diminishing returns seen with uniform pricing. The decrease in time afforded by more IAs becomes more of an issue when high-degree nodes are more expensive per timestep. This reinforces our previous finding, that it is better to maximise the number of IAs within the system as this reduces the overall cost of destabilisation.

When considering the degree-based pricing mechanisms for dynamic net-

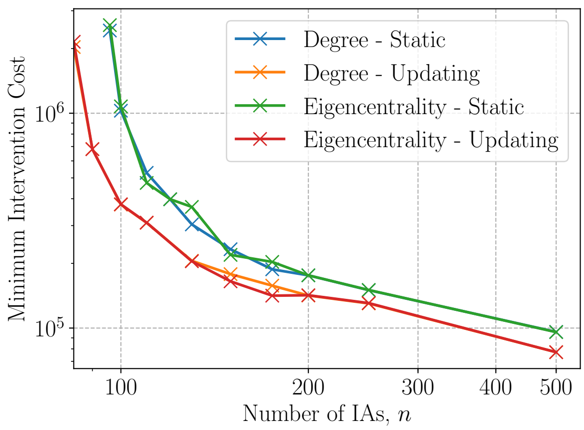

Figure 5.15: Number of IAs vs. the minimum intervention cost needed to cause destabilisation in the Ichinose TR network. IAs are placed by the metrics indicated with vertex costs equal to degree.

works we must consider what the cost of an IA means for static and updating heuristics as the cost of the agent is likely to change over time. We utilise a mechanism where the cost of an IA is what the system *believes* the degree to be at any given moment. Thus in static metric placement the cost of an IA is set at its selection and remains that way until it is no longer an IA. We believe this best models the situation from the placement perspective as the static metrics work under the assumption that the value of the metric is unchanging and hence they do not reconsider the IAs. We feel this models real-world situations where the price is "locked-in" and subject to value fluctuations, both good and bad, that do not affect the amount the agent is paid. We again find that the costs in the dynamic networks are mostly identical in both range and features and so we include Figure 5.15 as a representative case showing the effect of degree-based cost in the Ichinose TR network. The changes in the dynamic networks are less than those in the static networks. The same difference in performance between the static and updating placement metrics still present indicating that the cost benefit of these is important even across multiple pricing mechanisms. However we see the same effect as in the static networks where, when pricing by degree, increasing the number of IAs always causes a decrease in cost, lending support to the notion that this is a generalised rule applicable across many domains.

5.6 Cost-Based Placement

We have thus far shown that minimum interventions exist for given topologies and number of IAs. We have also shown that increasing the number of IAs reduces the cost of minimum intervention by making the presence of the IAs be required for less time. However, our previous experiments have assumed that it is possible to calculate the length of a minimum intervention *a priori* by running multiple simulations with increasing lengths and finding the minimum that still exhibits destabilisation. In real-world scenarios, this is impractical or in many cases impossible as the situations where we might want to utilise destabilisation are not amenable to multiple attempts. Instead we must consider an alternative, on-line notion of minimum intervention that still facilitates destabilisation of the dominant convention and allows us to know when IAs can be removed from the system with little chance of the system rebounding.

When considering the minimum cost of intervening we examine the idea of the minimum length of time that a given number of agents must remain in the system in order for irreversible destabilisation to occur. To quantify this we introduce a new measure: the crossover ratio χ_{co}. The crossover ratio is defined as:

$$\chi_{co} = \frac{\textit{size of IA strategy convention membership}}{\textit{size of dominant strategy convention membership}} \tag{5.1}$$

We can thus describe minimum interventions as the minimum amount of time that a given number of IAs must be introduced to cause χ_{co} to exceed some threshold, γ_{co}. In this section we set $\gamma_{co} = 1.5$ such that the convention of the IA strategy must become 50% larger than the previously dominant strategy to be classed as destabilisation.

This notion of destabilisation is different from that which we have used throughout the rest of this thesis and instead allows real-time evaluation of when

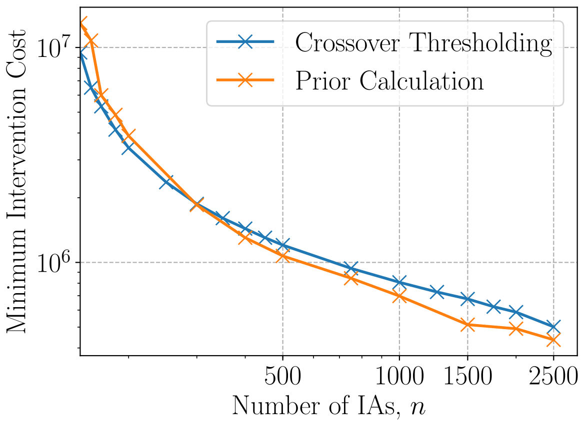

Figure 5.16: Comparison of the minimum cost of intervention produced using multiple methods of calculation in the 5000 node scale-free network. IAs are placed by the degree with vertex costs equal to degree.

a convention is likely to be destabilised and replaced even if IAs are removed. This allows us to monitor the minimum length of intervention knowing that the force of precedence and the self-reinforcing nature of conventions is likely to take the system the rest of the way and allow it to reach the 90% Kittock criteria. By defining the threshold relative to the dominant convention, rather than in the absolute sense of the number of agents in the population, we can be assured that the destabilisation has occurred and that the IA strategy is already emerging beyond it.

To quantify that this notion of minimum intervention is similar to the one used previously we must establish whether they produce similar costs of minimum intervention. We consider the case for the 5000 node scale-free network and produce minimum intervention costs for a range of IAs using both the old *a priori* method and the new crossover method. For the crossover method we monitor the χ_{co} value within the simulation at each timestep. When this exceeds γ_{co} destabilisation has occurred, the IAs are removed and the simulation terminated. The cost up to this point represents the cost of a minimum intervention. If this condition is not met by the end of the simulation then the run is deemed unlikely to destabilise and is marked as invalid. We require 2/3 of the runs to be valid for the minimum interventions to be considered representative and the average minimum cost over the valid runs is then used. In both cases (the *a priori* and crossover methods) we assume placement by degree and that nodes are priced by degree as they were above. The results of this are shown in Figure 5.16. As can be seen, the costs of minimum intervention generated by both methods are nearly perfectly aligned with one another, lending support to the notion that these definitions of minimum intervention are nearly equivalent. Crossover thresholding has a tendency to underestimate the cost at lower numbers of IAs and to overestimate it at higher values but both of these are likely exacerbated by the level of granularity available in the prior calculation method. Overall though, this result (and similar ones in other topologies) give us confidence that these two approaches will produce roughly similar outcomes.

In the previous sections we have also assumed that information about the topology and agent characteristics, such as degree, is readily available. We now consider the situation where such information is hidden, and all that is known is an *advertised cost* which may or may not be indicative of an agent's influence. In the following experiments, IAs are placed at *high cost* locations, without assuming knowledge of degree. This is similar to our work in Chapter 4 but we assume here that there is no method that can be used to acquire accurate information. We feel this models the unknown nature of influence in the real-world where additional factors beyond just a node's degree allow it to have far more or far less influence than the metrics would imply.

We begin by considering the effect of pricing (and hence placing) agents completely at random, to explore the nature of minimum intervention cost when no information is available on agent influence. We assign each node a random cost between $[0, 1]$ and place IAs at the highest cost locations. We perform 100 runs for each number of IAs and use the crossover thresholding described above to find the minimum cost of intervention for each, only considering the runs valid if 66% of them exhibit the crossover.

Figure 5.17 shows the results of this in the static scale-free and small-world

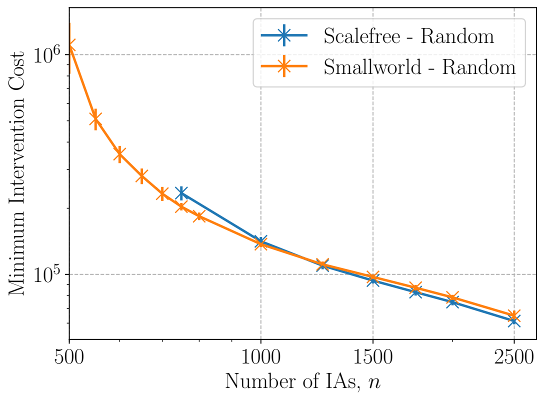

Figure 5.17: Number of IAs vs. the minimum intervention cost needed to cause destabilisation in the static networks when IAs are placed by cost and cost is random.

networks. The results show two main things. Firstly, that, even when placed completely at random, increasing the number of IAs reduces the cost of minimum intervention. This allows us to conclude that by far the predominant factor is the *number* of IAs rather than how they are placed. Whilst placing by metrics increases the efficacy of IAs, increasing the number is the primary driver to facilitate rapid destabilisation of established conventions. This notion is of particular importance when considering interventions that have a fixed budget rather than a fixed number of IAs and we explore this concept further in Section 5.7. Secondly, that random placement is more effective in small-world topologies than in scale-free, the latter not reaching an acceptable level of minimum intervention until much higher numbers of IAs. This threshold of performance indicates that metric-based placement is still beneficial, allowing minimum interventions at ranges that purely random placement cannot.

We additionally consider high-cost placement in the dynamic topologies. As before we must address the nature of dynamism and its effect on cost. We take the same approach as before with degree-based costing such that the cost for agents is not fixed but the cost that is recorded each timestep is that which the placement approach last observed for the IA in question. Thus, for static

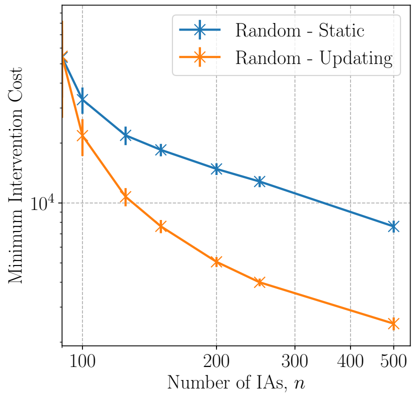

(a) González

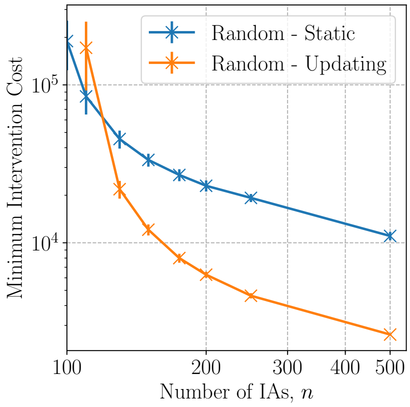

(b) Ichinose RR

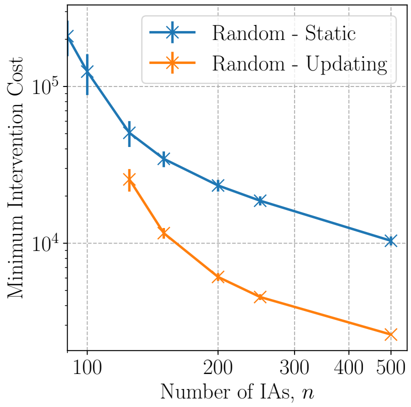

(c) Ichinose RR

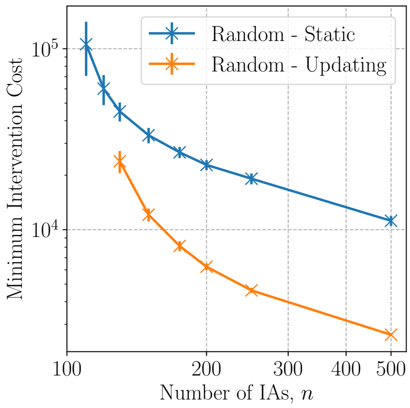

(d) Ichinose RR

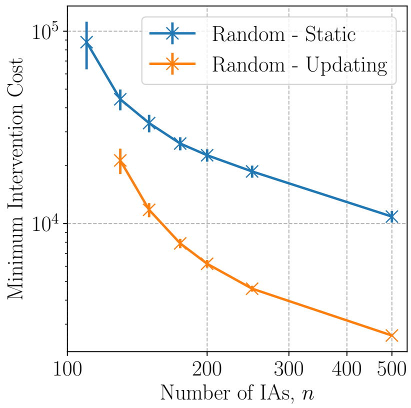

(e) Ichinose RR

Figure 5.18: The minimum cost of intervention for differing numbers of IAs when placed at high-cost locations for the different dynamic topologies. Costs are determined randomly.

high-cost placement, the cost is fixed once the IA has been selected whereas in updating high-cost placement new costs and hence IAs will be chosen every timestep, capturing the random nature of cost in two ways. Figure 5.18 shows the costs of minimum intervention in each of the dynamic topologies as where IAs are selected by highest-cost in both static and updating manners. The overall results are the same in each topology with updating high-cost placement resulting in substantially lower costs compared to static high-cost placement. This is likely due to the IAs being placed nearly continuously at different locations, allowing the destabilisation to spread much faster. However, the number of IAs needed to facilitate this is generally higher with several of the topologies showing that lower numbers of IAs don't allow updating high-cost random placement to meet the threshold required. This is likely due to exactly the same reason that it performs better when it does exceed the threshold; with fewer IAs having them constantly change location gives enough chance for the dominant convention to recover and hence destabilisation is less likely to occur. Again, however, we find the same effect in all topologies: maximising IA numbers continuously reduces cost, decreasing the time taken for minimum interventions even in the crossover threshold model.

Finally we examine the situation where the advertised cost of an agent is an imperfect indication of their degree (and hence influence). This pricing mechanism is useful in domains where agents may be asked to estimate their own influence or domains with unreliable information and covers the ground between the fully random costings just explored and known, precise degree-based placement. We model this by selecting each agent's advertised cost from a Gaussian distribution:

$$cost(v) = \mathcal{N}\left(deg(v), (deg(v) \times \text{noise})^2\right) \tag{5.2}$$

such that higher degree nodes have a wider range of values they might report compared to lower degree. We base this on the notion that it becomes harder and

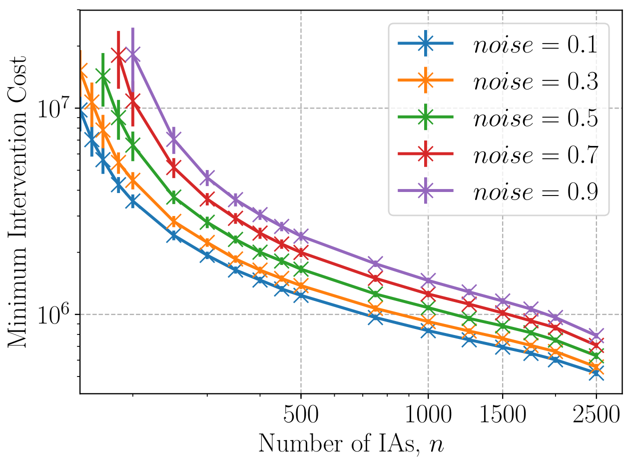

Figure 5.19: Number of IAs vs. the minimum intervention cost needed to cause destabilisation in the 5000 node scale-free network when IAs are placed by highest-cost. Cost is a noisy estimate of the degree of a node.

harder to accurately gauge influence as it increases whereas knowing the links of a minimally connected individual allows a much more nuanced estimation on their influence.

Figure 5.19 shows the results of IAs being placed at high-cost locations with cost being a noisy indicator of degree. We vary the noise level between 0.1 and 0.9 in steps of 0.2 to see how the cost of minimum intervention reacts to increasing levels of noise. We find that the behaviours shown here are the same as in the small-world network with the same constraints and so we focus only on the scale-free network for our discussion. As before, we measure minimum interventions using the crossover threshold model and perform 100 runs for each setting, presenting the averaged cost of the runs if more than 66% exhibited crossover.

The results show that increasing noise has a marked influence on the cost of minimum interventions with increases in noise increasing the cost similarly. This is to be expected as the IAs are placed at high-cost locations but the regular amounts by which the increases in cost occur shows that the noisy nature of the cost function is not causing the destabilisation efforts of the IAs to change dramatically; they are still able to facilitate the removal of the dominant convention even when placed sub-optimally at locations whose influence does not necessarily match their advertised cost. However, this is not to say it has no effect as at the higher levels of noise low numbers of IAs are unable to meet the threshold required for our crossover model, indicating that at these levels the noisy nature is causing destabilisation efforts to fail more frequently. This effect disappears rapidly however with ~200 IAs enough to allow destabilisation to occur frequently again regardless of the noise level. We find that, regardless of noise, the same general pattern exists with small increases in the number of IAs causing dramatic decreases in the average cost of intervention. Indeed, as IAs increase, the difference between the various levels of noise become less and less substantial indicating that even with noisy information, more IAs are able to overcome this and reduce the time needed for intervention. We can thus conclude that destabilisation efforts in the static networks are resilient against noisy or misleading information with the relationship between minimum cost and number of IAs.

Figure 5.20 shows similar effects in the González network as a representative example of how noisy degree costing affects minimum interventions in the dynamic networks. We find that the patterns here are repeated universally amongst the Ichinose models and so we only include the González version here. Static high-cost placement is fundamentally the same as in the static networks with increases in noise causing a corresponding increase in cost but otherwise not affecting the fundamental nature of destabilisation by minimum intervention. Indeed, the dynamic networks seem more resilient to the noisy nature, particularly at fewer numbers of IAs, with no noticeable pattern of destabilisation efforts being more or less likely to succeed; static placement with noise 0.1 has the same minimum starting IAs as with noise 0.9. However the behaviour of updating high-cost placement is quite distinctive with a wide ranging minimum cost at low levels of IAs and a number of higher levels of noise causing some numbers of IAs to be unable to facilitate destabilisation to the threshold

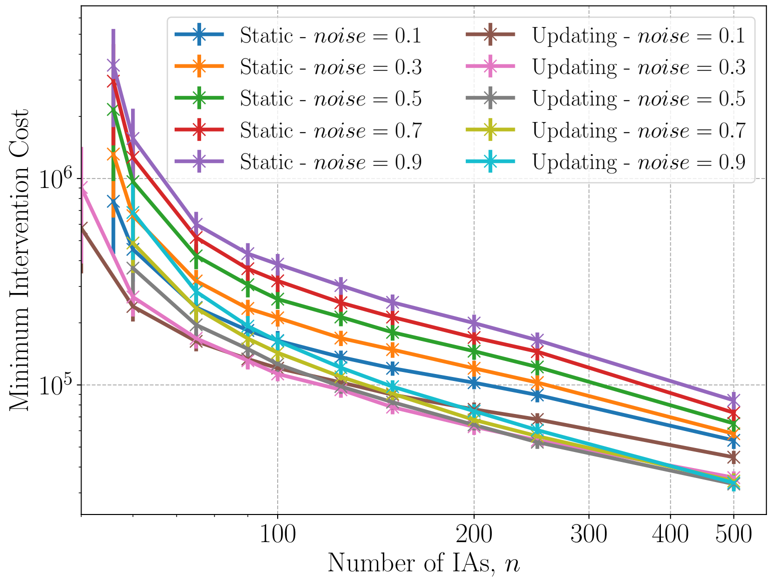

Figure 5.20: Number of IAs vs. the minimum intervention cost needed to cause destabilisation in the González network when IAs are placed by highest-cost. Cost is a noisy estimate of the degree of a node.

level. However, as the number of IAs increases this behaviour rapidly changes, unlike in the static placement case, with the intervention costs for all values of noise converging together. This indicates that, beyond the initial number of IAs updating placement is actually even more resilient to noise than the static equivalent, likely due to the placement mechanism not being locked-in to its decisions for long, meaning that an incorrect assessment is likely to be rapidly balanced out. Overall however, we see again this rapid decrease in minimum cost from even minor increases in the number of IAs. In both static and dynamic networks the noisy nature of degree information is not enough to overcome the benefit of utilising additional IAs to cause rapid destabilisation.

Algorithm 5 Budgeted Placement heuristic

```
procedure BudgetedPlacement(G = (V, E), b, p)
    degRanking ← vertices ranked from highest to lowest degree
    highDeg ← first |V|/2 entries from deg ranking
    lowDeg ← reversed ordering of last |V|/2 entries from deg ranking
    lowBudget, highBudget ← p × b, (1 − p) × b
    create empty set, selected
    for all vertices v in highDeg do
        cost ← deg(v)
        if highBudget − cost < 0 then
            continue
        else
            add v to selected
            highBudget ← highBudget − cost
    lowBudget ← lowBudget + highBudget
    for all vertices v in lowDeg do
        cost ← deg(v)
        if lowBudget − cost < 0 then
            continue
        else
            add v to selected
            lowBudget ← lowBudget − cost
    return selected
```

5.7 Budgeted Placement in Temporary Interventions

In the previous sections we have shown that the cost of intervention can be minimised by including more agents but for a non-linear amount of shorter time. We now seek to apply this understanding to the concept of *budgeted interventions*. In a budgeted intervention, instead of being able to assign a set number of IAs, as we have considered previously, the intervention is instead given a set budget which it can use to acquire agents to act as IAs. As before, each agent has a cost and we consider the case where the cost is equal to their degree, without noise, as a good measure of the influence capabilities of each agent.

We seek the best way to utilise the given budget to maximise the likelihood of destabilisation of the established convention. The inclusion of the budget adds an additional constraint from the previous section as we cannot simply add more and more high-degree nodes but instead must consider how the inclusion of large numbers of cheap but lowest-degree nodes might be used to facilitate destabilisation.

We introduce a new placement heuristic, BudgetedPlacement, which is shown in Algorithm 5. The basic function of BudgetedPlacement is, given a graph, G, a budget, b, and a minimum proportion of this budget that must be spent on low degree nodes, p, select a set of nodes to act as IAs whose total cost is less than or equal to the budget. It does this by acquiring a ranking of all nodes by degree, splitting this list in two and then greedily adding high-degree nodes from the first list that are within the still available budget for high-degree nodes and doing the same but in ascending order from the low-degree list and for the low-degree budget. These budgets are found by dividing up the total budget, b, based on the proportion, p, and this approach ensures that there is a maximum amount that can be spent on either type of node. Any unused high-degree budget is additionally allocated to be spent on low-degree nodes.

BudgetedPlacement allows us to see how the overall budget affects the ability to cause destabilisation as well as how varying the budget spending between high-degree and low-degree nodes affects the outcome. To establish this we vary the budget available and the proportion assigned to low-degree nodes and utilise the IAs selected by BudgetedPlacement to attempt to destabilise the dominant convention in both the scale-free and small-world topologies. To allow easier comparisons, we use 1000 node variants of each of the graphs but our experimental setup is otherwise the same as has been used throughout this chapter. We perform 100 runs for each setting combination and measure the proportion that result in the replacement of the dominant convention with the one assigned to IAs. We find that, due to the large number of IAs selected at higher values of p, when using the $K_{90\%}$ threshold the scale-free simulations exhibit the problem identified in Section 2.3.1 where IAs still reporting the dominant strategy (due to being unable to unlearn it) are too numerous and mean that 90% thresholding is impossible even if all non-IAs are members of the new convention. As such, for those results we utilise the 80% Kittock criteria.

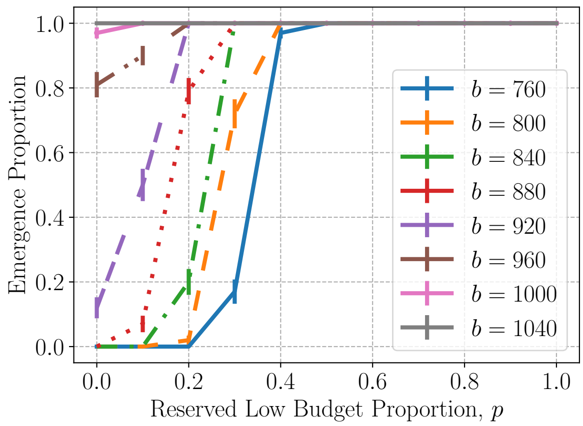

Figure 5.21: Effectiveness of BudgetedPlacement when used for destabilisation in the 1000 node scale-free topology.

Figure 5.21 shows the effectiveness of this strategy in the 1000 node scale-free network. The budget ranges shown represent the limits where, with $p = 0.0$, destabilisation occurs not at all or is guaranteed to do so. With ~3000 edges in the graph (and hence a total degree of ~6000), these budgets represent between 12.6% and 17.3% of the maximum budget possible and show that, similar to when using a fixed number of IAs, destabilisation is possible with only small proportions of the network being made into IAs. The change between no effect and fully effective again occurs over a narrow range, with only an increase of 280 (or 4.6%) needed to cause it, similar to the narrow range of effectiveness change found when targeting specific numbers of IAs. This highlights the notion of the "critical value", which destabilisation attempts need to exceed, and doing so by even relatively small amounts will nearly guarantee success.

More interesting however is the effect that increasing p has. For all budgets, even those which do not exhibit destabilisation when $p = 0$, increasing p to 0.5 or beyond means that they instead cause destabilisation 100% of the time. Indeed, for most budgets, increasing p by any amount causes a commensurate increase in the effectiveness of BudgetedPlacement at that budget. This indicates that our previous findings regarding the number of IAs is not limited to larger

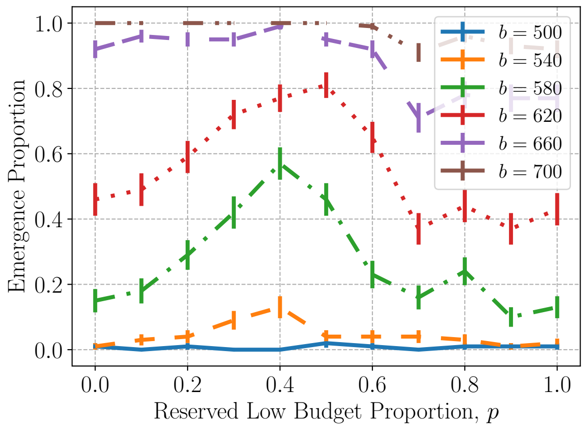

Figure 5.22: Effectiveness of BudgetedPlacement when used for destabilisation in the 1000 node small-world topology.

numbers of high-degree nodes. Including more IAs, regardless of their degree, is beneficial in the scale-free network and makes destabilisation easier as the desired convention spreads to other nodes more rapidly. Additionally it shows that the ability to destabilise is not tied to simply increasing the total degree of the IAs (represented by increasing budgets) but *how* the IAs are spread out, which has major implications for convention destabilisation and gives credence to the effectiveness of "grassroots" movements.

Figure 5.22 shows the results of using BudgetedPlacement in the 1000 node small-world network. It shows a similar transition from no effectiveness to full effectiveness over a narrow range of budgets as in scale-free. However, the actual budget needed to guarantee destabilisation with $p = 0$ is substantially less than that needed in the scale-free networks, despite the number of edges in both networks being ~3000. This runs contrary to our previous destabilisation findings where small-world topologies typically required much larger numbers of IAs compared to scale-free in order to guarantee an effect. The greedy selection of BudgetedPlacement allows more IAs to be selected due to the lack of a power-law degree distribution in small-world topologies. Whilst a given budget may allow a few high-degree nodes in scale-free distributions the same budget will allow many more nodes in small-world topologies due to the less skewed degree distribution. This further reinforces the point from the scale-free network that multiple nodes of slightly lower degree (and hence cost) are more effective than singular high-degree nodes.

The value of p again has a major effect on the performance of BudgetedPlacement but it is quite distinctive in small-world networks compared to scale-free. Continuously increasing p does not produce an equivalent increase in effectiveness. Instead, the optimal value of p peaks at between 0.4 and 0.6 for each budget with further increases detracting from the performance. In small-world networks there is a balance between higher degree nodes and lower degree nodes that must be found to optimise destabilisation. This again can be attributed to the much less skewed nature of the degree distribution in small-world networks; whereas in scale-free networks BudgetedPlacement must choose between a few high-degree nodes or many low-degree nodes, in small-world the choice is between many higher degree nodes or many lower degree nodes without the necessity of extremes that comes from the power-law degree distribution in scale-free networks. However, it still shows that a slight increase in the minimum budget that is allocated to low-degree nodes has beneficial effects.

We similarly apply the notion of a budgeted intervention to the dynamic networks. The function of BudgetedPlacement is identical in this domain except for two differences: (i) when a selected node is removed as an IA the cost paid for it is returned to the relevant budget and a new node of the same type (high-degree or low-degree) is selected, and (ii) any unused high-degree budget is not added to the low-degree budget. This is to prevent the "leaking" of budgets from one to the other over time as it is unlikely that a node of exactly the same cost will be selected to replace one. We also consider both static and updating BudgetedPlacement, as we have done with other dynamic placement heuristics, where the former has its choices set as IAs until they are removed from the graph itself whilst the latter removes all IAs each timestep

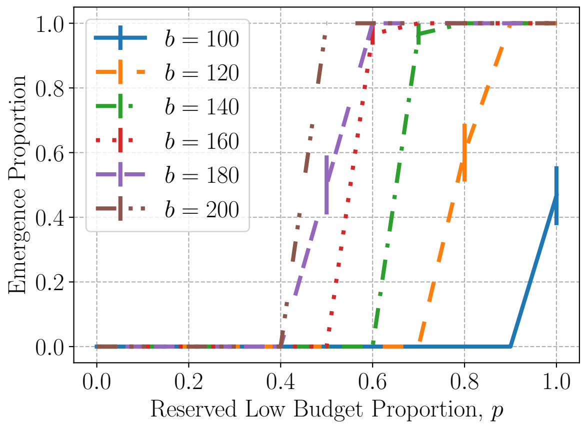

Figure 5.23: Effectiveness of BUDGETEDPLACEMENT when used for destabilisation in the 1000 node González, static placement.

(and hence starts the BUDGETEDPLACEMENT again). We use the same dynamic networks as have been used throughout with the IAs being introduced at $t = 1500$ and the simulation running for 5000 timesteps beyond this to give enough time for destabilisation to occur and calculating the proportion of 30 runs that emerge the desired convention.

Figure 5.23 shows the effect of BUDGETEDPLACEMENT when used in a static manner in the 1000 González model. Immediately, there are dramatic differences between the static networks and this one, primarily the range of budgets that exhibit change. With budgets as low as 100 (representing 0.016% of the possible total degree) BUDGETEDPLACEMENT is able to cause destabilisation with high values of p. Indeed, we find that the range of budgets that exhibit destabilisation when $p = 0.0$ is substantially higher than the range that exhibits it when $p = 1.0$ or even when $p = 0.5$ being closer to the numbers found in the static networks. However, within the dynamic networks, including more low-degree nodes has a significant effect even at very low budgets. Investigation reveals that this is because there are a number of 1-2 degree nodes within the dynamic networks due to the edge churn and nature of them that simply aren't present in the static networks where each node has a minimum degree of 3 (in the scale-free

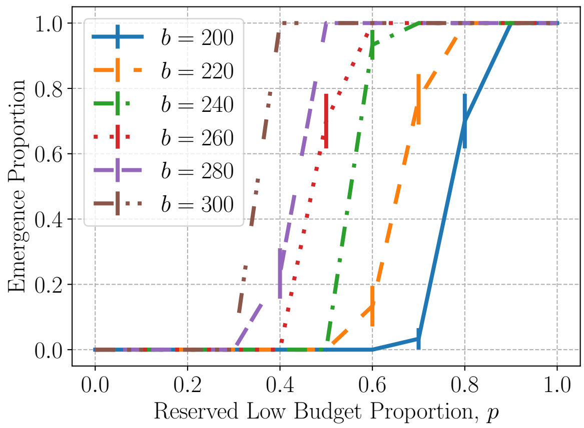

Figure 5.24: Effectiveness of BudgetedPlacement when used for destabilisation in the 1000 node González, updating placement.

network) and 5 (in the small-world network). The large numbers of these low degree nodes mean that many of them can be acquired as IAs even with low budget and, as we saw in the previous section, despite not being topologically influential the dominating factor is their number. This highlights again the primary finding of this chapter, that, wherever possible, you should maximise the number of IAs in order to most readily cause destabilisation. The approach of BudgetedPlacement as a method of selecting these IAs is best placed to enable this aspect of destabilisation.

We also utilise BudgetedPlacement in an updating manner the results of which are shown in Figure 5.24. As we see here, utilising BudgetedPlacement in an updating manner is actually detrimental to its effectiveness with high values of p, requiring much larger budgets to exhibit the same level of effect as is present in the static version. This is likely due to the updating placement heuristic constantly changing which low-degree nodes it selects as IAs and highlights one of the issues with low-degree nodes over high-degree nodes: they are more readily influenced by their own environment rather than influencing it themselves as they are connected to only 1-2 other nodes and hence limited in their interactions. If they are unable to convert their neighbours, which is unlikely due to the neighbours being statistically likely to have more edges, their influence is limited and they will tend to whatever action that neighbour is choosing to reduce clashes. As such, for budgeted interventions it is better to select low-degree nodes and leave them as IAs. This "multiplies" their influence due to not changing their action due to clashes, which forces their limited neighbours to do so instead and allows them to spread the convention rapidly. We see nearly identical behaviour in the Ichinose models, which also exhibit very low degree nodes due to the changing connections, and hence do not include them here. The primary difference between the results for them and the results for the González model is that the Ichinose models are require slightly higher budgets across the board, closer to $b = 200$ to exhibit static placement change and closer to $b = 400$ to exhibit change when using updating BudgetedPlacement. Both of these increases are likely due to the reduced churn of the Ichinose models which results in fewer degree 1 nodes than there are in the González model.

Overall we have shown the effectiveness with which BudgetedPlacement can cause destabilisation in situations where we are provided with a finite budget, b. We have shown that increasing the amount of this budget that we reserve for low-degree nodes, p, is nearly universally beneficial and that choosing a value of $p = 0.5$ should be sufficient to provide generally applicable optimal behaviour.

5.8 Conclusions

Using IAs, agents that continuously apply a given fixed strategy regardless of the consequences, have previously been shown to help direct and encourage convention emergence. In Chapter 3 we had previously expanded the state-of-the-art in using these agents to facilitate convention emergence as well as to use them to destabilise existing conventions to enable them to be replaced.

In this chapter we expanded further on that work by considering the notion of how time-limited inclusion of IAs affects their ability to enact the desired changes. We introduce the notion of *temporary interventions* where IAs are only placed in the system for a finite, heavily-limited time. Being able to effect a permanent change in the system with minimal intervention is highly desirable and understanding temporary interventions is a necessity to this end. The self-reinforcing nature of convention emergence and destabilisation means that with an initial "nudge" in the right direction, even a temporaneous one, we believe that the same level of effectiveness as when IAs are included permanently can be achieved. Finding the thresholds and boundaries of effectiveness that these temporary interventions have is important as it allows groups such as campaigners and marketers to focus efforts on when these interventions would be most effective rather than having to enact a permanent inclusion.

We began by studying temporary initial interventions; the amount of time that IAs must be present at the start of the simulation in order to cause a permanent shift that guarantees the emergence of the desired strategy as convention. We showed that, across all topologies and network types IAs need only be present for a very short time period (<150 timesteps) before there is no difference in their efficacy than if they had been included permanently. We studied the effect that additional IAs beyond the minimum required had on this minimum inclusion time and found that increasing the number of IAs reduced the minimum though in a manner of diminishing returns.

We then considered the effectiveness of temporary interventions when not applied at the very beginning of the simulation, so-called *staggered temporary interventions*, to establish the change in effectiveness amongst IAs when introduced later. We found that IA effectiveness rapidly decreases outside of starting at time $t = 0$ and surmise that the early stages of convention emergence are amongst some of the most important as the force of precedence rapidly grows and then will require additional effort to overcome. This lends credence to the notion of groups being highly resistant to change once something has become established and has implications for those seeking to change the public mindset over any issues, for instance political ones.

We have also shown that temporarily inserting IAs can also cause destabilisation, *temporary late interventions*, and that there exists a minimum length of time that they must be present in order to cause this. Removing IAs prior to this minimum duration will cause the established convention to return to near previous levels. We showed that increasing the number of IAs even slightly has a dramatic effect on reducing this minimum time with further increases causing a diminishing return in reduction. In systems that are closely modelled by our approach here this means that so-called “grassroots” efforts can be highly effective, more so than deferring to authorities within the system.

Next we considered the cost of these interventions, and show that, independent of whether cost is uniform or linked to degree, the cost of minimum intervention is inversely related to the number of IAs. However, the relationship is also one of diminishing returns. As such, placing as many IAs as possible into the system is beneficial but the additional effect generated reduces substantially after ∼10% of the population. Wherever possible however, our findings indicate that increasing the number of IAs causes faster and more robust destabilisation.

We then explored the effect of placing IAs by cost and monitoring destabilisation in real-time. The same relationship between number of IAs and cost was found to hold regardless of pricing/placement mechanism although higher numbers of IAs may be needed to sufficiently guarantee destabilisation. The effect of noise on the degree-based pricing mechanism was also considered. It was found, for all topologies, that the effect of noise was to increase the overall cost of minimum interventions but to not affect the relationship between cost, the number of IAs, and the duration of minimum intervention. We conclude from this that placing by advertised cost would offer reasonable results, even with near-random pricing and that destabilisation efforts are resilient to misinformation as long as the number of IAs is sufficient.

We then utilised our findings to investigate the best strategy for budgeted interventions where there is no limit on the number of IAs but rather a specific cost associated with each node that must be accounted for within a finite budget. We introduced a placement heuristic BudgetedPlacement and showed that guaranteeing a certain proportion of the budget to be spent on low-degree agents rather than high-degree dramatically increases the effectiveness of the intervention by increasing the number of IAs even if they not highly-ranked locations based on influence. We find that the *number* of IAs is as important if not more so than the location they are placed and that, in scenarios where the number of IAs is not set, maximising it is paramount.

Overall we have shown that encouraging emergence and destabilisation and replacement of an established convention is possible and that minimum criteria exist in order to cause this. We have also presented a number of ways of evaluating how much an intervention might cost using various pricing methods and demonstrated the relationship between the number of IAs and cost.

CHAPTER 6

Conclusion

In this thesis we have explored a number of unique and novel aspects concerning the emergence of conventions in multi-agent systems (MAS) and how this emergence might be directed to facilitate rapid and robust convergence to coordinated behaviour. We focused, broadly, on the use of Intervention Agents (IAs) as fixed strategy agents in order to provide guidance and influence to the rest of the population in order to encourage the system to emerge the desired convention under a range of topological conditions. In this chapter we review the contributions made by this thesis and evaluate them within the framework of the original aims of this body of work. We then consider some of the ways in which limitations of the work presented could be expanded upon in future research before presenting our final thoughts.

Whilst drawing from large parts of the state-of-the-art, the work in this thesis builds heavily on three primary pieces of the literature: the work of Franks et al. [2013], Sen & Airiau [2007], and Villatoro et al. [2009]. We will briefly compare and contrast the work in this thesis with each of these before moving on to the contributions in more detail. Sen & Airiau [2007] was one of, if not the, earliest works utilising fixed strategy agents to inform and manipulate convention emergence. We have extended their early work in a number of directions but the primary difference is that of using these fixed strategy agents not just to direct initial convention emergence but later in the convention lifecycle to manipulate and change already established conventions, showing the differences inherent in these two different applications. Building on the work of Sen & Airiau, Villatoro et al. [2009] is some of the seminal work in considering the underlying network topology and its effect on convention emergence. Their analysis is rather limited, only using small, synthetically-generated networks, but showed the major effect topology could have on the simulation dynamics. We have expanded and extended this work in a number of directions, considering larger scale synthetic networks, real-world networks and dynamic, time-varying networks that Villatoro et al. never touched upon. We have shown that this effect of topology is universal and can have a dramatic effect on convention emergence but that conventions emerge nonetheless. The work of Franks et al. [2013] comes closest to the work of this thesis in a lot of ways, considering more metrics and different network structures than any of the work before. We have continued in the same vein, exploring similar metrics and their applicability to the coordination game and extending the Influencer Agents of Franks et al. into the realm of destabilisation and general intervention. The exploration of dynamic networks extends the work of Franks et al. into previously unseen domains and our findings there have expanded the state-of-the-art.

Each of the chapters presented in this thesis has explored an important part of our original aims for research. In Chapter 3 we extended the knowledge on convention emergence and how it might be influenced as well as garnering an understanding of convention stability and how they might be destabilised in a range of static and dynamic topologies. In Chapter 4, we considered realistic restrictions on the amount of topological knowledge available when trying to place IAs and examined the effect this had on efficacy, providing solutions to mitigate this problem. Finally, Chapter 5 provided insight into the minimal nature of intervention that can be used to elicit change and how to utilise this information to increase performance, fulfilling our final objective.

6.1 Contributions

This thesis has made a number of contributions in its exploration of conventions, extending our understanding of them and how IAs may be used to manipulate them. We have shown that destabilisation is possible and how it can be achieved and investigated in both static and dynamic topologies. Additionally we have considered different constraints that might be placed upon convention emergence, such as observability and time, and explored their impact. In more detail, this thesis makes the following contributions:

- **Introduction and analysis of the concept of destabilisation of an existing convention and techniques to facilitate this.**

  In Chapter 3, we introduced the notion of *destabilising* an existing convention, causing it to collapse in support amongst the population and to be replaced by another. We demonstrated that established conventions are more resilient to outside influence than those earlier in their life cycle due to the self-reinforcing nature of conventions and the force of precedence the dominant convention enjoys. Small proportions of the population were utilised as IAs, successfully facilitating the emergence of a secondary convention to reduce the precedence of the dominant one and persuade other agents to switch. The results indicate that this proportion is relatively stable even across multiple scales of network size and identified ways in which the complex structure of the underlying topology impacts destabilisation efforts. We showed that *where* IAs are located within this topology heavily impacts the number required to cause destabilisation and that there is a critical number of IAs which exhibits a sharp transition between no effectiveness and full effectiveness at causing destabilisation. An evaluation of the performance of IAs when placed by a number of topological metrics was undertaken and showed that, universally, degree or eigencentrality performed best. From this we can conclude that these metrics offer a strong measure of influence and act as integral locations to maximise the disruption caused to the dominant convention. These findings were evaluated on a range of static networks, both real and synthetic, showing their general applicability. Additionally, we introduced the concept of *passive destabilisation*, removing an existing convention without directly replacing it with another and show this is possible using IAs of multiple different strategies. Our findings show that facilitating destabilisation in this way requires a larger effort and proportion of the population as IAs. The investigations from this chapter are slightly limited however, due to the same interaction model, that of the social learning coordination game, being used in all settings and would require validation against a wider range of models and games.

- **An exploration of convention emergence in dynamic topologies and the creation of placement metrics to encourage convention emergence and destabilisation in these topologies.**

  Additionally in Chapter 3, we introduce the usage of dynamic topologies, where edges and nodes may change over time, to model the links between the individual agents in the population. These are known to induce different system dynamics than those caused in static networks and thus we explored the nature of convention emergence in these types of networks. The results show that the changing nature of the edges within the graph allows conventions to emerge rapidly, unaided within the populations as long as minimum levels of connectivity and agent longevity are met. We showed that these populations are able to be directed using small numbers of IAs and that the numbers required were in general comparable to those needed to do the same in static networks. Two methods of applying the traditional placement metrics were developed, Static and Updating, and the efficacy of these in influencing convention emergence was explored. Additionally, a new heuristic, Life-Degree, was created to explore the impact that considerations of agent longevity have on viability of agents as IAs. This showed that longevity was not a primary concern and that the measure of influence an agent exhibits is of higher priority in their efficacy as IAs.

  We additionally applied the concept of destabilisation to dynamic networks. The results show that, in general, dynamic networks are more resilient to destabilisation attempts, requiring a larger proportion of IAs to elicit the same level of effect as in static networks. The updating placement metrics were found to substantially outperform their static equivalents indicating the importance of up-to-date measures of influence when attempting to destabilise the established convention. Across all topologies, we showed that placement by updating degree and eigencentrality consistently offered better performance than the naïve static placement approach.

We explored the effect that size of convention space and payoff matrix may have on convention emergence in dynamic networks showing that they are sensitive to both with larger convention spaces making convention emergence difficult. The efficacy of Life-Degree when considering the destabilisation of conventions within these large convention spaces was shown, highlighting the importance of agent longevity in these instances.

Our work here is limited by two key factors: (i) whilst we explore a number of different dynamic network models, the population size is the same in all of them (1000 nodes) to allow comparison. Larger populations must be examined to assure ourselves of the general applicability of these findings. (ii) the synthetic nature of the dynamic topologies may well cause them to differ from real-world dynamic networks; establishing the applicability of our findings in this domain is a separate problem.

- **Development of placement algorithms that find influential locations in topologies with restricted observability.**

  In Chapter 4 we consider the impact of partial observability of the network topology on the use of IAs for influencing convention emergence. We show that this restriction can have a large negative impact on the efficacy of IAs and that careful use of a limited number of observations can help to mitigate this by identifying influential individuals based only on local information. A placement algorithm, PO-Place, was developed that utilises finite observations to approximate globally important locations. We show that PO-Place, with observations of only small proportions of the overall network topology, is able to frequently achieve performances comparable to those when given full graph observability. An analysis of the effect of a number of different constraints on the performance of PO-Place was performed and show it to be robust in a range of scenarios. PO-Place was then utilised for placing IAs for convention emergence and destabilisation in a range of real-world and synthetic networks and found to elicit similar levels of convention emergence to the previous, unrestricted case with small levels of observation.

We additionally created DynaPO to solve the equivalent limited observation problem in dynamic topologies, identifying a range of different issues and concerns that arise when considering partial observability in these topologies and analysing their likely effects. We showed that the potential variants of DynaPO offer different levels of impact in initial and late intervention and analyse how DynaPO performs under a number of different settings. The results show that even fewer observations are needed than in the case of PO-Place for equivalent levels of convention emergence and destabilisation with only slight increases in the proportions of IAs needed than in the fully observable case.

- **Mechanisms for assessing minimal interventions, their costs and an investigation into their effectiveness at manipulating conventions.**

  In Chapter 5 we show that temporary rather than permanent interventions using IAs are able to elicit the same level of convention emergence and destabilisation. An exploration of the minimum time that such interventions must take place in order to guarantee convention emergence is undertaken and shows that the initial period of convention emergence is the most important with IAs only required for a short period to direct the permanent change of the population. We highlight the importance of this initial period again by exploring intervention at differing times in the early stages of convention emergence and show a rapid decline in efficacy. An exploration of the notion of minimum interventions for *destabilisation* of conventions is performed and shows that there are both minimum numbers of IAs required to cause destabilisation and a minimum length they must be present to prevent rebounding of the dominant convention. The concept of the *cost* of an intervention is introduced and we use this mechanism to analyse the effect of increasing IAs on the minimum length required for destabilisation, showing that, universally, increasing the number of IAs had a more than commensurate decrease in the amount of time taken for destabilisation. We established a real-time, on-line notion of minimum intervention and showed the resilience of our conclusions to noise in the system. This knowledge of minimum interventions was then used to develop an approach for BudgetedPlacement of IAs and showed the effectiveness of our method at even low levels of budget.

6.2 Directions for Future Work

We have established and explored convention emergence and destabilisation in a number of different settings and domains. However, a number of unanswered questions still exist on aspects of convention emergence that are open to manipulation, change and efficiency increases. In particular, there are considerations along a number of different axes for methodologies that might increase the effectiveness of IAs in effecting convention emergence. We present some of these below as potential directions for expansion of the work presented in this thesis.

6.2.1 Discount Heuristics for Intervention Agent Placement

Identifying influential locations for IAs placement can be done easily using the degree placement metric. However, in highly connected networks, it is likely that many high-degree nodes will share neighbours. As such, their influence is slightly redundant in nature and in order to facilitate rapid convention emergence we may be better placing IAs at diverse locations in order to maximise the number of unique entities that they can reach. In the related work of influence spread [Kempe et al., 2003], Chen et al. [2009] identify similar concerns regarding redundancy amongst multiple agents. They produce a degree discount heuristic which takes this redundancy into consideration when selecting additional agents and they show that this causes a marked improvement in outcomes. Whilst the underlying models differ substantially we believe the high-level concepts may be equally applicable and may allow better maximisation of efficacy when considering placement of a set number of IAs.

6.2.2 Predicting Influential Nodes

The use of dynamic topologies in the study of convention emergence is a particular focus of this thesis. Within these dynamic topologies the influence of a node is able to change over time as additional edges come and go. Being able to predict which nodes are likely to become or remain influential would facilitate targeting these nodes early so that they act as propagators of the desired convention without the need to create them as IAs later and at lower cost than doing so as well. A number of methods have been proposed for the identification of nodes likely to reach prominence [Yang et al., 2014] but it is unclear whether these approaches would benefit the study of convention emergence and how the information they provide may be best utilised to reduce costs of intervention.

6.2.3 Graph Partitioning and Community Identification

Throughout this thesis we have noted the difference and resilience of small-world topologies to emerging conventions compared to simpler and faster convention emergence in other topologies. This is due to the locally clustered and connected nature of the small-world network [Franks, 2013; Kleinberg, 2000b] and highlights that for many topologies with local clustering, placing at highly-influential locations outside of this community structure will have less efficacy at converting this portion of the population than placing inside the community. The importance and identification of these community structures has been previously explored in the notion of cooperation in multi-agent systems [O'Riordan & Sorensen, 2008] but its use to ensure convention emergence or destabilisation has not been considered. A number of techniques exist for community identification [Danon et al., 2005] and investigating the effectiveness of these could increase the efficacy of IAs. Related to this problem is the concept of graph partitioning, finding subsets of the graph such that the edges between them are minimised. This has been well-understood in parallel computing for decades [Karypis & Kumar, 1998; Simon & Teng, 1997] but its use in identifying subsections of the graph that are least likely to interact with each other could be used to highlight which areas need focused effort to spread conventions to.

6.2.4 Real-world Dynamic Networks

We have used a number of dynamic networks in this thesis to investigate the effect of dynamism of convention emergence. However all of these are synthetic dynamic networks. Whilst efforts have been made to ensure they closely model features from the real-world [González et al., 2006a; González et al., 2006b], it has previously been noted that static synthetic networks differ substantially from real-world ones [Franks, 2013; Pujol et al., 2005] and we expect similar discrepancies are likely between synthetic dynamic networks and real ones. As such, exploring the nature of convention emergence on these real-world dynamic networks is a logical extension and would provide assurances to the general applicability of the results found in course of this work.

6.3 Final Thoughts and Remarks

In this thesis we have examined the ways in which conventions emerge amongst agent-agent interactions in MAS and how they can be manipulated and directed in order to support coordinated behaviour amongst agent populations. We have extended the state-of-the-art on the theory of conventions and introduced mechanisms and tools that can be used to manipulate conventions in the form of Intervention Agents (IAs). We have shown these to be highly effective in the common constraints of many forms of MAS and able to be utilised at multiple stages in the evolution of a convention to bring about differing results as shown by our work in Chapter 3.

Throughout this thesis, one of the most important considerations that impacts convention emergence dynamics is that of the topological structure that agents find themselves in. The form of this interconnecting topology has varied effects from the rate of convention emergence to its stability and dramatically changes the nature of the agent population. With the introduction of dynamic networks we have considered an additional dimension to the topological structures and seen the considerations that such dynamism brings to the fore. In Chapter 4 we see how the limitation of local information can significantly reduce the efficacy of intervention efforts and yet this is how most networks appear to us, without a top-down overview and only able to see a few hops away from ourselves.

Even with these additional concerns and constraints, we have shown that it is possible to encourage rapid and robust convention emergence amongst populations of agents. We have explored the stability of conventions and shown how to replace them if necessary without drastic intervention in the system. We believe these aspects to be a key step in developing a comprehensive model of conventions within MAS and that the work in this thesis contributes to that goal.

Bibliography

[1] Stéphane Airiau, Sabyasachi Saha, and Sandip Sen. "Evolutionary Tournament-Based Comparison of Learning and Non-Learning Algorithms for Iterated Games". In: *Journal of Artificial Societies and Social Simulation* 10.3 (2007), pp. 1–7.

[2] Stéphane Airiau, Sandip Sen, and Daniel Villatoro. "Emergence of conventions through social learning". In: *Autonomous Agents and Multi-Agent Systems* 28.5 (Sept. 2014), pp. 779–804.

[3] George A. Akerlof. "A Theory of Social Custom, of Which Unemployment May be One Consequence". In: *The Quarterly Journal of Economics* 94.4 (1980), pp. 749–775.

[4] Reid Andersen et al. "Local Computation of PageRank Contributions". In: *Algorithms and Models for the Web-Graph: 5th International Workshop, WAW 2007, San Diego, CA, USA, December 11-12, 2007. Proceedings*. Ed. by Anthony Bonato and Fan R. K. Chung. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 150–165.

[5] K.J. Arrow. *The Limits of Organization*. Fels Center of Government Series. Norton, 1974.

[6] W. Brian Arthur. "Inductive Reasoning and Bounded Rationality". In: *The American Economic Review* 84.2 (1994), pp. 406–411.

[7] K. Avrachenkov et al. "Quick Detection of High-Degree Entities in Large Directed Networks". In: *2014 IEEE International Conference on Data Mining*. Dec. 2014, pp. 20–29.

[8] R Axelrod and WD Hamilton. "The evolution of cooperation". In: *Science* 211.4489 (1981), pp. 1390–1396.

[9] Robert Axelrod. "An Evolutionary Approach to Norms". In: *American Political Science Review* 80 (04 Dec. 1986), pp. 1095–1111.

[10] Albert-László Barabási and Réka Albert. "Emergence of Scaling in Random Networks". In: *Science* 286.5439 (1999), pp. 509–512.

[11] Alex Bavelas. "Communication Patterns in Task-Oriented Groups". In: *The Journal of the Acoustical Society of America* 22.6 (1950), pp. 725–730.

[12] Cristina Bicchieri, Richard C. Jeffrey, and Brian Skyrms. *The Dynamics of Norms*. Cambridge Studies in American Literature and Culture. Cambridge University Press, 1997, pp. 47–66.

[13] Guido Boella and Leendert van der Torre. "Substantive and procedural norms in normative multiagent systems". In: *Journal of Applied Logic* 6.2 (2008). Selected papers from the 8th International Workshop on Deontic Logic in Computer Science, pp. 152–171.

[14] Elhanan Borenstein and Eytan Ruppin. "Enhancing autonomous agents evolution with learning by imitation". In: *Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour* 1.4 (2003), pp. 335–348.

[15] Stephen P Borgatti. "Centrality and network flow". In: *Social networks* 27.1 (2005), pp. 55–71.

[16] Tilman Börgers and Rajiv Sarin. "Learning Through Reinforcement and Replicator Dynamics". In: *Journal of Economic Theory* 77.1 (1997), pp. 1–14.

[17] Christian Borgs et al. "Influence Maximization in Social Networks: Towards an Optimal Algorithmic Solution". In: *CoRR* abs/1212.0884 (2012).

[18] Christian Borgs et al. "The Power of Local Information in Social Networks". In: *Proc. of the 8^{th} International Workshop on Internet and Network Economics*. Ed. by Paul W. Goldberg. 2012, pp. 406–419.

[19] Michael Bowling and Manuela Veloso. "Rational and convergent learning in stochastic games". In: *Proceedings of the 17th International Joint Conference on Artificial Intelligence*. Vol. 2. 2001, pp. 1021–1026.

[20] Robert Boyer and André Orléan. "How do conventions evolve?" In: *Journal of Evolutionary Economics* 2.3 (Sept. 1992), pp. 165–177.

[21] Steven J Brams. *Game theory and politics*. Courier Corporation, 2011.

[22] Ulrik Brandes. "A faster algorithm for betweenness centrality". In: *Journal of mathematical sociology* 25.2 (2001), pp. 163–177.

[23] Hannelore Brandt and Karl Sigmund. "Indirect reciprocity, image scoring, and moral hazard". In: *Proceedings of the National Academy of Sciences of the United States of America* 102.7 (2005), pp. 2666–2670.

[24] Michael Brautbar and Michael J Kearns. "Local Algorithms for Finding Interesting Individuals in Large Networks". In: *Proc. of Innovation in Computer Science*. 2010, pp. 188–199.

[25] M.A.R. de Cara, O. Pla, and F. Guinea. "Competition, efficiency and collective behavior in the "El Farol" bar model". In: *The European Physical Journal B - Condensed Matter and Complex Systems* 10.1 (May 1999), pp. 187–191.

[26] Duan-Bing Chen, Rui Xiao, and An Zeng. "Predicting the evolution of spreading on complex networks". In: *Scientific Reports* 4 (2014), p. 6108.

[27] Wei Chen, Yajun Wang, and Siyu Yang. "Efficient Influence Maximization in Social Networks". In: *Proc. of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Paris, France, 2009, pp. 199–208.

[28] Colin Cooper, Tomasz Radzik, and Yiannis Siantos. "A Fast Algorithm to Find All High Degree Vertices in Power Law Graphs". In: *Proceedings of the 21st International Conference on World Wide Web*. WWW '12 Companion. Lyon, France: ACM, 2012, pp. 1007–1016.

[29] Leon Danon et al. “Comparing community structure identification”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2005.09 (2005), P09008.

[30] Jordi Delgado. “Emergence of social conventions in complex networks”. In: *Artificial Intelligence* 141.1–2 (2002), pp. 171–185.

[31] Jordi Delgado, Josep M. Pujol, and Ramón Sangüesa. “Emergence of Coordination in Scale-Free Networks”. In: *Web Intelligence and Agent Systems* 1.2 (Apr. 2003), pp. 131–138.

[32] Sebastiano A. Delre et al. “Will It Spread or Not? The Effects of Social Influences and Network Topology on Innovation Diffusion”. In: *Journal of Product Innovation Management* 27.2 (2010), pp. 267–282.

[33] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of Networks: From Biological Nets to the Internet and WWW (Physics)*. New York, NY, USA: Oxford University Press, Inc., 2003.

[34] Edmund H. Durfee. “Challenges to Scaling-Up Agent Coordination Strategies”. In: *An Application Science for Multi-Agent Systems*. Ed. by Thomas A. Wagner. Boston, MA: Springer US, 2004, pp. 113–132.

[35] Edmund H. Durfee. “Distributed Problem Solving and Planning”. In: *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. Ed. by Gerhard Weiss. Cambridge, MA, USA: MIT Press, 1999, pp. 121–164.

[36] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. New York, NY, USA: Cambridge University Press, 2010.

[37] Joshua M. Epstein. “Learning to Be Thoughtless: Social Norms and Individual Computation”. In: *Computational Economics* 18.1 (Aug. 2001), pp. 9–24.

[38] Facebook Developers. *Facebook: Graph API Reference*. Aug. 2017. URL: https://developers.facebook.com/docs/graph-api/reference.

[39] Facebook Developers. *Facebook: Rate Limiting - Graph API*. Aug. 2017. URL: https://developers.facebook.com/docs/graph-api/advanced/rate-limiting.

[40] Christina Fang et al. "On Adaptive Emergence of Trust Behavior in the Game of Stag Hunt". In: *Group Decision and Negotiation* 11.6 (Nov. 2002), pp. 449–467.

[41] Michael P. Fay and Michael A. Proschan. "Wilcoxon-Mann-Whitney or t-test? On assumptions for hypothesis tests and multiple interpretations of decision rules". In: *Statist. Surv.* 4 (2010), pp. 1–39.

[42] Henry Franks. "Supporting cooperation and coordination in open multi-agent systems". PhD thesis. University of Warwick, Mar. 2013.

[43] Henry Franks, Nathan Griffiths, and Sarabjot Singh Anand. "Learning Agent Influence in MAS with Complex Social Networks". In: *Autonomous Agents and Multi-Agent Systems* 28.5 (Sept. 2014), pp. 836–866.

[44] Henry Franks, Nathan Griffiths, and Arshad Jhumka. "Manipulating convention emergence using influencer agents". In: *Autonomous Agents and Multi-Agent Systems* 26.3 (May 2013), pp. 315–353.

[45] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein. "Cooperation Without Communication". In: *Proceedings of the Fifth AAAI National Conference on Artificial Intelligence*. AAAI'86. Philadelphia, Pennsylvania: AAAI Press, 1986, pp. 51–57.

[46] Minas Gjoka et al. "Practical recommendations on crawling online social networks". In: *IEEE Journal on Selected Areas in Communications* 29.9 (2011), pp. 1872–1892.

[47] Marta C. González, Pedro G. Lind, and Hans J. Herrmann. "Networks based on collisions among mobile agents". In: *Physica D: Nonlinear Phenomena* 224.1–2 (2006). Dynamics on Complex Networks and Applications, pp. 137–148.

[48] Marta C. González, Pedro G. Lind, and Hans J. Herrmann. "System of Mobile Agents to Model Social Networks". In: *Phys. Rev. Lett.* 96 (8 Mar. 2006), p. 088702.

[49] Sanjeev Goyal and Maarten C. W. Janssen. "Non-Exclusive Conventions and Social Coordination". In: *Journal of Economic Theory* 77.1 (Nov. 1997), pp. 34–57.

[50] Nathan Griffiths and Sarabjot Singh Anand. "The Impact of Social Placement of Non-learning Agents on Convention Emergence". In: *Proc. of the 11th International Conference on Autonomous Agents and Multi-agent Systems*. Valencia, Spain, 2012, pp. 1367–1368.

[51] Mohammad R. Hasan, Sherief Abdallah, and Anita Raja. "Topology Aware Convention Emergence". In: *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*. AAMAS '14. Paris, France: International Foundation for Autonomous Agents and Multiagent Systems, 2014, pp. 1593–1594.

[52] Mohammad Rashedul Hasan. "Topology-aware approach for the emergence of social norms in multiagent systems". PhD thesis. The University of North Carolina at Charlotte, 2014.

[53] Jerry L. Hintze and Ray D. Nelson. "Violin Plots: A Box Plot-Density Trace Synergism". In: *The American Statistician* 52.2 (1998), pp. 181–184.

[54] Michael N. Huhns and Larry M. Stephens. "Multiagent Systems and Societies of Agents". In: *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. Ed. by Gerhard Weiss. Cambridge, MA, USA: MIT Press, 1999, pp. 79–120.

[55] Genki Ichinose, Yuto Tenguishi, and Toshihiro Tanizawa. "Robustness of cooperation on scale-free networks under continuous topological change". In: *Physical Review E* 88 (Nov. 2013), 052808. 6 pages.

[56] Paul Jaccard. "The Distribution of the Flora in the Alpine Zone". In: *New Phytologist* 11.2 (1912), pp. 37–50.

[57] Nick R. Jennings. "Commitments and conventions: The foundation of coordination in multi-agent systems". In: *The Knowledge Engineering Review* 8.3 (1993), pp. 223–250.

[58] Stephen R. G. Jones. *The economics of conformism / Stephen R.G. Jones*. English. B. Blackwell Oxford [Oxfordshire] ; New York, NY, 1984, x, 123 p. :

[59] Michihiro Kandori. "Social Norms and Community Enforcement". In: *The Review of Economic Studies* 59.1 (1992), pp. 63–80.

[60] U Kang et al. "Centralities in Large Networks: Algorithms and Observations". In: *Proceedings of the 2011 SIAM International Conference on Data Mining*. 2011, pp. 119–130.

[61] George Karypis and Vipin Kumar. "Multilevel k-way Partitioning Scheme for Irregular Graphs". In: *J. Parallel Distrib. Comput.* 48.1 (Jan. 1998), pp. 96–129.

[62] David Kempe, Jon Kleinberg, and Amit Kumar. "Connectivity and Inference Problems for Temporal Networks". In: *Journal of Computer and System Sciences* 64.4 (2002), pp. 820–842.

[63] David Kempe, Jon Kleinberg, and Éva Tardos. "Maximizing the Spread of Influence Through a Social Network". In: *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '03. Washington, D.C.: ACM, 2003, pp. 137–146.

[64] Hyoungshick Kim and Ross Anderson. "Temporal node centrality in complex networks". In: *Phys. Rev. E* 85 (2 Feb. 2012), p. 026107.

[65] James Kittock. "Emergent Conventions and the Structure of Multi–Agent Systems". In: *Lectures in Complex Systems: Proc. of the 1993 Complex Systems Summer School*. 1995, pp. 507–521.

[66] James E. Kittock. "The Impact of Locality and Authority on Emergent Conventions: Initial Observations". In: *Proceedings of the Twelfth National Conference on Artificial Intelligence (Vol. 1)*. AAAI '94. Seattle, Washington, USA: American Association for Artificial Intelligence, 1994, pp. 420–425.

[67] Jon Kleinberg. "The Small-world Phenomenon: An Algorithmic Perspective". In: *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing*. STOC '00. Portland, Oregon, USA: ACM, 2000, pp. 163–170.

[68] Jon M. Kleinberg. "Authoritative Sources in a Hyperlinked Environment". In: *J. ACM* 46.5 (Sept. 1999), pp. 604–632.

[69] Jon M. Kleinberg. "Navigation in a small world". In: *Nature* 406.6798 (2000), pp. 845–845.

[70] V. Kostakos. "Temporal graphs". In: *Physica A Statistical Mechanics and its Applications* 388 (Mar. 2009), pp. 1007–1023.

[71] Krzysztof Kułakowski and Przemysław Gawroński. "To cooperate or to defect? Altruism and reputation". In: *Physica A: Statistical Mechanics and its Applications* 388.17 (2009), pp. 3581–3584.

[72] Kiran Lakkaraju and Les Gasser. "Norm emergence in complex ambiguous situations". In: *Proceedings of the AAAI workshop on coordination, organizations, institutions and norms AAAI, Chicago*. 2008.

[73] Glenn Lawyer. "Understanding the influence of all nodes in a network". In: *Scientific reports* 5 (2015), p. 8665.

[74] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. "Graph evolution: Densification and shrinking diameters". In: *ACM Transactions on Knowledge Discovery from Data (TKDD)* 1.1 (2007), p. 2.

[75] Jure Leskovec and Andrej Krevl. *SNAP Datasets: Stanford Large Network Dataset Collection*. `http://snap.stanford.edu/data`. June 2014.

[76] Jure Leskovec et al. "Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters". In: *Internet Mathematics* 6.1 (2009), pp. 29–123.

[77] Jure Leskovec et al. "Statistical Properties of Community Structure in Large Social and Information Networks". In: *Proceedings of the 17th International Conference on World Wide Web*. WWW '08. Beijing, China: ACM, 2008, pp. 695–704.

[78] David Lewis. *Convention: A Philosophical Study.* Harvard University Press, 1969.

[79] Ted G. Lewis. "Networks". In: *Critical Infrastructure Protection in Homeland Security.* John Wiley & Sons, Inc., 2006, pp. 71–106.

[80] Zhang Lian-Ming et al. "Degree and connectivity of the Internet's scale-free topology". In: *Chinese Physics B* 20.4 (2011), p. 048902.

[81] Samhar Mahmoud. "Metanorms, Topologies, and Adaptive Punishment in Norm Emergence". PhD thesis. King's College London, 2013.

[82] Samhar Mahmoud et al. "Establishing norms with metanorms over interaction topologies". In: *Autonomous Agents and Multi-Agent Systems* (Apr. 2017).

[83] Arun S. Maiya and Tanya Y. Berger-Wolf. "Online Sampling of High Centrality Individuals in Social Networks". In: *Advances in Knowledge Discovery and Data Mining: 14th Pacific-Asia Conference, PAKDD 2010, Hyderabad, India, June 21-24, 2010. Proceedings. Part I.* Ed. by Mohammed J. Zaki et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 91–98.

[84] H. B. Mann and D. R. Whitney. "On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other". In: *Ann. Math. Statist.* 18.1 (Mar. 1947), pp. 50–60.

[85] James Marchant and Nathan Griffiths. "Convention emergence in partially observable topologies". In: *Proceedings of the International Workshop on Coordination, Organizations, Institutions and Norms in Agent Systems (COIN@AAMAS 2017)*. 2017.

[86] James Marchant and Nathan Griffiths. "Limited observations and local information in convention emergence". In: *Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems. 2017, pp. 1628–1630.

[87] James Marchant and Nathan Griffiths. "Manipulating Conventions in a Particle-Based Topology". In: *Coordination, Organizations, Institutions, and Norms in Agent Systems XI*. Ed. by Virginia Dignum et al. Vol. 9628. LNCS XI. 2016, pp. 242–261.

[88] James Marchant, Nathan Griffiths, and Matthew Leeke. "Convention Emergence and Influence in Dynamic Topologies". In: *Proc. of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. Istanbul, Turkey, 2015, pp. 1785–1786.

[89] James Marchant, Nathan Griffiths, and Matthew Leeke. "Destabilising Conventions: Characterising the Cost". In: *Proc. of the 8th IEEE International Conference on Self-Adaptive and Self-Organizing Systems*. 2014.

[90] James Marchant et al. "Destabilising Conventions Using Temporary Interventions". In: *Coordination, Organizations, Institutions, and Norms in Agent Systems X*. Vol. 9372. LNCS. 2015, pp. 148–163.

[91] Richard H McAdams. "Beyond the Prisoners' Dilemma: Coordination, Game Theory, and Law". In: *S. Cal. L. Rev.* 82 (2008), p. 209.

[92] Julian J McAuley and Jure Leskovec. "Learning to Discover Social Circles in Ego Networks." In: *Neural Information Processing Systems*. 2012, pp. 548–56.

[93] Paul Meara. "Emergent Properties of Multilingual Lexicons". In: *Applied Linguistics* 27.4 (2006), pp. 620–644.

[94] Othon Michail. "An Introduction to Temporal Graphs: An Algorithmic Perspective". In: *CoRR* abs/1503.00278 (2015).

[95] Shodai Mihara, Sho Tsugawa, and Hiroyuki Ohsaki. "Influence Maximization Problem for Unknown Social Networks". In: *Proc. of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining.* Paris, France, 2015, pp. 1539–1546.

[96] Mihail Mihaylov, Karl Tuyls, and Ann Nowé. "A decentralized approach for convention emergence in multi-agent systems". English. In: *Autonomous Agents and Multi-Agent Systems* 28.5 (2014), pp. 749–778.

[97] Alan Mislove et al. "Measurement and analysis of online social networks". In: *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement.* ACM. 2007, pp. 29–42.

[98] P. Mukherjee, S. Sen, and S. Airiau. "Emergence of Norms with Biased Interactions in Heterogeneous Agent Societies". In: *2007 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology - Workshops.* Nov. 2007, pp. 512–515.

[99] Declan Mungovan, Enda Howley, and Jim Duggan. "Norm Convergence in Populations of Dynamically Interacting Agents". English. In: *Artificial Intelligence and Cognitive Science.* Ed. by Lorcan Coyle and Jill Freyne. Vol. 6206. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2010, pp. 219–230.

[100] R.B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, 1991.

[101] John F. Nash. "Equilibrium points in n-person games". In: *Proceedings of the National Academy of Sciences* 36.1 (1950), pp. 48–49.

[102] M. E. J. Newman. "The Structure and Function of Complex Networks". In: *SIAM Review* 45.2 (2003), pp. 167–256.

[103] Vincenzo Nicosia et al. "Graph Metrics for Temporal Networks". In: *Temporal Networks*. Ed. by Petter Holme and Jari Saramäki. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 15–40.

[104] Martin A Nowak. "Five rules for the evolution of cooperation". In: *science* 314.5805 (2006), pp. 1560–1563.

[105] E. Nudelman et al. "Run the GAMUT: a comprehensive approach to evaluating game-theoretic algorithms". In: *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004.* July 2004, pp. 880–887.

[106] Colm O'Riordan and Humphrey Sorensen. "Stable Cooperation in the N-Player Prisoner's Dilemma: The Importance of Community Structure". In: *Adaptive Agents and Multi-Agent Systems III. Adaptation and Multi-Agent Learning: 5th, 6th, and 7th European Symposium, ALAMAS 2005-2007 on Adaptive and Learning Agents and Multi-Agent Systems, Revised Selected Papers*. Ed. by Karl Tuyls et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 157–168.

[107] Lawrence Page et al. *The PageRank citation ranking: Bringing order to the web*. Tech. rep. Stanford InfoLab, 1999.

[108] Raj Kumar Pan and Jari Saramäki. "Path lengths, correlations, and centrality in temporal networks". In: *Phys. Rev. E* 84 (1 July 2011), p. 016105.

[109] Sen Pei, Shaoting Tang, and Zhiming Zheng. "Detecting the influence of spreading in social networks with excitable sensor networks". In: *PloS One* 10.5 (2015), e0124848.

[110] Jürgen Pfeffer and Kathleen M. Carley. "k-Centralities: Local Approximations of Global Measures Based on Shortest Paths". In: *Proceedings of the 21st International Conference on World Wide Web*. WWW '12 Companion. Lyon, France: ACM, 2012, pp. 1043–1050.

[111] Josep M. Pujol et al. “The Role of Clustering on the Emergence of Efficient Social Conventions”. In: *Proceedings of the 19th International Joint Conference on Artificial Intelligence*. IJCAI'05. Edinburgh, Scotland: Morgan Kaufmann Publishers Inc., 2005, pp. 965–970.

[112] Martin Purvis et al. “Mechanisms for Cooperative Behaviour in Agent Institutions”. In: *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. IAT '06. Washington, DC, USA: IEEE Computer Society, 2006, pp. 121–124.

[113] Frederick W. Rankin, John B. Van Huyck, and Raymond C. Battalio. “Strategic Similarity and Emergent Conventions: Evidence from Similar Stag Hunt Games”. In: *Games and Economic Behavior* 32.2 (2000), pp. 315–337.

[114] Norman Salazar, Juan A. Rodriguez-Aguilar, and Josep L. Arcos. “Robust coordination in large convention spaces”. In: *AI Communications* 23.4 (2010), pp. 357–372.

[115] F. C. Santos, J. M. Pacheco, and Tom Lenaerts. “Evolutionary dynamics of social dilemmas in structured heterogeneous populations”. In: *Proceedings of the National Academy of Sciences of the United States of America* 103.9 (2006), pp. 3490–3494.

[116] Bastin Tony Roy Savarimuthu, Rexy Arulanandam, and Maryam Purvis. “Aspects of Active Norm Learning and the Effect of Lying on Norm Emergence in Agent Societies”. In: *Agents in Principle, Agents in Practice*. Ed. by David Kinny et al. Vol. 7047. LNCS. Springer, 2011, pp. 36–50.

[117] Bastin Tony Roy Savarimuthu et al. “Norm Emergence in Agent Societies Formed by Dynamically Changing Networks”. In: *Proceedings of the 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. IAT '07. Washington, DC, USA: IEEE Computer Society, 2007, pp. 464–470.

[118] Bastin Tony Roy Savarimuthu et al. "Role Model Based Mechanism for Norm Emergence in Artificial Agent Societies". English. In: *Coordination, Organizations, Institutions, and Norms in Agent Systems III*. Ed. by Jaime Simão Sichman et al. Vol. 4870. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2008, pp. 203–217.

[119] Thomas C Schelling. *The strategy of conflict*. Harvard university press, 1980.

[120] Andrew Schotter. *The Economic Theory of Social Institutions*. Cambridge University Press, 1981.

[121] Sandip Sen and Stéphane Airiau. "Emergence of Norms Through Social Learning". In: *Proc. of the 20^{th} International Joint Conference on Artifical Intelligence*. Hyderabad, India, 2007, pp. 1507–1512.

[122] Yoav Shoham and Moshe Tennenholtz. "Emergent conventions in multi-agent systems: Initial experimental results and observations". In: *Proc. of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*. 1992, pp. 225–231.

[123] Yoav Shoham and Moshe Tennenholtz. "On Social Laws for Artificial Agent Societies: Off-line Design". In: *Artif. Intell.* 73.1-2 (Feb. 1995), pp. 231–252.

[124] Yoav Shoham and Moshe Tennenholtz. "On the emergence of social conventions: modeling, analysis, and simulations". In: *Artificial Intelligence* 94.1–2 (1997), pp. 139–166.

[125] Yoav Shoham and Moshe Tennenholtz. "On the Synthesis of Useful Social Laws for Artificial Agent Societies". In: *Proceedings of the Tenth National Conference on Artificial Intelligence*. AAAI'92. San Jose, California: AAAI Press, 1992, pp. 276–281.

[126] Mile Šikić et al. "Epidemic centrality–is there an underestimated epidemic impact of network peripheral nodes?" In: *The European Physical Journal B* 86.10 (2013), p. 440.

[127] Horst D. Simon and Shang-Hua Teng. "How Good is Recursive Bisection?" In: *SIAM J. Sci. Comput.* 18.5 (Sept. 1997), pp. 1436–1445.

[128] Duncan Snidal. "Coordination versus Prisoners' Dilemma: Implications for International Cooperation and Regimes". In: *The American Political Science Review* 79.4 (1985), pp. 923–942.

[129] Luc Steels. "The Origins of Ontologies and Communication Conventions in Multi-Agent Systems". In: *Autonomous Agents and Multi-Agent Systems* 1.2 (Oct. 1998), pp. 169–194.

[130] Toshiharu Sugawara. "Emergence and Stability of Social Conventions in Conflict Situations". In: *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One*. IJCAI'11. Barcelona, Catalonia, Spain: AAAI Press, 2011, pp. 371–378.

[131] J. Tang et al. "Small-world behavior in time-varying graphs". In: *Physical Review E* 81.5, 055101 (May 2010), p. 055101.

[132] Nick Tinnemeier, Mehdi Dastani, and John-Jules Meyer. "Programming Norm Change". In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*. AAMAS '10. Toronto, Canada: International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 957–964.

[133] R. Toivonen et al. "Broad lifetime distributions for ordering dynamics in complex networks". In: *Phys. Rev. E* 79 (1 Jan. 2009), p. 016109.

[134] Jeffrey Travers and Stanley Milgram. "An Experimental Study of the Small World Problem". In: *Sociometry* 32.4 (1969), pp. 425–443.

[135] R. Tuomela. *Social norms, tasks, and roles.* Dept. of Philosophy, University of Helsinki, 1992.

[136] Raimo Tuomela. *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford University Press, 1995.

[137] Twitter Developers. *Twitter: API Rate Limits.* Aug. 2017. URL: `https://dev.twitter.com/rest/public/rate-limiting`.

[138] Twitter Developers. *Twitter: REST APIs.* Aug. 2017. URL: `https://dev.twitter.com/rest/reference`.

[139] Edna Ullmann-Margalit. *The emergence of norms.* Vol. 11. Oxford University Press, 1977.

[140] P. Urbano and H. Coelho. "From consensus to consensus: random evolution of collective choices". In: *2005 portuguese conference on artificial intelligence.* Dec. 2005, pp. 274–279.

[141] Paulo Urbano et al. "Force Versus Majority: A Comparison in Convention Emergence Efficiency". In: *Coordination, Organizations, Institutions and Norms in Agent Systems IV : COIN 2008 International Workshops, COIN@AAMAS 2008, Estoril, Portugal, May 12, 2008. COIN@AAAI 2008, Chicago, USA, July 14, 2008. Revised Selected Papers.* Ed. by Jomi Fred Hübner et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 48–63.

[142] Daniel Villatoro, Jordi Sabater-Mir, and Sandip Sen. "Social Instruments for Robust Convention Emergence". In: *Proc. of the 22^{nd} International Joint Conference on Artificial Intelligence.* Barcelona, Catalonia, Spain, 2011, pp. 420–425.

[143] Daniel Villatoro, Sandip Sen, and Jordi Sabater-Mir. "Topology and memory effect on convention emergence". In: *Proc. of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology.* 2009, pp. 233–240.

[144] Daniel Villatoro et al. "Dynamic Sanctioning for Robust and Cost-Efficient Norm Compliance". In: *Twenty-Second International Joint Conference on Artificial Intelligence.* Ed. by Toby Walsh. IJCAI/AAAI. Barcelona: IJCAI/AAAI, 2011, pp. 414–419.

[145] Thuc Vu, Rob Powers, and Yoav Shoham. "Learning Against Multiple Opponents". In: *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*. AAMAS '06. Hakodate, Japan: ACM, 2006, pp. 752–759.

[146] Bart de Vylder. "The evolution of conventions in multi-agent systems". PhD thesis. Artificial Intelligence Lab, Vrije Universiteit Brussel, 2007.

[147] Adam Walker and Michael Wooldridge. "Understanding the Emergence of Conventions in Multi-Agent Systems". In: *International Conference on Multi-Agent Systems*. 1995, pp. 384–389.

[148] Xiaojie Wang et al. "Effective identification of multiple influential spreaders by DegreePunishment". In: *Physica A: Statistical Mechanics and its Applications* 461 (2016), pp. 238–247.

[149] Christopher John Cornish Hellaby Watkins. "Learning from Delayed Rewards". PhD thesis. Cambridge, UK: King's College, May 1989.

[150] Duncan J. Watts and Steven H. Strogatz. "Collective dynamics of 'small-world' networks". In: *nature* 393.6684 (1998), p. 440.

[151] Duncan Whitehead. *The El Farol Bar Problem Revisited: Reinforcement Learning in a Potential Game*. Tech. rep. Edinburgh School of Economics, University of Edinburgh, 2008.

[152] Yang Yang, Yuxiao Dong, and Nitesh V Chawla. "Predicting node degree centrality with the node prominence profile". In: *Scientific reports* 4 (2014), p. 7236.

[153] H Peyton Young. "The Evolution of Conventions". In: *Econometrica* 61.1 (Jan. 1993), pp. 57–84.

[154] Hobart Peyton Young. "The Economics of Convention". English. In: *The Journal of Economic Perspectives* 10.2 (1996), pp. 105–122.

[155] H. Zhuang et al. "Influence Maximization in Dynamic Social Networks". In: *2013 IEEE 13th International Conference on Data Mining*. Dec. 2013, pp. 1313–1318.