Economic essays on marine invasive species and international fisheries agreements

Adam N. Walker

Thesis Committee

Promotor
Prof. Dr E.C. van Ierland
Professor of Environmental Economics and Natural Resources
Wageningen University

Co-promotors
Dr R.A. Groeneveld
Associate Professor, Environmental Economics and Natural Resources Group
Wageningen University

Dr H.-P. Weikard
Associate Professor, Environmental Economics and Natural Resources Group
Wageningen University

Other members
Prof. Dr M. Finus, University of Bath, England
Dr K.G. Hamon, LEI, The Hague
Prof. Dr A.G.J.M Oude Lansink, Wageningen University
Prof. Dr K. Eisenack, Carl von Ossietzky Universität Oldenburg, Germany

This research was conducted under the auspices of the WASS graduate school

Economic essays on marine invasive species and international fisheries agreements

Adam N. Walker

Thesis
submitted in the fulfilment of the requirements for the degree of doctor
at Wageningen University
by the authority of the Rector Magnificus
Prof. Dr A.P.J. Mol,
in the presence of the
Thesis Committee appointed by the Academic Board
to be defended in public
on Monday 15 February 2016
at 13.30 in the Aula.

Adam N. Walker
Economic essays on marine invasive species and international fisheries agreements,
170 pages.

PhD thesis, Wageningen University, Wageningen, NL (2016)
With references, with summary in English

ISBN 978-94-6257-630-8

Contents

Acknowledgements v

Introduction 1

1.1 Topics and research questions 1

1.1.1 Part A: Invasive species 1

1.1.2 Part B: International fisheries agreements 5

1.2 Methods 7

1.3 Novel contributions in each chapter 10

Optimal ballast water management standards: implications of Allee effects and minimum viable populations 15

2.1 Introduction 16

2.2 The model 20

2.2.1 The PD-R and technology functions 20

2.2.2 The hazard rate function 21

2.2.3 The expected damage function 22

2.2.4 The loss function 23

2.3 Analysis of optimal standards 24

2.4 Conclusions 29

Appendix 2.1 Derivatives of the hazard function 31

Appendix 2.2 Derivatives of the damage function 33

Invasive species control in a one-dimensional metapopulation network 37

3.1 Introduction 38

3.2 The model 41

3.3 Results from a two-patch model 46

3.3.1 No containment 46

3.3.2 Introducing the possibility for containment 50
3.4 Extension to three patches 56
3.5. Discussion and conclusion 57
Appendix 3.1 Parameterisations for Section 3.3 61
Appendix 3.2 Parameterisations for Section 3.4 61
Farsightedness, changing stock location and the stability of international fisheries agreements 65
4.1 Introduction 66
4.2 The bioeconomic model 68
4.3 The fisheries game 69
4.3.1 Choosing effort levels 70
4.3.2 Stability, solution concepts and sharing rules 72
4.3.3 Some established results 76
4.4 Analysing FDS coalitions with *n* symmetric players 77
4.5 Analytics of asymmetry and changing stock location 79
4.6 Numerical and sensitivity analyses 81
4.6.1 Changes in stock location with three players 82
4.6.2 Changes in stock location with four players 88
4.7 Discussion and conclusion 90
Appendix 4.1 Equilibrium strategy for the All Singletons structure. 91
Appendix 4.2a Derivation of Inequality (4.12) 92
Appendix 4.2b Comparison of (4.13) and (4.15) 93
Appendix 4.3 The independence of coalition stability from *r* 93
Appendix 4.4 The relative advantage of cost versus catchability asymmetries 94
The rise and fall of the great fish pact under endogenous risk of stock collapse 97
5.1 Introduction 98
5.2 Bio-economic model 100
5.3 Grand Coalition stability 103

5.4 Results of the standard two-stage game ... 104
5.5 Including transition payoffs ... 113
5.5.1 Method for including transition payoffs ... 113
5.5.2 Results of including transition payoffs ... 114
5.6 Conclusions ... 115
Appendix 5.1 Utility function and numerical accuracy ... 117
Appendix 5.2 Deriving the envelope condition ... 119
Appendix 5.3 The properties of harvest functions. ... 120
Synthesis ... 123
6.1 Summary and answers to research questions ... 123
6.1.1 Ballast Water Management ... 123
6.1.2 Invasive species control ... 124
6.1.3 Changing fish stock location ... 126
6.1.4 Endogenous risk of fish stock collapse ... 128
6.2. Evaluation ... 129
6.2.1 Evaluation of Part A ... 129
6.2.2 Evaluation of Part B ... 131
6.3 Reflections and policy relevance ... 134
6.3.1 Discussion of Part A ... 134
6.3.2 Discussion of Part B ... 137
6.4 Concluding remarks ... 140
References ... 143
Summary ... 161

Acknowledgements

I extend my sincerest thanks to everyone who has been involved in this PhD, both inside and outside of the university. Thanks go first and foremost to Ekko van Ierland, Rolf Groeneveld and Hans-Peter Weikard who formed my supervisory team. Many thanks to you for the discussions, many revisions to articles, and in no small measure, your patience. Co-authors also made substantial contributions to this thesis. Silke Gabbert's support, interest and steadfastness with Chapter 2 has been deeply appreciated. In addition to his role in preparing the manuscript for Chapter 3, Jan-Jaap Poos's programming skills saved me several months by drastically reducing computation time. Chapter 5 benefited from Andries Richter's pursuit of awesomeness in scientific exposition.

Thanks go also to my other colleagues in the Environmental Economics and Natural Resources (ENR) Group. The mutual support of a good group of fellow PhD students sure can make a difference in moral. Additionally, having some more senior people to turn to also helps a lot. Edwin van der Werf's council proved particularly helpful on several occasions, and I thank him for taking the time to talk to me. Also within ENR, thanks go to Wil and Gré for all their administrative support. Within the Department of Economics, I received many useful comments on my work, particularly in seminars. Again, I extend my thanks. Outside of the university, thanks go to the many members of the VECTORS project, who took an interest in my work and provided much of the inspiration for the contents of this thesis.

Friends and family play a huge role in life regardless of whether one is doing a PhD. Gratitude to them should be shown every day. So, on the day you read this, THANKS!

1 Introduction

1.1 Topics and research questions

There are many changes in marine ecology taking place in European seas. This thesis contributes to the need to understand the economic implications of these changes and to formulate management strategies. To do so, the thesis is divided into two parts which consider marine ecological change within two distinct sub-topics. Firstly, in Part A, this thesis focuses on economic analysis of marine invasive species. Invasive species are taxa that have been introduced outside of their native range (IUCN, 2000). These introduced taxa may then be detrimental to economic output and lead to biodiversity loss and reduced ecological services (Frésard and Boncoeur, 2006). Secondly, Part B concerns International Fisheries Agreements (IFAs). IFAs aim to ensure profitable and sustainable fishing when multiple countries have an interest in a given fish stock (Hilborn et al., 2005). Part B focuses on the impacts of changing fish stock location and possible fish stock collapse on the success of IFAs. The topics in Parts A and B represent important contemporary concerns in Europe, as reflected by the focus on these topics within the EU Seventh Framework Program project named VECTORS (www.marine-vectors.eu), to which the research in this thesis has contributed.

1.1.1 Part A: Invasive species

European seas are reported to contain 879 multicellular invasive species (Galil et al., 2014) and 176 marine invasive species are known to have an economic impact (Vilà et al., 2010). Human health, ecosystem services and biodiversity can all suffer due to invasions (Frésard and Boncoeur, 2006; Pimentel et al., 2005; Scalera, 2010). Managers concerned with these impacts have two problems which must be considered in tandem. These are the prevention problem and the control problem

(Olson and Roy, 2005; Burnett et al., 2006; Kim et al., 2006; Finnoff et al., 2007; Finnoff et al., 2010b; Sanchirico et al. 2010; Burnett et al. 2012). The prevention problem refers to determining appropriate management actions to attempt to prevent the establishment of invasive species in a given area. The control problem refers to determining appropriate management actions should an invasion become established in a given area. Chapter 2 in this thesis is concerned with the prevention problem and Chapter 3 with the control problem. Note that both of these problems have a spatial aspect. The spread of an invasive species from its native range to a non-native range is a spatial process. Further, the spread of an invasive species within its new non-native range can also be considered as a spatial process.

Invasive species prevention

Chapter 2 addresses the prevention problem by considering Ballast Water Management (BWM). Ballast water is a major vector for the spread of invasive species (Ruiz et al., 1997). Ballast water is pumped into ships in order to increase their weight and balance them when the ship is unloaded. The ship then travels to another area and pumps out the unneeded ballast water upon loading cargo. The ballast water is thus a vector which transports species globally. The BWM convention has been set up by the International Maritime Organization (IMO, 2004) to manage this vector. The stated aim of the BWM convention is to eliminate the risks from IAS transported by ballast water (Gollasch et al., 2007), to be achieved by a combination of measures including ballast water treatment. Ballast water treatment is the process of removing organisms from ballast water, and thus reducing the probability that an invasive species will become established. This is normally achieved with a combination of technologies, such as filtration, biocides and UV light (Dobbs and Rogerson, 2005). These technologies are used in ballast water treatment systems which are units installed into ships. Ballast water treatment systems are designed to operate to a specific standard, which is set in the BWM convention. The standard is defined in terms of the concentration of organisms in treated ballast water. Lower (stricter) concentration standards result in lower probabilities that invasions will become established.

Purchasing an individual ballast water treatment system can cost between $640,000 and $950,000 (King et al., 2012), but the damage prevented could also be substantial. The number of potential invasive species is large (Molnar et al., 2008; Galil et al., 2014) and the arrival of a single invasive species can lead to large economic damages. For example, the arrival of the comb jelly in the Black Sea led to a US$16.7 million reduction in the present value of anchovy fishery rents (Knowler and Barbier, 2000). Ballast water treatment can also avoid the cost of controlling established invasions. For example, the cost of controlling the invasive slipper-limpet in order to protect the scallop fisheries in the Bay of St-Brieuc, France, is estimated at circa €1 million per year (Frésard and Boncoeur, 2006). BWM is thus an economic problem of choosing a ballast water treatment standard which minimises the sum of treatment costs and expected damages. The ballast water treatment standard which minimizes this sum is the optimal standard. There are many features of BWM which need to be accounted for to determine the optimal standard. The theoretical analysis in Chapter 2 considers two specific features of BWM, namely Minimum Viable Populations (MVPs) of the invasive species, and Allee effects. This leads to Research Question 1:

Research Question 1: *What are the implications of minimum viable populations and Allee effects for optimal ballast water management standards?*

An MVP is the smallest number of individual organisms which is sufficient to ensure that the population can sustain itself upon arrival (i.e. become established). It is inefficient to reduce the population in ballast water to more than marginally below the MVP, because the probability of invasion establishment is zero for all population sizes less than marginally below the MVP. An Allee effect occurs when smaller invasive populations face disproportionately greater probabilities of extinction due to poor resilience to fluctuations in birth and death rates or environmental shocks (Williamson, 1989). In the context of the conservation of endangered species, greater probabilities of extinction are considered to be a bad thing, but in the context of invasive species, any effects which increase the probability of extinction can be exploited in order to mitigate against invasions in the most effective manner.

Invasive species control

The BWM convention is not yet in force (Scriven et al., 2015) and there are many other vectors via which invasive species can be transported. Therefore, Chapter 3 addresses the optimal management of established invasions, i.e. the control problem. The control problem concerns the trade-off between the damages resulting from the invasion and the costs of controlling the invasion. Costs are incurred by the efforts to limit the size of the invasive species population or its spread within a non-native range. Determining whether such interventions are worth the cost and exactly the best way to implement these interventions is a complex management problem which has many different aspects. In particular, a better understanding of the spatial aspects of invasion control is needed (Albers et al., 2010; Epanchin-Niell and Hastings, 2010; Savage and Renton, 2014). Chapter 3 aims to contribute towards the understanding of the spatial aspects of the control problem in determining optimal management.

Space can be treated as either continuous or discrete. Continuous space is considered most recently by Finnoff et al. (2010a) and Carrasco et al. (2010a). Discrete space, which we adopt in Chapter 3, divides the non-native range into "patches", as is done most recently by Carrasco et al. (2012), McDermott et al. (2013), Fenichel et al. (2014) and Kovacs et al. (2014). Each patch may contain an invasive species population. When considering all patches together, we are considering a population of populations, or a "metapopulation". The metapopulation approach has the advantage that it allows for certain patches of the non-native range to be invaded or uninvaded, but also, crucially for the novelty of Chapter 3, that it allows patches to contain varying populations of the invasive species. This modelling approach is relevant for several pertinent real-world cases, which are detailed in Chapter 3. We therefore construct a model which allows for varying invasive population sizes in patches and for removal of any amount of the population from any patch (cf. Burnett et al., 2007), as opposed to treating patches as either invaded or non-invaded (as in and Epanchin-Niell and Wilen, 2012). This leads to Research Question 2:

Research Question 2: *What are the implications for optimal spatial control of invasive species when the invasive population is modelled with varying invasive population sizes within patches?*

Modelling varying invasive population sizes within patches increases the level of modelling detail. We can therefore expect the model to produce more detailed optimal management interventions. For example, when a patch can only be invaded or non-invaded, the management choice for that patch is restricted to doing nothing or destroying the invasive population in that patch. When the invasive population size within a patch can vary, management can remove a proportion of the invasion from that patch. We can also expect that there is greater scope for timing of interventions. For example, when the invasive population size can vary within a patch, the management decisions depend not only upon whether the invasion has arrived in that patch, but also the length of time that the invasion has been present in that patch, i.e. how much time the invasive population within that patch has had to grow.

1.1.2 Part B: International fisheries agreements

Part B focuses on the effects of ecological changes on cooperation in IFAs. It has long been recognized that international fisheries management is a problem of cooperation (Crutchfield, 1964). Cooperation leads to larger fish stocks and greater total profits (Clark, 2010). Cooperation is maximized when all fishing nations decide on how much to fish, not to maximize their own profit from the fish stock, but to maximize the sum of the profit of all nations who wish to fish the stock. Such a state of affairs is referred to as a Grand Coalition. If a nation decides to fish to maximize its private profit from the fish stock, with no concern for how this might affect the profits of other nations, then that nation is said to be free-riding. Not only does free-riding reduce the profit of other nations, it also reduces the size of the fish stock, which has implications for the sustainability of the fish stock. Determinants of successful IFAs are therefore of scientific interest in general (Hilborn, 2007) as is the potential for IFAs to be affected by ecological changes or the possibility of ecological changes (Munro, 2008). Of particular concern are the impacts of climate change on the location of fish stocks (Cheung et al., 2009) and the risk of stock collapse caused by overfishing (Mullon et al., 2005). Therefore, Chapter 4 considers changing stock location and Chapter 5 considers a risk of stock collapse.

Fisheries management under changing stock location

Chapter 4 analyses the effects of changing stock location on cooperation in IFAs. Changes in stock location are likely due to climate change (Cheung et al., 2009). For example, mackerel stocks in the North East Atlantic have recently shifted northwards (Jansen and Gislason, 2011). This has led to unilateral setting of national fishing quota by Iceland which violated the existing IFA agreement and resulted in increased exploitation of the fish stock (Haraldsson and Carey, 2011; Arnason, 2012). This shows that changes in fish stock location may have a destabilising effect on IFAs. It is therefore important to understand the conditions under which coalitions can maintain their stability in the face of changing stock location. This leads to

Research Question 3:

Research Question 3: *To what extent can farsightedness stabilize IFAs in the face of changing fish stock location?*

As we have established, cooperation in IFAs is a strategic problem and therefore we employ coalition theory. Coalition theory requires assumptions about how other players will respond to the cooperative choices of others. These assumptions imply particular behaviours on behalf of the players. We are therefore interested in analysing plausible behaviours in terms of their implications for whether the game is played in a socially optimal, i.e. cooperative, way. Specifically, we are interested in behaviours which determine how players respond to a choice to not cooperate (i.e. to deviate) by other another player. One such plausible behaviour is embodied by the Nash conjecture. Under the Nash conjecture, players do not change their cooperative choice in response to such a change by another player. In this respect, players are "shortsighted". The shortsightedness of the Nash conjecture has been criticized by Harsanyi (1974). An alternative, and more plausible behavioural assumption, is that players may respond to a deviation by another player by changing their choice of whether to cooperate. This is the farsightedness concept of Chwe (1994). Players who adopt farsighted conjectures may respond to a change in the cooperative choice of another player by changing their own cooperative choice. Different behavioural assumptions imply that cooperative agreements will be affected differently by changes in fish stock location. We therefore address Research Question 3 by analysing whether cooperation in the face of changing stock location can be better maintained under farsighted conjectures rather than under Nash (or shortsighted) conjectures.

IFAs under a risk of fish stock collapse

Chapter 5 analyses the effect of the possibility for fish stock collapse on the stability of Grand Coalitions. Globally, from 1950 to 2000, 366 fisheries collapsed and the collapses are generally attributed to over-fishing (Mullon et al., 2005). The risk of fish stock collapse is likely to be important in determining the strategic harvest choices of fishing nations (Hannesson, 2014). Avoiding fish stock collapse is an important motivation for the formation of IFAs (Hilborn et al., 2005). A risk of fish stock collapse can be considered as exogenous or endogenous. Exogenous risk exists when the actions of agents do not affect the probability that an event occurs. A suitable example is a tsunami because there are no actions available which can reduce the probability of a tsunami occurring. Endogenous risk exists when the actions of agents affect the probability of the event occurring. This is particularly relevant for fisheries agreements because lower fish stocks are at greater risk from collapse (Mullon et al., 2005). This leads to Research Question 4:

Research Question 4: *What are the implications of an endogenously determined risk of fish stock collapse on cooperation in IFAs?*

The intensity of renewable resource exploitation can either increase or decrease due to endogenous risk (Ren and Polasky, 2014; Sakamoto, 2014). Of particular interest in Chapter 5 are the implications for cooperation of such changes in exploitation. It is important to consider these implications because theoretical modelling of IFAs has yet to successfully replicate empirical observations regarding the number of players for whom cooperation can be successfully maintained (Breton and Keoula, 2014). This is highlighted by Hannesson (2014) who demonstrates that it is not possible to reconcile theoretical predictions regarding resource exploitation in the North East Atlantic mackerel fishery with empirical observations. Existing theoretical insights suggest that cooperation can only be sustained for small numbers of players, whereas in reality, much larger coalitions can be observed. This divergence between theory and empirical observations is referred to as the “puzzle of small coalitions” (Breton and Keoula, 2014). We are therefore particularly interested in the effects of endogenous risk in terms of the extent to which they provide insight into the puzzle of small coalitions.

1.2 Methods

This thesis employs diverse methods to address the research questions. The approach is always theoretical, but employs both analytical and numerical methods. Chapter 2 adopts a purely analytical modelling approach. The model consists of a “loss function” which includes the cost of BWM and the expected damage under different standards. The properties of the loss function are analysed to gain insight into the determinants of optimal standards. Our modelling approach respects the irreversibility in invasion establishment. An irreversible invasion is one for which it is impossible to reduce the size of the IAS population to zero. This is the most general assumption as far as marine invasive species are concerned (Vitousek et al., 1997; Parker et al., 1999). Accordingly, this study relates in terms of methodology to the work on optimal resource management in the face of irreversible events stemming from Tsur and Zemel (1998). More directly, this study builds on Knowler and Barbier (2005), Kim et al. (2006) and Burnett et al. (2012), who treat invasions as irreversible.

In Chapter 3, optimal spatial control policies for an invasive species are determined using Stochastic Dynamic Programming (SDP). SDP is frequently applied to study the optimal management of invasive species and can be used to analyse various sources of stochasticity. For example, the growth of the invasive population can be treated as stochastic (Olson and Roy, 2002), as can the detection of new invasions (Mehta et al., 2007) or the damages of the invasion (Sims and Finnoff, 2013). Chapter 3 uses SDP to model the stochastic spread of an invasion through a landscape of discrete patches (the analysis focuses on a two-patch landscape). This approach is closely related to that of Carrasco et al. (2010b), who employ a stochastic spread process in continuous space. The method is also closely related to studies considering deterministic spread between discretized patches (Blackwood et al., 2010; Epanchin-Niell and Wilen, 2012).

Specifically, we model the case where the invasion can grow within a patch and spread to neighbouring patches, in which it can also grow. The spread between the areas is stochastic and the probability of spread into a non-invaded patch increases in the size of the invasive population in adjacent invaded patches. Patches are arranged in a one-dimensional network, which consists of a series of patches connected in a line. One-dimensional networks have been used to analyse invasive species control by Chadès et al. (2011) and are relevant for invasive species spreading along coastlines such as the pacific oyster in the Wadden Sea (Troost, 2010). The model allows for the invasion to be controlled by reducing the size of the invasion in a given patch or by implementing a barrier which reduces the probability of spread between patches without affecting the invasive populations within invaded patches. Optimal application of these interventions is determined using value function iteration and the optimal interventions are presented graphically. Optimal interventions are presented for specific parameterizations to demonstrate the implications of allowing for varying invasive populations within patches.

Chapters 4 and 5 employ game theory to analyse IFAs. Game theory is the study of multi-agent decision problems (Gibbons, 1992). Fisheries can be conceptualized as a common-pool resource whereby multiple nations (players) can exploit the stock (Munro, 1979). As such, nations must decide how intensively to exploit the stock and in making this decision, take the actions of other nations into account. Game theory is therefore an appropriate approach. Chapters 4 and 5 employ game theory and coalition theory to analyse cooperation, following in the tradition of Mesterson-Gibbons (1992). Coalition theory allows nations to cooperate, with the aim of maximizing the sum of the benefits from exploiting the resource for all nations in a coalition. Thus, coalition theory allows a group of nations to behave as one single player. If all nations with an interest in the fish stock are members of an agreement, then a "Grand Coalition" exists. In fisheries games, if a Grand Coalition exists then optimal management of the stock ensues (Clark, 2010). Thus coalition theory is a useful lens enabling us to analyse the common-pool resource problem in terms of IFAs.

Chapter 4 analyses IFAs under changing fish stock location using the classic Gordon-Schaefer model and a farsighted stability concept. We analytically determine the implications of farsightedness in the simplest case where stock location is constant and players are symmetric. We then proceed to analyse changing stock location. In the Gordon-Schaefer model, the fish stock is conceptualized as existing at a single point in space. We define the location of this point in relation to the locations of the fishing nations, which are also conceptualized as points in space. If fish stock location changes then the point giving the location of the fish stock changes such that some nations will be closer to the fish stock and others become further away. We analyse the effects of changing stock location on the stability of Grand Coalitions using numerical sensitivity analysis. Accordingly, we analyse the extent to which farsightedness can ensure cooperation, relative to shortsightedness, in the face of changing stock location.

In Chapter 5, we adopt the Levhari and Mirman (LM) model (Levhari and Mirman, 1980) to analyse the effect of an endogenous risk of stock collapse on Grand Coalition stability. We adopt this model because it employs a more suitable conceptualization of the externality for our purposes than the Gordon-Schaefer model. In the Gordon-Schaefer model, the externality of fishing is that if one nation fishes more, then it becomes harder for other nations to catch a given amount of fish. In the LM model, the externality is dynamic, in that if one nation fishes more now, then there will be fewer fish in the future[1]. The conceptualization of the externality of fishing in the LM model is more suitable if we wish to extend the model to include an endogenous risk of fish stock collapse. In that case, the externality is two-fold. If one nation fishes more now, then there will be less stock in the future and also a greater risk that the stock will collapse such that fishing in the future will not be possible.

We extend the LM model by including a risk of stock collapse which increases at an increasing rate as fish stock size reduces (i.e. the function is convex). This approach is similar to that of Ren and Polasky (2014), although they only assume that risk increases as stock size reduces, not that the function is convex. Our approach is most similar to Nikuiya et al. (2014), who assume that the risk is convex in fish stock size. Under these conditions, optimal harvest is non-linear in stock and therefore, the model becomes analytically intractable (Antoniadou et al., 2013). Therefore, we adopt a numerical approach. Our model calculates steady state fish stock sizes across a parameter space of different growth and discount rates. Payoffs are then calculated and used to determine the implications of endogenous risk on the stability of Grand Coalitions.

1.3 Novel contributions in each chapter

Chapter 2 of this thesis provides the first general theoretical insights into the economics of BWM. BWM has been studied only by Fernandez (2006), who adopts a game theoretic perspective to study what standards individual countries would set while taking the standards of other countries into account. However, this is not a relevant specification given the current institutional context, in which the International Maritime Organization chooses the standard. The strategic consideration is whether a nation should ratify the convention given the ratification choices of other nations. In terms of results, Chapter 2 of this thesis is novel in that it provides the first well-grounded (albeit theoretical) evaluation of BWM standards. Gollasch et al. (2007) highlight that there is some dissatisfaction with the manner in

[1] The externality in the Gordon-Schaefer model can also be modelled dynamically. However, when the Gordon-Schaefer model is used to analyse IFAs, the analyses predominantly focus on the static externality (cf. Pintassilgo et al., 2010). This is partly because of the complexity of analysing the Gordon-Schaefer model dynamically.

which BWM standards were derived, but there has been no formal grounding for this critique other than that more information upon which to base the choice of standards would have been preferable, and that some relevant stakeholders and experts were not involved in the process. In Chapter 2, by formally analysing BWM in an appropriate setting, we generate insights into the drivers of optimal BWM standards and hence provide a formal basis with which current standards can be evaluated, and with which possible new standards can be considered.

Chapter 3 of this thesis makes a predominantly methodological contribution to the invasive species literature. The literature on invasive species control has progressed from deterministic models where the invasion is expressed as a single variable, depicting invasive population size or extent of spread (e.g. Eiswerth and Johnson, 2002) to include various stochastic elements (e.g. Olson and Roy, 2002) and uncertainty about parameters of the model (e.g. D'Evelyn et al., 2008). Recently, more focus has been given to the spatial aspects of invasive species control (e.g. Blackwood et al., 2010). All these different directions have resulted in valuable insights into the invasive species control problem. The insights provided by our modelling approach essentially stem from combining previous modelling approaches. Within a patch, our model is exactly the same as the basic models of a deterministically growing invasive species. The novelty stems from linking such patches via a stochastic spread process. This is an extension of models where patches can be only invaded or uninvaded (e.g. Epanchin-Niell and Wilen, 2012). Chapter 3 demonstrates how this extension facilitates novel and more detailed optimal management policies.

Chapters 4 and 5 make contributions to the recent literature on applying coalition theory to fisheries, most recently by Pintassilgo et al. (2010), Long and Flaaten (2011), Bjørndal and Lindroos (2012), Breton and Keoula (2012), Rettieva (2012) and Breton and Keoula (2014). Chapter 4 of this thesis makes novel contributions through the modification of the farsighted stability concept of Chwe (1994). Stability concepts are rules which are applied to determine the choices of a player regarding whether or not to cooperate. A stability concept embodies behavioural assumptions about how players will respond to the choices of others. In order to answer the research question, modifications are made to the farsighted stability concept such that it can be operationalized in games with asymmetric players and transfer payments. This is required because the original farsightedness stability concept can result in a situation where, ceteris paribus, the cooperative choices of players continuously change. We term this a "cycle". Cycles happen when a cooperative agreement collapses due to progressive deviations and then rebuilds itself again, only to infinitely repeat the process. Our modification to the farsighted stability concept, which we term Farsighted Downwards Stability (FDS) prevents these cycles from occurring. Furthermore, Chapter 4 is novel in the way in which changing stock location is modelled. Changing stock location has been considered in a variety of settings by Brandt and Kronbak (2010), Ekerhovd (2010), Ellefsen (2012) and Ishimura et al. (2012). Chapter 4 provides a more general approach, based on the Gordon-Schaefer model, by using the variable determining harvesting technology (normally given by q) to represent fish stock location. The idea is that having fishing grounds closer to the home port increases the efficiency of fishing effort in much the same way as more efficient fishing technology.

The innovative contribution of Chapter 5 is two-fold. Firstly, Chapter 5 contributes to the literature by analysing endogenous risk in renewable resource games, as in Sakamoto (2014), Ren and Polasky (2014) and Nikuiya et al. (2014), from a coalition theory perspective. Secondly, Chapter 5 contributes to the literature by relaxing a common assumption, namely that payoffs are determined in steady states. A steady state refers to a fish stock size at which total harvest is equal to the growth of the stock such that, once this fish stock size is achieved, it remains constant. Payoffs calculated in these steady states inform the decisions of players regarding whether or not to cooperate. Chapter 5 relaxes this assumption by including payoffs in the transition between steady states. Transitions between steady states result from a change in the cooperative decisions of players. For example, if all players are cooperating, then a deviation by one player will mean that fishing will increase. The new steady state will therefore be lower. Transition payoffs occur in the time that it takes for the fishery to adjust to increased fishing and to settle at its new steady state. The importance of transition payoffs is highlighted by Sakamoto (2014), but the implications of transition payoffs for cooperation have not yet been elucidated. Insights generated regarding the effects of endogenous risk and transition payoffs are useful in furthering the debate on the puzzle of small coalitions.

Part A: Marine invasive species

2 Optimal ballast water management standards: implications of Allee effects and minimum viable populations

ABSTRACT

The stated aim of the Ballast Water Management (BWM) convention is to eliminate the risk from invasive species transported via ballast water by enforcing a ballast water treatment standard. Accordingly, a standard has been set with the aim of reducing the size of invasive populations to below their Minimum Viable Population (MVP) and thus prevent their establishment. This study develops a theoretical model of irreversible invasions to study the determinants of optimal BWM standards. The analysis suggests that a BWM standard which aims to reduce the size of invasive populations to marginally below their MVP can only be optimal if the function determining the hazard rate of invasion establishment is non-continuously differentiable around the MVP. We proceed to analyse the conditions under which the hazard function would be non-continuously differentiable. Non-continuous differentiability of the hazard function may not hold in the presence of an Allee effect. An Allee effect occurs when the probability of successful invasion establishment increases at an increasing rate in the number of individuals of the invasive species emitted in ballast water. We conclude therefore that the presence of an Allee effect fundamentally determines whether the current BWM standard is optimal.

This chapter is based on the paper: Walker, A.N., Groeneveld, R.A., Gabbert, S. Optimal ballast water management standards: implications of Allee Effects and minimum viable populations. Submitted to a peer reviewed journal.

2.1 Introduction

Ballast Water Management (BWM) standards are used to manage the risks from invasive species transported in ballast water. This risk is being addressed by the International Maritime Organisation's BWM Convention for the Control and Management of Ship's Ballast Water and Sediments (IMO, 2004). The stated aim of the BWM convention is to eliminate risks from invasive species transported by ballast water (Gollasch et al., 2007) achieved by treating ballast water such that the concentration of organisms is reduced to a given concentration standard. The concentration standard in the BWM convention aims to reduce invasive population sizes below their Minimum Viable Population (MVP) size, and thus ensure failure of establishment (Gollasch et al., 2007). Such a policy possesses intuitive appeal because it precludes the potentially significant damages from invasive species transported in ballast water. Invasive species affect commercial fish stocks (Leppäkoski et al., 2002) and aquaculture (Switzer et al., 2011) and are a global problem. Indeed, only 16% of marine ecoregions have no reported invasions (Molnar et al., 2008).

However, BWM comes at a cost. Ballast water treatment systems are designed to achieve a given ballast water treatment standard. Purchasing an individual ballast water treatment system can cost between $640,000 and $950,000 (King et al., 2012). Given the number of ships which will be required to fit such systems, costs are clearly significant. Determining an optimal BWM standard is thus an economic problem of balancing costs and expected damages. The current BWM standard has been arrived at via a process of deliberation (Gollasch et al., 2007) and hence, not by empirical economic analysis of the costs and reduced expected damages resulting from different standards. Such an empirical derivation of an optimal BWM standard is challenging due to the global and complex nature of BWM. In Molnar et al. (2008), aggregation of data sets identified 329 invasive species, ranging from fish to plants, and algae to molluscs and crustaceans. For each species, diverse types of information must be collected. In non-exclusive general terms, data should include ecological and economic impact, geographic extent, invasive potential (how "invasive" a species is) and potential management options (for if the species does become established) including their costs and efficacy. The challenges of empirically deriving an optimal BWM standard necessitate a sound theoretical understanding of the mechanisms behind such standards. In the first instance, as will be the focus of this paper, such a theoretical analysis can provide insights into the conditions required for the current BWM standard, which aims to eliminate risk, to be optimal. In turn, such an understanding serves to focus natural science research on the most policy relevant questions. In the second instance, a theoretical understanding of the mechanisms behind an optimal BWM standard can aid future deliberative decision making.

Theoretical insights into the determinants of an optimal BWM standard must respect that marine alien species invasions are best conceptualised as being irreversible (Vitousek et al., 1997; Parker et al., 1999), such that once an invasion becomes established in an area, it cannot be completely eradicated. While there has been little literature which specifically considers BWM, the economics of BWM is grounded in the economics of irreversible events as in Tsur and Zemel (1998, 2004). Irreversible invasions have been considered in contexts other than BWM by Knowler and Barbier (2005), Kim et al. (2006) and Burnett et al. (2012). Knowler and Barbier (2005) consider the risks from importing exotic plant species. Kim et al. (2006) consider the trade-off between measures to reduce the risk of invasion establishment and measures to control an invasion should it become established. They do so in a general setting given uncertain discovery times of an invasion. Burnett et al. (2012) study two specific terrestrial invasive species on the island of Hawaii. Related literature considers issues such as tariffs (e.g. Costello and McAusland, 2003) and inspections (e.g. Springborn, 2014), or considers invasive species establishment or arrival as random, i.e. Bernoulli processes (Horan et al., 2002; Leung et al., 2002; Burnett et al., 2006; Finnoff et al., 2007; Adams and Lee, 2012; Hyytiäinen et al., 2013). Therefore, this study contributes to the literature by analysing the optimal management of irreversible events in the specific context of BWM, with the aim of providing theoretical insights into the determinants of an optimal BWM standard. This specifies the most general modelling framework (as provided by Finnoff et al., 2010b) for the particular case of BWM.

Our method is to build and analyse a model of irreversible invasions using a hazard function, which determines the hazard rate of successful invasion establishment as a function of ballast water treatment effort. We find that if the hazard function is continuously differentiable, then a standard which reduces the hazard rate to zero cannot be optimal. This is interesting because the current BWM standard is chosen to achieve a hazard rate of zero. Therefore, whether or not the hazard rate is continuously differentiable has important implications for the evaluation of the current BWM standard. Accordingly, we consider in detail the differential properties of the hazard function which depends on the technology function, the Propagule Dose-Response (PD-R) function and network effects, each of which we shall now outline.

Technology determines the relationship between treatment effort and the number of propagules (individual organisms) in ballast water. Reductions in the number of propagules is normally achieved via a combination of technologies, such as filtration, biocides and UV light (Dobbs and Rogerson, 2005). The design and testing of ballast water treatment systems is an important topic in marine engineering (see Tsolokai and Diamadopolous (2010) for a review). The PD-R function determines the relationship between the number of propagules in a single emission of ballast water and the probability of invasion establishment. The PD-R function is a central topic of study in invasion ecology (Lockwood et al., 2005; Simberloff, 2009; Blackburn et al., 2015). The PD-R function includes a Minimum Viable Population (MVP) below which the hazard rate is zero. At and above the MVP, the PD-R function can be approximated as being concave (Leung et al., 2004) but may also display convexity at lower numbers of propagules. This convexity occurs due to an Allee effect[2] (Leung et al., 2004). In natural resource economics, an Allee effect is a property of the growth function. This study does not employ a growth function. Instead, we follow the definition of Williamson (1989) who defines an Allee effect as occurring when smaller populations face disproportionately greater probabilities of extinction due to poor resilience to fluctuations in birth and death rates or environmental shocks. This is the common meaning of the term in invasion ecology (see, for example, Blackburn et al., 2015). Evidence for such an Allee effect has been found for many marine species (Kramer et al., 2009). Allee effects are therefore likely to be relevant for marine invasive species (Williamson, 1989; Lee et al., 2013; Blackburn et al., 2015).

Network effects recognise the complex interactions between areas from which invasive species emanate, and to which they can spread (Hulme, 2009; Keller et al.,

[2] Note that there is some inconsistency in the use of the term Allee effect. Sometimes it is used to refer to minimum viable populations. Allee effects are also often employed in a deterministic context to refer to convexity in the function mapping current population size to future population size. In this chapter, we use the term Allee effect solely to refer to convexity in the PD-R function as in Leung et al. (2004) and Lee et al. (2013).

2011; Liu and Tsai, 2011). A central aspect of this complexity is that the more countries which have been invaded by a particular species, the greater the probability that another country will be invaded. This implies that there is a network effect from setting a stricter standard. If the standard is made stricter, then individual countries benefit directly, because they are less likely to be invaded, and also indirectly, because it is less likely that other countries will become invaded and function as stepping stones for the invasion.

The analysis of the technology function, the PD-R function and network effects shows that continuous differentiability of the hazard function depends principally on whether an Allee effect exists. If an Allee effect exists then it is possible that the hazard function is continuously differentiable such that a standard other than one which achieves a probability of establishment of zero (as is aimed for by the BWM convention) must be optimal. This demonstrates the importance of Allee effects for BWM and the need to empirically investigate the PD-R functions of invasive species. While the PD-R function is indeed a central topic in invasion ecology, the issue of continuous differentiability of the PD-R function is not a central concern. This study therefore demonstrates that, from an economic perspective, the continuous differentiability of the PD-R function is of central importance. This study also contributes by conceptualising the roles of technology, the PD-R function, network effects and irreversibility in the context of BWM, which facilitates an understanding of the theoretical determinants of optimal BWM standards.

Section 2.2 describes how the PD-R function, technology function and network effects determine the hazard rate of successful invasion establishment. This section then uses the hazard rate to determine the damage function and shows how the damage function, along with the costs of BWM, constitute the loss function, which is minimised to determine an optimal standard. Section 2.3 analyses the loss function, demonstrates the significance of differential continuity in the PD-R and analyses the role of an Allee effect in terms of the differential continuity of the loss function. Section 2.4 discusses the generalisability of the results given our modelling assumptions and concludes.

2.2 The model

The description of the model begins by specifying the PD-R and technology functions. Next, we consider how the PD-R and technology functions are included, along with network effects, in the hazard function. We proceed to show how the hazard function is used in the expected damage function. Finally, the objective function (the loss function) is defined, which contains the expected damage and cost functions.

2.2.1 The PD-R and technology functions

Let x be the amount of ballast water treatment effort. Ballast water treatment effort reduces the number of propagules of a single invasive species in a single emission of treated ballast water. The number of propagules is given by q, which is determined by the technology function[3] $Q(x)$, i.e. $q = Q(x)$. The number of propagules determines the probability of invasion establishment from a single emission event and is given by $p = P(q)$, which is the PD-R function. Thus the probability of invasion establishment from a single emission event is given by $p = P\big(Q(x)\big)$. We proceed to specify the technology and PD-R functions.

Ballast water treatment systems are predominantly designed to meet the standards set in the BWM convention. Therefore there are too few data points to gain insights into the form of the technology function. We make the assumption that marginal effectiveness of effort declines in treatment effort such that removing the n^{th} propagule from ballast water requires more effort than removing the first. Formally, we assume that the technology function is downward sloping and convex, i.e. $\frac{\partial Q(x)}{\partial x} < 0$ and $\frac{\partial^2 Q(x)}{\partial x^2} > 0$.

Let us define the Maximum Non-Viable Population (MN-VP) as the number of propagules marginally below the MVP. The MN-VP is therefore the largest number of propagules for which the probability of establishment is zero. The PD-R function $P(q)$ includes an MN-VP at $\bar{q}$ propagules, where $\bar{q} \equiv Q(\bar{x})$, such that $P(q \leq \bar{q}) = 0$.

[3] Note that throughout this chapter, when both upper and lower case letters are used, the lower case letters denote the outputs of functions and upper case letters denote the functions themselves.

For $q > \bar{q}$, $P(q)$ is increasing in q. The function $P(q)$ may or may not be continuously differentiable (smooth) at $\bar{q}$. Whether the function $\bar{q}$ increases at an increasing or decreasing rate depends on whether an Allee effect is present for the species in question. If no Allee effect exists then $P(q)$ increases at a decreasing rate, i.e. it is concave. If an Allee effect is present then there exists some interval of $q \in (\bar{q}, b]$, where $b > \bar{q}$, such that $\left.\frac{\partial^2 P(q)}{\partial q^2}\right|_{q<b} > 0$, $\left.\frac{\partial^2 P(q)}{\partial q^2}\right|_{q=b} = 0$ and $\left.\frac{\partial^2 P(q)}{\partial q^2}\right|_{q>b} < 0$. We will consider the cases in which an Allee effect is present and not present in the function $P(q)$ in order to understand the role of an Allee effect in determining an optimal standard. These cases are shown in Figure 2.1.

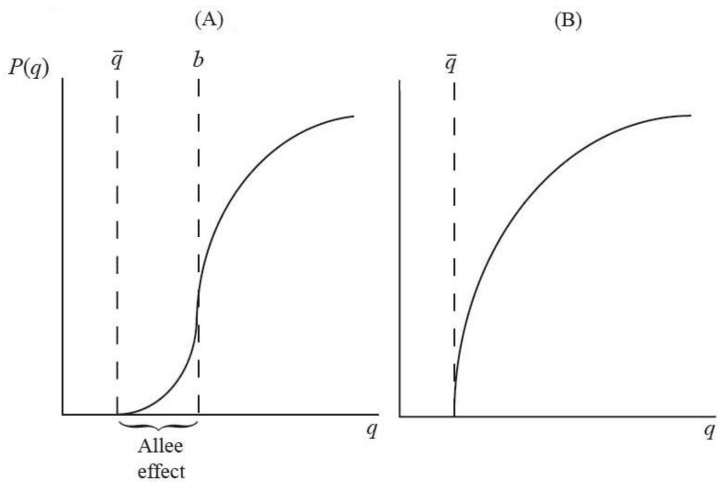

Figure 2.1: Illustrations of the PD-R function with an Allee effect (Panel A) and without an Allee effect (Panel B) where $\bar{q}$ is the MN-VP.

The PD-R and technology functions thus determine the probability of invasion establishment from a single emission of ballast water. This probability is used to determine the hazard rate of invasion establishment.

2.2.2 The hazard rate function

The function $p = P\big(Q(x)\big)$ is used to determine the hazard rate of invasion establishment at time t, along with the number of invaded countries at time t, which is given by n_t. The number of invaded countries will increase over time at a rate which depends on the probability of establishment from a single emission event. Therefore $n_t = N\left(t, P\big(Q(x)\big)\right)$. The hazard rate of invasion establishment at time t is given by r_t, which is determined by the hazard function $R\left(P(Q(x)), N\big(t, P(Q(x))\big)\right)$. Hence, the hazard rate is determined directly by $P\big(Q(x)\big)$ and also indirectly, via the effect of $P\big(Q(x)\big)$ on the number of invaded countries. To shorten the length of equations, we will often denote $R\left(P(Q(x)), N\big(t, P(Q(x))\big)\right)$ simply as $R(p, n_t)$.

The MN-VP is an important determinant of the form of the hazard function. Because $R\left(P(Q(x \geq \bar{x})), N\big(t, P(Q(x \geq \bar{x}))\big)\right) = 0$, there is no benefit in reducing the number of propagules to below the MN-VP and as such, the domain of x is $[0, \bar{x}]$. Appendix 2.1 shows that the first derivative of the hazard function with respect to x is negative, i.e. $\left.\frac{\partial R(p,n_t)}{\partial x}\right|_{x<\bar{x}} < 0$. Whether or not the hazard function is continuously differentiable depends solely on whether the PD-R function is continuously differentiable. When an Allee effect is not present (Figure 1(B)), it is clear that that $\left.\frac{\partial P(q)}{\partial q}\right|_{q=\bar{q}}$ is not defined. The PD-R function is therefore not continuously differentiable and in turn, the hazard function is not continuously differentiable. The hazard function is continuously differentiable if it holds that $\lim_{\bar{q}\leftarrow q} P(q) = 0$. It can only hold that $\lim_{\bar{q}\leftarrow q} P(q) = 0$ in the presence of an Allee effect. However, an Allee effect is not sufficient for $\lim_{\bar{q}\leftarrow q} P(q) = 0$. An Allee effect is therefore a necessary condition for differential continuity in the hazard function.

The hazard function is therefore determined by the technology and PD-R functions and network effects. The form of the hazard function, and in particular, whether it is continuously differentiable, effects the form of the expected damage function.

2.2.3 The expected damage function

Given the hazard function, we now proceed to specify the expected damage function. Let Θ be the time at which the invasion occurs. We assume that Θ has an exponential probability density function given by $r_t exp\big(-\int_0^t r_\tau d\tau\big)$. In other words, the probability that the invasion will happen at time t is determined by the hazard rate r_t, conditional upon the invasion not having happened in the period up to time

t, given by $exp\left(-\int_0^t r_\tau d\tau\right)$. Expected damage is determined by the annual damage, should the invasion arrive, given by y and the discount rate, δ, such that the present value of the damage from an invasion (evaluated from $t = 0$ to ∞) is given by $\frac{y}{\delta}$. The present value of damage is discounted according to the discount factor for the time in which the invasion establishes, $exp(-\delta t)$. Therefore, expected damage is given by

$$Z(x) = \frac{y}{\delta}\int_o^\infty exp(-\delta t)R(p, n_t)exp\left(-\int_0^t R(p, n_\tau)d\tau\right)dt, \tag{2.1}$$

or

$$Z(x) = \frac{y}{\delta}\int_o^\infty R(p, n_t)exp\left(-\delta t - \int_0^t R(p, n_\tau)d\tau\right)dt. \tag{2.2}$$

The expected damage function is therefore determined by the hazard function and the irreversibility of an invasion. The expected damage function can now be incorporated into the objective function, which we term the loss function.

2.2.4 The loss function

The loss function is the objective function of a benevolent social planner who is concerned with a single invasive species. The loss is the sum of the cost of implementing the BWM standard and the expected damage under the standard, z, for the total number of countries *m,* which we assume to be symmetric in that they face the same damages from an established invasion and the same costs of treatment. For simplification, we assume that all countries export and import the same amount of ballast water and that the export from a given country (and thus also import to a given country) is evenly distributed across all countries. Further, we assume that all vessels carry the same amount of ballast water on all journeys.

The objective of the social planner is to minimise expected loss by choosing treatment effort x, which is the control variable. Effort is chosen at $t = 0$ and remains constant thereafter. This is appropriate because of the impracticality of frequently changing the BWM standard. The BWM convention will only come into force if enough states ratify the agreement such that 35% of global shipping is covered by the agreement. Attempts to achieve a sufficient number of ratifications have been ongoing since 2004, but enough ratifications have not as yet been achieved (Scriven et al., 2015). While it is technologically possible to change treatment effort frequently, this is clearly impractical given the amount of time it takes to ratify the current standard. Therefore, we assume that effort is chosen once and then that it remains at that level.

The loss function is determined by the expected damage function (and thus the hazard function), and the cost of treatment effort. Effort has a unit cost k, where $k > 0$, such that kx gives the costs per country of a given level of treatment. Remember that the technology function displays decreasing marginal returns to effort. Therefore, there are increasing marginal costs of reducing the number of propagules. Expected damages and costs of treatment are used in the loss function, which is given b

$$L(x) = m\big(kx + Z(x)\big), \tag{2.3}$$

where m is the total number of countries. We therefore assume that all countries have not yet suffered an established invasion. We also assume that countries must continue ballast water treatment after they suffer an established invasion. This is reasonable because the BWM standard is binding after ratification regardless of whether an invasion has become established.

2.3 Analysis of optimal standards

This section proceeds to analyse the loss function in terms of an optimal BWM standard. The First Order Condition (FOC) for an extremum and the Second Order Condition (SOC) for a minimum are $\frac{\partial L(x)}{\partial x} = 0$ and $\frac{\partial^2 L(x)}{\partial x^2} > 0$ respectively. Given that

$$\frac{\partial L(x)}{\partial x} = m\left(k + \frac{\partial Z(x)}{\partial x}\right) \tag{2.4}$$

and

$$\frac{\partial^2 L(x)}{\partial x^2} = m\frac{\partial^2 Z(x)}{\partial x^2}, \tag{2.5}$$

the FOC and SOC can be represented as $\frac{\partial Z(x)}{\partial x} = -k$ and $\frac{\partial^2 Z(x)}{\partial x^2} > 0$. Appendix 2.2 demonstrates that $\frac{\partial Z(x)}{\partial x} < 0$. Appendix 2.2 also provides $\frac{\partial^2 Z(x)}{\partial x^2}$ and shows that $\frac{\partial^2 Z(x)}{\partial x^2}$

is ambiguous in sign. Therefore, the damage function may be either concave or convex. Further, the damage function may display convexity for part of the domain of x and concavity for another part.

We have established that the hazard function may or may not be continuously differentiable in Section 2.2.1. Further, Section 2.2.1 demonstrated that differential continuity of the hazard function requires differential continuity of the PD-R function. This has implications for the damage function. To repeat, continuous differentiability of the hazard function means that $\lim_{x \to \bar{x}} R(p, n_t) = 0$. Considering Equation (2.2), we see that $\lim_{x \to \bar{x}} R(p, n_t) = 0$ means that $\lim_{x \to \bar{x}} Z(x) = 0$, which means that the damage function is continuously differentiable. Further, $\lim_{x \to \bar{x}} Z(x) = 0$ implies that the damage function is convex as it approaches $\bar{x}$. Formally, let us define a convex sub-domain of the function $Z(x)$ as the interval $[a, \bar{x}]$ where $0 \leq a < \bar{x}$ and where a is equal to the smallest value of x for which $\frac{\partial^2 Z(x)}{\partial x^2} > 0$. In the interval $[a, \bar{x}]$, it holds that $\left.\frac{\partial^2 Z(x)}{\partial x^2}\right|_{x > a} > 0$. Because $\lim_{x \to \bar{x}} R(x) = 0$ implies $\lim_{x \to \bar{x}} Z(x) = 0$, continuous differentiability of the hazard function implies that there must exist such a convex sub-domain in $Z(x)$. Given that the hazard function will be continuously differentiable if the PD-R function is continuously differentiable, $\lim_{\bar{q} \leftarrow q} P(q) = 0$ implies $\lim_{x \to \bar{x}} Z(x) = 0$. Therefore, a convex sub-domain exists in $Z(x)$.

Given this understanding of the properties of the damage function, we proceed to analyse the implications for the optimal standard, as determined by the FOC and the SOC under the assumption that the PD-R is continuously differentiable. Figure (2.2) illustrates two possible forms for $Z(x)$, depending on the range of the convex sub-domain, and two possible marginal costs k_1 and k_2. The implications of these functions for the loss function are also illustrated.

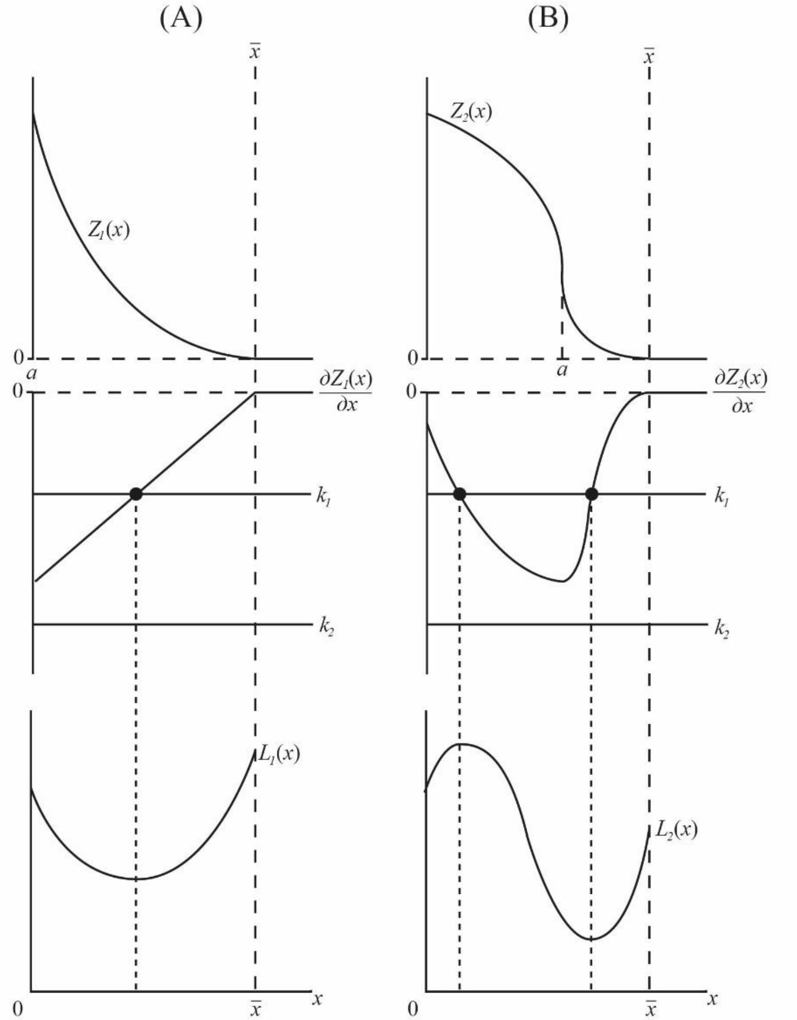

Figure 2.2: Diagrammatic illustration of forms of the damage function and the marginal damage function if the PD-R function is continuously differentiable, and the implications thereof for minimisation of the loss function.

In Panel (A), the convex sub-domain $[a, \bar{x}]$ is equal to the domain $[0, \bar{x}]$. The damage function is therefore decreasing at an increasing rate for all x such that the SOC is satisfied for all x. If marginal cost is equal to k_1, an interior minimum exists (the FOC is satisfied), as shown by the loss function. When marginal cost is equal to k_2, no interior minimum exists (because the FOC cannot be satisfied) and the optimal standard is a corner solution with zero treatment effort. In Panel (B), the damage function decreases at a decreasing rate for $x < a$ and then at an increasing rate for

$x \geq a$. The convex sub-domain is therefore a subset of the domain $[0, \bar{x}]$, i.e. $[a, \bar{x}] \neq [0, \bar{x}]$ because $a \neq 0$. If the damage function intersects the marginal cost where marginal damage is decreasing then an interior maximum exists (i.e. the FOC is satisfied but the SOC is not). An interior minimum exists in Panel (B) where marginal cost k_1 intersects the marginal damage function in the convex sub-domain. Such an interior minimum may not be a global minimum: a BWM standard of $x = 0$ may be a global minimum.

Under the assumption that the PD-R function is continuously differentiable, we see that the solution is driven by whether or not $[a, \bar{x}] = [0, \bar{x}]$ and whether or not there exists some x such that $\frac{\partial Z(x)}{\partial x} = -k$. We summarise the possible outcomes in terms of these two drivers in Table 2.1.

Table 2.1: Possible solutions under the assumption that the PD-R function is continuous.

	$[a, \bar{x}] = [0, \bar{x}]$	$[a, \bar{x}] \neq [0, \bar{x}]$
$\frac{\partial Z(x)}{\partial x} = -k$ is satisfied for some x.	A global interior solution exists.	A global interior minimum exists **or** an interior local minimum exists with a global minimum at $x = 0$.
$\frac{\partial Z(x)}{\partial x} \neq -k \ \forall \ x$.	A global corner solution exists at $x = 0$.	A global corner solution exists at $x = 0$.

Table 2.1 shows all possible outcomes and therefore demonstrates that, given the assumption that $k > 0$, a minimum where $x = \bar{x}$ is not possible if the PD-R is continuously differentiable. The reason for this is that treatment effort always has a strictly positive marginal cost ($k > 0$). No matter how low the marginal cost is, there will always be a treatment level where the FOC is satisfied. Further, at very low marginal cost, the FOC would be satisfied in the convex sub-domain, where the SOC is satisfied also.

Let us proceed to analyse the case where the PD-R function is not continuously differentiable. In this case, it is possible that $\bar{x}$ is the optimal amount of treatment effort. If the PD-R function is not continuously differentiable around $\bar{x}$, then the hazard function is not continuously differentiable around $\bar{x}$ and therefore marginal damage is not continuous around $\bar{x}$. This implies that a marginal cost k can exist such that $-k > \frac{\partial Z(x)}{\partial x}$ $\forall$ x. This means that a global corner solution can exist where $x = \bar{x}$. This is illustrated in Figure 2.3(A). Alternatively, a global corner solution can exist where $x = \bar{x}$ because no convex sub-domain exists, as illustrated in Figure 2.3(B). If the PD-R function is not continuously differentiable then it is not necessary that a convex sub-domain exists.

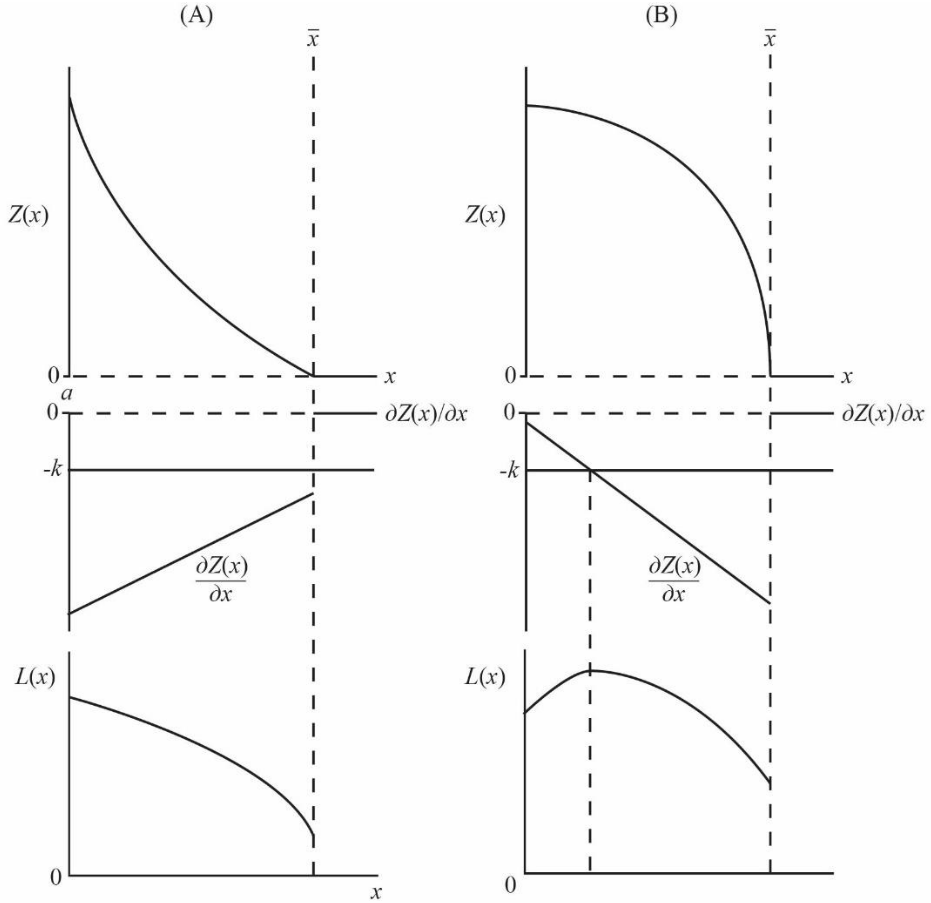

Figure 2.3: Diagrammatic illustration of the possibility for an optimal standard where $x = \bar{x}$ when the hazard function not continuously differentiable. In Panel (A), $x = \bar{x}$ is optimal because $-k > \frac{\partial Z(x)}{\partial x}$ $\forall$ x. In Panel (B), $x = \bar{x}$ is optimal because the SOC is never satisfied and, in this example, loss is not minimised where $x = 0$.

We summarise our results up to now in Result 2.1.

Result 2.1: *A non-continuously differentiable PD-R function is a necessary condition for the optimality of a level of treatment effort equal to $\bar{x}$.*

Furthermore, we can also establish the following result.

Result 2.2: *The existence of an Allee effect is a necessary condition for continuous differentiability in the hazard function.*

For proof of Result 2.2, consider that for the PD-R function to be continuously differentiable around the MN-VP, it must hold that $\lim_{\bar{q}\leftarrow q} P(q) = 0$. In turn, this implies that a sub-domain $[\bar{q}, b]$ must exist where $\left.\frac{\partial^2 P(q)}{\partial q^2}\right|_{q\in[\bar{q},b]} \geq 0$, i.e. where the PD-R function is convex. This is definition of an Allee effect. Panel (A) of Figure 2.1 shows the case where no Allee effect is present, for which it is clear that $\lim_{\bar{q}\leftarrow q} P(q) = 0$ cannot hold. However, even if an Allee effect does exist, this does not guarantee that $\lim_{\bar{q}\leftarrow q} P(q) = 0$, i.e. the PD-R function may still be not be continuously differentiable around the MN-VP. An Allee effect is therefore a necessary but not a sufficient condition for continuity in the hazard function.

2.4 Conclusions

This study analyses a model of the optimality of Ballast Water Management (BWM) standards. The objective of a benevolent social planner is to minimise loss, which is the sum of the expected damages from an invasion and the costs of treatment effort. Expected damage is calculated in accordance with an irreversible invasion with an exponential probability density function. Expected damage depends on the hazard rate of invasion establishment. The hazard function determines the hazard rate and includes the technology function, Propagule Dose-Response (PD-R) function and network effects. The PD-R function includes a Minimum Viable Population (MVP) and may include an Allee effect. We analyse the possibility for a BWM standard which eliminates the risk of invasion establishment (as is aimed for by the BWM convention) to be optimal. This is achieved by reducing the number of propagules in ballast water emissions to below the MVP. We find that such a standard can only be optimal if the PD-R function is not continuously differentiable around the MVP. Furthermore, we find that the PD-R function will not be continuously differentiable around the MVP in the absence of an Allee effect.

These results have both shorter term implications regarding future research as well as longer term value regarding decision making on BWM. In the shorter term, this study contributes by demonstrating that PD-R functions and in particular, the prevalence of Allee effects, are vital information needs for policy making. In addition to the importance of Allee effects in invasion ecology, this study demonstrates the Allee effects have important policy implications. What then are the prospects for gaining more information about Allee effects to inform decision making? It is well established that Allee effects play an important role in determining the probabilities of establishment of small populations (Blackburn et al., 2015). However, Allee effects are only part of the story. Demographic and environmental stochasticity as well as genetic effects all play a role. Disentangling the role of Allee effects requires empirical study such as that by Duncan et al. (2014). Duncan et al. (2014) analyse data from 55 experimental releases of a non-native terrestrial insect species into New Zealand. They fit the data to different PD-R functions. They find that the probability of survival was best explained by a function combining demographic stochasticity plus Allee effects. However, the best fit PD-R function does not converge to zero as the number of propagules approaches zero. This is evidence that the PD-R function for this particular species is not continuously differentiable. The analysis of Duncan et al. (2014) demonstrates that, if data is available, statistical analysis can provide evidence regarding the differentiable continuity of the PD-R function.

The longer term implications result from our conceptualisation of the key mechanisms behind optimal BWM standards. This demonstrates how the mechanisms can be incorporated into an optimisation framework. In turn, this provides the theoretical foundations for BWM standards from an economic perspective, which has, up to now, received little attention. The combination of the theoretical foundation provided in this study and the collection of more data to be analysed in the same manner as Duncan et al. (2014), can be employed in an evaluation of the current BWM standard. Evaluation of the current BWM standard is particularly necessary because of the problems in achieving sufficient signatories to ratify the current agreement. The analysis of Scriven et al. (2015) shows that the rate of ratification increased in the five years to 2009, but thereafter has slowed dramatically. They argue that there may still be a long period before the BWM convention is put into force. Our study shows that the current BMW standard can only be optimal under specific conditions. Empirical insights into whether these conditions are met will help to evaluate the current standards. An evaluation of current standards may provide evidence which encourages nations to ratify the convention, or alternatively, for the IMO to modify the convention such that more nations will ratify it.

Our findings are subject to a number of simplifying assumptions. The most significant assumption relates to our conceptualisation of BWM as a single species problem. BWM is inherently a multispecies problem. The most important issue to deal with when considering the multispecies case is that the PD-R function will vary between species. In this case, the extent of continuous differentiability in the expected damage function will depend on the PD-R functions of all invasive species. The generalisability of our results to the multispecies case is an important issue for future research. Future research should also consider that the standard set in the BWM convention actually consists of three standards for different size classes of organism. Organisms are therefore either classed as small, medium or large sized and different concentrations of organisms in are permitted for each size class. It may be the case, for example, that Allee effects are more prominent for one size class than for another, and this would need to be taken into account for standards, by size class, to be optimal.

Appendix 2.1 Derivatives of the hazard function

This appendix shows that the hazard function is downward sloping in effort and, for completeness, derives the second derivative of the hazard function with respect to effort. To do so, we omit the arguments of functions in order to save space.

The first derivative of $R\left(P(Q(x)), N\big(t, P(Q(x))\big)\right)$ is

$$\frac{\partial R}{\partial x} = \frac{\partial R}{\partial p}\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x} + \frac{\partial R}{\partial n_t}\frac{\partial N}{\partial p}\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x}. \tag{2.i}$$

Given that $\frac{\partial R}{\partial p}$, $\frac{\partial P}{\partial q}$, $\frac{\partial R}{\partial n_t}$ and $\frac{\partial N}{\partial p}$ are positive in sign and that $\frac{\partial Q}{\partial x}$ is negative in sign, it holds that $\left.\frac{\partial R}{\partial x}\right|_{x<\bar{x}} < 0$.

The second derivative is given by

$$\frac{\partial^2 R}{\partial x^2} = \frac{\partial}{\partial x}\left(\frac{\partial R}{\partial p}\right)\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x} + \frac{\partial}{\partial x}\left(\frac{\partial P}{\partial q}\right)\frac{\partial R}{\partial p}\frac{\partial Q}{\partial x} + \frac{\partial}{\partial x}\left(\frac{\partial Q}{\partial x}\right)\frac{\partial R}{\partial p}\frac{\partial P}{\partial q} + \frac{\partial}{\partial x}\left(\frac{\partial R}{\partial n_t}\right)\frac{\partial N}{\partial p}\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x} + \frac{\partial}{\partial x}\left(\frac{\partial N}{\partial p}\right)\frac{\partial R}{\partial n_t}\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x} + \frac{\partial}{\partial x}\left(\frac{\partial P}{\partial q}\right)\frac{\partial R}{\partial n_t}\frac{\partial N}{\partial p}\frac{\partial Q}{\partial x} + \frac{\partial}{\partial x}\left(\frac{\partial Q}{\partial x}\right)\frac{\partial R}{\partial n_t}\frac{\partial N}{\partial p}\frac{\partial P}{\partial q} \quad (2.ii)$$

$$= \frac{\partial^2 Q}{\partial x^2}\frac{\partial P}{\partial q}\left(\frac{\partial R}{\partial p} + \frac{\partial R}{\partial n_t}\frac{\partial N}{\partial p}\right) + \frac{\partial}{\partial x}\left(\frac{\partial P}{\partial q}\right)\frac{\partial Q}{\partial x}\left(\frac{\partial R}{\partial p} + \frac{\partial R}{\partial n_t}\frac{\partial N}{\partial p}\right) + \frac{\partial}{\partial x}\left(\frac{\partial R}{\partial p}\right)\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x} + \frac{\partial}{\partial x}\left(\frac{\partial R}{\partial n_t}\right)\frac{\partial N}{\partial p}\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x} + \frac{\partial}{\partial x}\left(\frac{\partial N}{\partial p}\right)\frac{\partial R}{\partial n_t}\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x} \quad (2.iii)$$

$$= \left(\frac{\partial R}{\partial p} + \frac{\partial R}{\partial n_t}\frac{\partial N}{\partial p}\right)\left(\frac{\partial^2 Q}{\partial x^2}\frac{\partial P}{\partial q} + \frac{\partial}{\partial x}\left(\frac{\partial P}{\partial q}\right)\frac{\partial Q}{\partial x}\right) + \frac{\partial P}{\partial q}\frac{\partial Q}{\partial x}\left(\frac{\partial}{\partial x}\left(\frac{\partial R}{\partial p}\right) + \frac{\partial}{\partial x}\left(\frac{\partial R}{\partial n_t}\right)\frac{\partial N}{\partial p} + \frac{\partial}{\partial x}\left(\frac{\partial N}{\partial p}\right)\frac{\partial R}{\partial n_t}\right). \quad (2.iv)$$

Let us establish the following four identities from Equation (2.iv).

$$\frac{\partial}{\partial x}\left(\frac{\partial P}{\partial q}\right) = \frac{\partial^2 P}{\partial q^2}\frac{\partial Q}{\partial x} \quad (2.v)$$

$$\frac{\partial}{\partial x}\left(\frac{\partial R}{\partial p}\right) = \frac{\partial P}{\partial q}\frac{\partial Q}{\partial x}\left(\frac{\partial^2 R}{\partial p^2} + \frac{\partial^2 R}{\partial n_t \partial p}\frac{\partial N}{\partial p}\right) \quad (2.vi)$$

$$\frac{\partial}{\partial x}\left(\frac{\partial R}{\partial n_t}\right)\frac{\partial N}{\partial p} = \frac{\partial P}{\partial q}\frac{\partial Q}{\partial x}\frac{\partial N}{\partial p}\left(\frac{\partial^2 R}{\partial n_t^2} + \frac{\partial^2 R}{\partial n_t \partial p}\right) \quad (2.vii)$$

$$\frac{\partial}{\partial x}\left(\frac{\partial N}{\partial p}\right)\frac{\partial R}{\partial n_t} = \frac{\partial R}{\partial n_t}\frac{\partial^2 N}{\partial p^2}\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x} \quad (2.viii)$$

Substituting Equations (2.v) through (2.viii) into (2.iv) gives

$$\frac{\partial^2 R}{\partial x^2} = \left(\frac{\partial R}{\partial p} + \frac{\partial R}{\partial n_t}\frac{\partial N}{\partial p}\right)\left(\frac{\partial^2 Q}{\partial x^2}\frac{\partial P}{\partial q} + \frac{\partial^2 P}{\partial q^2}\left(\frac{\partial Q}{\partial x}\right)^2\right) + \left(\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x}\right)^2\left(\frac{\partial^2 R}{\partial p^2} + \frac{\partial^2 R}{\partial n_t \partial p}\frac{\partial N}{\partial p} + \frac{\partial N}{\partial p}\left(\frac{\partial^2 R}{\partial n_t^2} + \frac{\partial^2 R}{\partial n_t \partial p}\right) + \frac{\partial R}{\partial n_t}\frac{\partial^2 N}{\partial p^2}\right) \quad (2.ix)$$

$$= \left(\frac{\partial R}{\partial p} + \frac{\partial R}{\partial n_t}\frac{\partial N}{\partial p}\right)\left(\frac{\partial^2 Q}{\partial x^2}\frac{\partial P}{\partial q} + \frac{\partial^2 P}{\partial q^2}\left(\frac{\partial Q}{\partial x}\right)^2\right) + \left(\frac{\partial P}{\partial q}\frac{\partial Q}{\partial x}\right)^2\left(\frac{\partial^2 R}{\partial p^2} + \frac{\partial N}{\partial p}\left(2\frac{\partial^2 R}{\partial n_t \partial p} + \frac{\partial^2 R}{\partial n_t^2}\right) + \frac{\partial R}{\partial n_t}\frac{\partial^2 N}{\partial p^2}\right). \quad (2.x)$$

We do not require an assumption about the sign of the second derivative of the hazard function because, as demonstrated in Appendix 2.2, the sign of the second derivative of the damage function is ambiguous in sign regardless of whether the hazard function is concave or convex.

Appendix 2.2 Derivatives of the damage function

This appendix provides the first derivative of the damage function and explains our assumption that the damage function is downward sloping in x. This appendix also provides the second derivative of the damage function and shows that it is ambiguous in sign. The damage function is

$$Z(x) = \frac{y}{\delta}\int_o^\infty R(p,n_t)exp\left(-\delta t - \int_0^t R(p,n_\tau)d\tau\right)dt. \tag{2.xi}$$

The first differential of the damage function is

$$\frac{\partial Z(x)}{\partial x} = \frac{y}{\delta}\int_o^\infty \begin{pmatrix} \frac{\partial R(p,n_t)}{\partial x} exp\left(-\delta t - \int_0^t R(p,n_\tau)d\tau\right) \\ -R(p,n_t)\int_0^t \frac{\partial R(p,n_\tau)}{\partial x} d\tau \, exp\left(-\delta t - \int_0^t R(p,n_\tau)d\tau\right) \end{pmatrix} dt, \tag{2.xii}$$

which can be simplified to

$$\frac{\partial Z(x)}{\partial x} = \frac{y}{\delta}\int_o^\infty exp\left(-\delta t - \int_0^t R(p,n_\tau)d\tau\right)\left(\frac{\partial R(p,n_t)}{\partial x} - R(p,n_t)\int_0^t \frac{\partial R(p,n_\tau)}{\partial x} d\tau\right)dt. \tag{2.xiii}$$

Note that $\frac{\partial R(p,n_t)}{\partial x} < 0$ and therefore $\frac{\partial Z(x)}{\partial x}$ is ambiguous in sign. For example, if

$$R(p,n_t)\int_0^t \frac{\partial R(p,n_\tau)}{\partial x} d\tau < \frac{\partial R(p,n_t)}{\partial x} \quad \forall \; t, \tag{2.xiv}$$

then it will hold that $\frac{\partial Z(x)}{\partial x} > 0$. We assume that $\frac{\partial Z(x)}{\partial x} < 0$ hold for the following reason, which relates to the "tens rule" and the effects of discounting. The tens rule of Williamson (1996) is a well-established general rule determining the probability of invasion establishment. It states that there is a 10% chance that a species will arrive in a given year, a 10% chance that the species will become established and a 10% chance that the species will cause economic damages. Hence, a reasonable approximation of the probability of invasion establishment in a given year is $Pr = 0.001$. The hazard rate for an event with probability of occurrence within a year of $Pr = 0.001$ is calculated by solving $0.001 = 1 - e^{-rt}$ for r where $t = 1$ which gives

$R(p,n_t) \approx 0.001$. This implies that $R(p,n_t)\int_0^t \frac{\partial R(p,n_\tau)}{\partial x}d\tau$ will be a small relative to $\frac{\partial R(p,n_t)}{\partial x}$ even if the integral $\int_0^t \frac{\partial R(p,n_\tau)}{\partial x}d\tau$ is taken over many periods. Further, (A2.4) is most likely to hold for higher values of t, because the integral is calculated over a longer time period. However, as t increases, discounting increases also. Discounting thus means that for points in time where (A2.4) is most likely to hold, the fact that it holds will have very little effect on the sign of $\frac{\partial Z(x)}{\partial x}$.

We obtain the second derivative from the first derivative. After simplifying, we obtain

$$\frac{\partial^2 Z(x)}{\partial x^2} = \frac{y}{\delta}\int_o^\infty exp\left(-\delta t - \int_0^t R(p,n_\tau)d\tau\right)\begin{bmatrix} R(p,n_t)\left(\int_0^t \frac{\partial R(p,n_\tau)}{\partial x}d\tau\right)^2 \\ -2\frac{\partial R(p,n_t)}{\partial x}\int_0^t \frac{\partial R(p,n_\tau)}{\partial x}d\tau \\ \frac{\partial^2 R(p,n_t)}{\partial x^2} - R(p,n_t)\int_0^t \frac{\partial^2 R(p,n_\tau)}{\partial x^2}d\tau \end{bmatrix}dt. \quad (2.xv)$$

We can analyse the sign of the second derivative by considering the three lines of terms within the square bracket (outside the square bracket, everything is positive in sign). Note that we have made no assumption regarding the second derivative of the hazard function. The first and second lines within the bracket are unaffected by the second derivative of the hazard function. The terms within the first and second lines are therefore positive in sign. However, the second line is preceded by a minus sign. Thus $\frac{\partial^2 Z(x)}{\partial x^2}$ is necessarily ambiguous in sign regardless of whether $\frac{\partial^2 R(p,n_t)}{\partial x^2}$ and $\int_0^t \frac{\partial^2 R(p,n_\tau)}{\partial x^2}d\tau$ are positive or negative in sign. Therefore, the damage function may either be concave or convex. Further, the damage function may display convexity for part of the domain of x and concavity for another part.

3 Invasive species control in a one-dimensional metapopulation network

ABSTRACT

The growth and spread of established Invasive Alien Species (IAS) cause significant ecological and economic damages. Minimising the costs of controlling, and the damages from, IAS depends on the spatial dynamics and uncertainty regarding IAS spread. This study expands on existing modelling approaches by allowing for varying stock sizes within patches and stochastic spread between patches. The objective of this study is to demonstrate the added value from this more detailed modelling approach. This is achieved in the context of coastal and riparian systems, which can be accurately modelled one-dimensional landscape, i.e. a series of patches connected in a line. The model allows for two types of intervention, namely (1) partial or complete removal of the population in within any patch; and (2) containment to reduce spread between patches. We analyse the general properties of the model using a two-patch setup to determine how the optimal policy depends on both the location and size of the invasion in patches. We find that allowing for varying stock sizes within patches facilitates optimal timing of the application of containment. We also identify two novel optimal policies: the combination of containment and removal to stop spread between patches and the application of up to four distinct policies for a single patch depending on the size of the invasion in that patch.

This chapter is based on the paper: Walker, A.N., Poos, J.-J., Groeneveld, R.A., 2015. Invasive species control in a one-dimensional metapopulation network. Ecological Modelling 316, 176-184.

3.1 Introduction

IAS (Invasive Alien Species) are species that proliferate, spread, and persist after introduction into a natural environment (Mack et al., 2000). IAS can cause dramatic changes in ecological systems and have profoundly altered terrestrial and marine ecosystems worldwide (Gurevitch and Padilla, 2004; Hulme, 2006). Although invasions are not necessarily human-driven, the number of invasions has grown substantially as a result of global travel and trade (Mack et al., 2000). Invasions can lead to significant losses in terms of human health, biodiversity, and ecological services (Frésard and Boncoeur, 2006; Pimentel et al., 2005; Scalera, 2010). These losses can be mitigated by appropriate management in response to invasions, informed by scientific decision support (Carrasco et al., 2010a). A better understanding of the costs and benefits of controlling IAS improves management efficiency (Genovesi, 2005). A particular aspect requiring more attention is our understanding of the spatial aspects of invasion control (Albers et al., 2010; Epanchin-Niell and Hastings, 2010; Savage and Renton, 2014).

Much of the literature concerning spatial dynamics is concerned with the interaction of multiple jurisdictions in response to both the invasive species and the actions of other jurisdictions. These include Huffaker et al. (1992), Albers et al. (2010), Sanchirico et al. (2010), Zhang et al. (2010), Carrasco et al. (2012), McDermott et al. (2013) and Fenichel et al. (2014). The literature considering single jurisdictions consisting of multiple spatial areas has the shortcoming that it either does not allow for varying stock sizes within areas (i.e. areas are modelled in binary terms: either invaded or not invaded) or restricts removal of invasions in a given area to complete eradiction only (Carrasco et al., 2010a; Finnoff et al., 2010a, Epanchin-Niell and Wilen, 2012; Epanchin-Niell et al., 2012). Restricting removal of the invasive population to complete eradication only is particularly problematic in the marine context because a policy of eradication is rarely pursued in practice (Vitousek et al., 1997).

The binary restriction (areas are modelled as either invaded or not invaded) limits modelling richness as it excludes within-patch density dependence of damages. Further, the binary restriction limits the set of potential management options. When patches are either invaded or not invaded, the set of management options in terms of reducing the size of the stock is restricted to doing nothing or completely eradicating the invasion in that area. This precludes the identification of optimal management policies which maintain an intermediate invasive population in a given patch.

We therefore construct a model which allows for varying stock sizes in patches and removal of any amount of the population from any patch (cf. Salinas et al., 2005 and Burnett et al., 2007), as opposed to being invaded or non-invaded in a binary sense (as in Epanchin-Niell and Wilen, 2012 and Chadès et al., 2011), in a single jurisdiction setting. Additionally, we allow for a second intervention which we term containment. Containment reduces the probability of spread between patches without affecting the population size within the patch. This paper therefore builds on Burnett et al. (2007), who consider varying population size within patches (but do not allow for measures to directly contain the spread of the invasion) by allowing for a containment intervention, such as employed by Sharov (2004).

Allowing for varying stock sizes in patches increases the dimensionality of the problem. In a network of two patches which can only be invaded or non-invaded there are only four possible states. However, if a patch can either be invaded, invaded at an intermediate population size, or fully invaded, (thus, three possible states for a given patch) then there are nine possible states for the network as a whole. Thus, the computational burden of modelling more complex systems can quickly become problematic. This burden is further increased by our use of two interventions; removal and containment. In this paper, we consider the case of a one-dimensional network, which limits the increased computational burden resulting from varying population size within patches. A one-dimensional network consists of a series of patches connected in a line. Chadès et al. (2011) refer to such a spatial arrangement as a line network and employ line-networks to analyse invasive species management. A one-dimensional network consists of two end patches which are linked to only one other patch, and all other patches are linked to only two other patches such that all the patches, visually, form a line. A one-dimensional network is therefore fully defined by the number of patches. We assume that an invasion can only spread between patches for which there is a connection. Hence, if there are three patches with Patch 1 and 3 as the end patches and Patch 1 is invaded, then Patch 3 can only become invaded after Patch 2 is invaded.

Invasions spreading in coastal and riparian systems are suitable to be modelled as one-dimensional networks. The modelling approach of this study is influenced in particular by two cases; that of the Pacific Oyster (Crassostrea gigas) in the Wadden Sea and the Chinese Mitten Crab (Eriocheir sinensis) in European rivers. The Pacific Oyster can affect commercial mussel yields and cause injury to recreationists due to its sharp shells (Troost, 2010). Further, the increase in substrate which may result from Pacific Oyster invasions can form a platform for the establishment of future invasions of other species (Haydar and Wolff, 2011). Barriers to spread between parts of the Wadden Sea exist due to the presence of tidal basins. Tidal basins are systems of coastal currents which form a barrier to the spread of the Pacific Oyster larvae and thus the spread of the invasion through the Wadden Sea (Kraft et al., 2010). The Chinese Mitten Crab causes damage to manmade structures such as flood defences via burrowing, damages nets and traps by feeding on the fish caught within them and increases the competition for food with native species (Herborg et al., 2003). In riparian habitats, the spread of the Chinese Mitten Crab can be impeded by installing traps at weirs (Herborg et al., 2003), although, this method is not totally effective at preventing further spread.

The two case studies considered above share a common theme: that of a barriers to spread. Barriers to spread imply that the rate at which patches are invaded is not constant. Instead it depends on the invasive population size in adjacent patches. The model therefore employs a stochastic spread process as an intuitive way to link the size of the invasion within a given patches to the probability of spread to an adjacent patch. Such a relationship can be conceptualised in two ways. Firstly, a stochastic spread processes conforms to the principle of propagule pressure, whereby the probability of a species becoming established in a new patch increases with the number of arrivals (Kolar and Lodge, 2001). Hence, a greater population in one patch leads to a great number of arrivals in a connected patch, and thus that the probability of successful establishment of the invasion in the new patch increases. Alternatively, a greater population in an invaded patch implies a greater number of possible attempts to cross the barrier, and thus a greater total probability of success.

In order to analyse optimal control of IAS with varying stock size within patches, we construct a model which is solved using Stochastic Dynamic Programming. We assume that it is always possible (if not necessarily optimal) to remove all or some of an invasion in specific patches. In practice then, the invasion can be harvested or destroyed in a given patch. We do not assume that there are always feasible methods to restrict the ability of the invasion to spread. For example, it is difficult to conceive a realistic containment technology to limit the spread of Pacific Oyster spat between tidal basins. It is however, reasonable to attempt to trap invasive Chinese Mitten Crab as they cross a weir. Therefore, unlike Epanchin-Niell and Wilen (2012), we do not assume that the spread can be prevented with certainty, rather that the probability of spread can only be reduced.

We construct a generalised model of N patches in one-dimensional space. Under the assumption that the invasion always arrives at one end of the line network, and spreads patch by patch through the network, a two-patch model is sufficient to analyse the optimality of removal, containment and combining both removal and containment. Two-patch models have been shown to provide useful insights in related settings by Salinas et al. (2005) and Sanchirico et al. (2010). We explore the effects of heterogeneity of damage costs between patches and the costs of interventions on optimal policies and thus demonstrate the added value from considering varying stock size within patches. We proceed to demonstrate how the invasion grows with, and spreads between, patches in a three-patch system under the optimal policy. This also demonstrates the generalisability of the modelling approach to larger systems.

3.2 The model

We consider the spread of an invasive species over time, indexed t, in a line network, with N patches, indexed by i. The state of the system in a given time period is described by the size of the invasion in each patch and is given by $\boldsymbol{S_t} = [s_{1,t}, s_{2,t} \dots, s_{N,t}]$. The values which $s_{i,t}$ can take (stock sizes) are determined by the set of values in the vector $\boldsymbol{Q} = [q_1, q_2, .., q_M]$ such that $s_{i,t} \in \boldsymbol{Q}\ \forall\ i, t$. The final element of $\boldsymbol{Q}$, q_M, is the maximum possible size of the stock in any given patch. Because we use a discrete approach, M gives the number of different values which stock in a given patch can take. Where j indexes the elements of $\mathbf{Q}$, the properties of $\mathbf{Q}$ are, firstly, $0 \leq q_j \leq 1$ and secondly, $q_1 = 0$. The second property means that if $s_{i,t} = q_1$ then $s_{i,t}$ is non-invaded.

The stock within a patch increases deterministically according to a vector, $\boldsymbol{G} = [g_1, g_2, \ldots, g_M]$. As described above, the state of any patch is equal to an indexed element of $\boldsymbol{Q}$. If the current stock size of a given patch is equal to the jth element of $\boldsymbol{Q}$ then the jth element of $\boldsymbol{G}$ gives stock size in the next period for that patch. To illustrate, let us consider the example of a single patch ($N = 1$) with state given by $s_t = 0.4$. Taking the example of $\boldsymbol{Q} = [0, 0.2, 0.4, 0.6, 0.8, 1]$, we see that $0.4 = q_3$. Hence, the jth element of interest is the 3rd element. If, $\boldsymbol{G} = [0, 0.4, 0.6, 0.8, 1, 1]$, then the 3rd element of $\boldsymbol{G}$ is 0.6, thus $s_{t+1} = 0.6$. The first element of $\boldsymbol{G}$ is always zero because stock sizes of zero cannot grow. The specification of $\boldsymbol{G}$ can represent various types of growth functions. In the above example, the growth rate is constant for non-zero stock sizes less than q_M. For non-constant growth rates, for example the specification $\boldsymbol{G} = [0, 0.6, 0.8, 1, 1, 1]$ could be used.

The state of invaded patches adjacent to a non-invaded patches determines the probability that the non-invaded patch will become invaded. We set the values of $\mathrm{q_j}$ such that $0 \leq q_j \leq 1$ for two reasons. Firstly, we choose the range of q_j in order to express the state of the invasion in a given patch in terms of the probability that an adjacent non-invaded patch will become invaded by a single invaded patch. Secondly, we choose the range of $\mathrm{q_j}$ such that it is never certain that a non-invaded patch adjacent to patch j will certainly become invaded (as would be the case if $q_j = 1$) or will certainly not become invaded (as would be the case if $q_j = 1$). Whether the elements of $\mathbf{Q}$ increase in constant or non-constant increments affects whether there is linearity in the relationship between the size of the invasion in a patch and the probability of invasion in an adjacent non-invaded patch. We assume a linear relationship for reasons of simplicity, i.e. $q_{j+1} - q_j = q_2 \; \forall \, j < M$. If a non-invaded patch, i, has only one invaded adjacent patch e.g. $i - 1$, then the probability that the non-invaded patch becomes invaded is given by

$$\Pr\left(s_{i,t+1} > 0 \,\middle|\, s_{i,t} = 0 \wedge s_{i-1,t} > 0 \wedge s_{i+1,t} = 0\right) = s_{i-1,t} \tag{3.1}$$

where $\Pr$ refers to the probability of a given event occurring. Hence, the probability that patch i becomes invaded in the next period, conditional upon that patch not currently being invaded, and that only one adjacent patch is invaded, is given by the stock size in the single adjacent patch. If a non-invaded patch i has two adjacent invaded patches, $i-1$ and $i+1$, then the probability that the non-invaded patch becomes invaded is given by

$$\Pr\left(s_{i,t+1} > 0 \middle| s_{i,t} = 0 \wedge s_{i-1,t} > 0 \wedge s_{i+1,t} > 0\right) = 1 - \left[\left(1 - s_{i-1,t}\right)\left(1 - s_{i+1,t}\right)\right]. \qquad (3.2)$$

Hence, non-invaded patches which are adjacent to two invaded patches are probabilistically invaded as a function of the stock size in both adjacent patches.
There are two options for controlling the invasion. The first control option is containment. Containment reduces the probability of spread to uninvaded adjacent patches without affecting the size of the invasion in patches which are already invaded. For example, barriers have been used at weirs in Germany to reduce the spread of Chinese Mitten Crab (Herborg et al., 2007). The second control option is removal. Removal reduces the stock size in a patch by a chosen amount.

The variable $\alpha_t = (0,1)$ determines whether containment has been implemented in time period t. When $\alpha_t = 1$, containment is implemented. Containment in time period t does not affect $\mathbf{S_t}$, but does affect $\mathbf{S_{t+1}}$ by reducing the probability of spread. The factor by which containment reduces the probability of spread between patches is given by ψ where $0 < \psi < 1$. Accordingly, when containment is possible, we modify Equation (3.1), which refers to the case when only one adjacent patch is invaded, in Equation (3.3).

$$\Pr\left(s_{i,t+1} > 0 \mid s_{i,t} = 0 \wedge s_{i-1,t} > 0 \wedge s_{i+1,t} = 0\right) = s_{i-1,t}(1 - \alpha_t) + \alpha_t \psi s_{i-1,t} \qquad (3.3)$$

We also modify Equation (3.2), which refers to the case when two adjacent patches are invaded, in Equation (3.4).

$$\Pr\left(s_{i,t+1} > 0 \mid s_{i,t} = 0 \wedge s_{i-1,t} > 0 \wedge s_{i+1,t} > 0\right) = (\alpha_t \psi - \alpha_t + 1)\left(s_{i-1,t} + s_{i+1,t}\right) + s_{i-1,t} s_{i+1,t} (\alpha_t \psi - \alpha_t + 1)^2 \qquad (3.4)$$

For modelling simplicity, we assume that containment cannot be targeted at specific patches and thus is either applied to all patches or no patches. Containment incurs a cost λ which is the annuity of the investment costs and the annual operation and maintenance costs incurred each year that the containment policy is enacted.

The second option for controlling the invasion is removal. Removal reduces the size of the invasion in any given patch according to the vector $\boldsymbol{K_t} = [k_{1,t}, k_{2,t}, \dots, k_{N,t}]$ where $k_{i,t} \in Q$ and $k_{i,t} \leq s_{i,t}$. The cost structure, expressed by a vector

$\boldsymbol{E_l} = [e_1, e_2, .. e_{M-1}]$, determines the incremental costs of removal (where M gives the number of different values which stock in a given patch can take). The first element, e_1, refers to the cost of removing the final unit of the invasion from a patch. The last element, e_{M-1}, refers to the cost of reducing the size of the invasion in a given patch from size q_M to size q_{M-1}. We parameterise two vectors of incremental costs $\boldsymbol{E_l}$ where $l \in (1,2)$. Both vectors entail non-linear cost structures. The cost of control increases as the size of the invasion becomes smaller. This approach follows Mehta et al. (2007), Burnett et al. (2007) and Carrasco et al. (2010b) . We assume that it is always possible to remove the last unit of invasion but that removing the last unit is more expensive than removing the first.

Specification $\boldsymbol{E_1}$ displays low levels of non-linearity and $\boldsymbol{E_2}$ displays higher levels of non-linearity. For clarity, we refer to $\boldsymbol{E_1}$ as the "flatter" incremental cost function and $\boldsymbol{E_2}$ as the "steeper" incremental cost function. We illustrate the two incremental cost functions in the following figure for the example where $M = 6$. For other values of M for a given specification, e_1 and e_{M-1} would retain the same value. Values for e_i where $1 < i < M$ are assigned via interpolation in order to retain the degree of non-linearity.

Control cost for a given $\boldsymbol{E_l}$ is given by

$$c_t = \sum_{i=1}^{N} \sum_{j=f(k_{i,t})}^{f(s_{i,t})-1} e_j. \tag{3.5}$$

The function f gives the position (value of j) of $y \in (s_{i,t}, k_{i,t})$ in vector $\mathbf{Q}$ and is given by

$$f(y) = \frac{y}{q_2} + 1 \tag{3.6}$$

Damage costs are proportional to the size of the invasion in all patches. This is the most general assumption according to Parker et al. (1999). Damage cost by patch is determined by the vector $\boldsymbol{\Gamma} = [\gamma_1, \gamma_2, \dots, \gamma_N]$ where $\gamma_i \geq 0 \ \forall\, i$ such that damage may vary between patches. Damage cost in a given period are calculated after any removal has taken place. Therefore, damage is given by

$$d_t = \boldsymbol{\Gamma}(\boldsymbol{S_t} - \boldsymbol{K_t})^T \tag{3.7}$$

where T indicates the transpose.

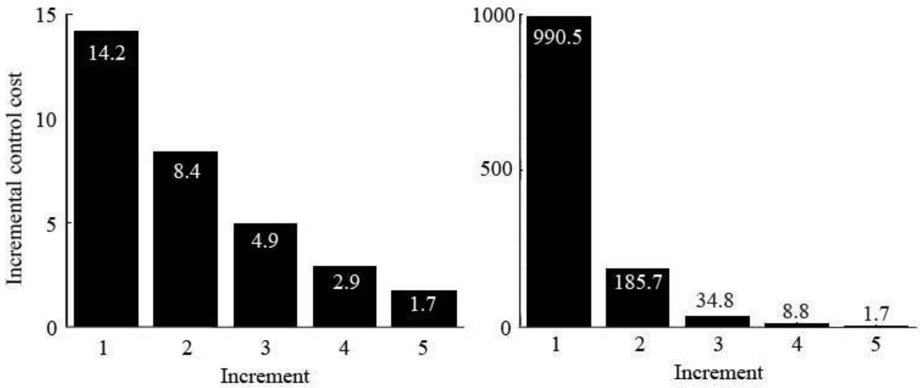

Figure 3.1: The two specifications of incremental control in a single patch. For the purposes of illustration, we take the example where $M = 6$. The first increment is the cost removing the last unit of invasion in a given patch. The last increment (5 in this example) is the cost of reducing the size of the invasion in a patch from its maximum size to its second to largest size (as determined by the vector $\mathbf{Q}$) . Incremental costs given in the figure are rounded to 1 decimal place.

The objective function minimises V_t by choosing a containment and removal policy. The policy affects the expected value of the z possible future states $\mathbf{X}_j$. The Bellman equation is therefore given by:

$$V_t(\mathbf{S}_t) = \min_{\mathbf{K}_t,\alpha_t} \left\{ \begin{array}{c} c_t(\mathbf{K}_t,\ \mathbf{S}_t\colon \mathbf{E}_l) + \boldsymbol{\Gamma}(\boldsymbol{S_t} - \boldsymbol{K_t})^T + \lambda\alpha_t \\ +\beta \sum_{j=1}^{z} Pr\big(\mathbf{S}_{t+1} = \mathbf{X}_j | \alpha_t \psi\big) V_{t+1}\big(\mathbf{X}_j\big) \end{array} \right\} \tag{3.8}$$

where β is the discount factor and λ is the cost of implementing the containment policy. Throughout the analysis we set β according to a 5% discount rate, such that $\beta \approx 0.95$.

We calculate solutions to Equation (3.8) using value function iteration (see Judd, 1998) to analyse the dynamic effects of all possible policies and find optimal policies for all states in all time periods.

The model is implemented in Matlab. The size of the state space determines the running time. Our two-patch model, used to analyse optimal interventions depending on the state of the invasion, has a smaller state space in comparison to our three-patch model, which is used to demonstrate how the invasion grows within, and spreads between, patches in a larger system. The size of the state space is given by M^N. The two-patch model uses $M = 21$ and $N = 2$ giving state space of size 441 which runs in under 5 minutes on a 2.50 GHz Intel Core i5 vPro with 4 GB of RAM. For a model calibrated to specific invasions, larger state spaces, and particularly, larger numbers of patches, would likely be required. Our three-patch model uses $M = 11$ and $N = 3$ giving state space of size 1331 which runs in 30 minutes.

3.3 Results from a two-patch model

Throughout this section, the parameterisations for $\psi, \mathbf{Q}$ and $\mathbf{G}$ remain constant. These values are given in Appendix 3.1. The values of $\mathbf{G}$ have been chosen to approximate a logistic growth function. We will begin by evaluating the case where containment is not possible for both cost specifications. We then introduce containment and analyse different containment costs. We will analyse optimal policies, which consist of one or more types of interventions. Interventions are applied for specific states of the invasion. The optimal policy determines which intervention to apply to the system for every given state. Removal can be employed as an optimal policy in several ways. These are: Immediate Eradication (IE), which removes all invasion from all patches immediately; Full Removal (FR), which removes all of the invasion from a single patch; and Partial Removal (PR), which removes of some but not all of the invasion in a given patch or patches; and No Removal (NR), which removes none of the invasion in any patch. The amount of PR within a patch is endogenously determined. We will first consider cases where containment is not possible and then cases where containment is possible. When containment is possible, containment can be used in addition to the above interventions.

3.3.1 No containment

We analyse optimal policies for the two cost specifications by first considering the case where containment is not possible.

Flatter cost specification $\mathbf{E}_1$

Olson and Roy (2002) established that optimal solutions in controlling IAS are often corner solutions. The corner solution is either No Removal (NR) or Immediate Eradication (IE). Small invasions are optimally immediately eradicated and larger invasions are often left to proliferate (i.e. nothing is done). The size of the invasion at which the optimal intervention switches from IE to NR is known as the "switching-point". In our model, the state of the system is not described by a single number, so we will refer to a switching threshold. Without the possibility for containment and using the incremental control costs given in specification E_1, Figure 3.2 shows how the switching threshold between IE or NR depends on the damage cost.

Figures 3.2(A) and 3.2(B) both show two steady states. The first is (0,0). This is the result of applying the IE intervention to all black states. The second is (0.95,0.95). This is the result of applying NR for all grey states. The border between the black (IE) and grey (NR) areas determines the switching threshold. Here, we can only see the effects of damage costs on the switching threshold. Other variables will also affect the switching threshold. The discount rate affects the present value of damages resulting from NR. The parameterisation of incremental removal costs affect the costs of IE. The damage function affects both the costs of NR and IE. We do not show analysis of the effects of the other determinants of the switching point because the general pattern seen in Figure 3.2 holds. Partial Removal (PR) is thus not optimal in this case. The convex shape of the switching threshold is due to our assumption that incremental reductions of large stocks in a given patch are less costly than incremental reductions of small stocks. This means that for a given sum of the stock across both patches, it is less costly to eradicate an invasion which is very unevenly distributed across the two patches.

We can also consider the, potentially more realistic, case of heterogeneous damage costs. In order to identify the effects of heterogeneity, we will take the damage costs from Figure 3.2(B) ($\Gamma = [4,4]$) and redistribute them between the patches. The results are shown in Figure 3.3.

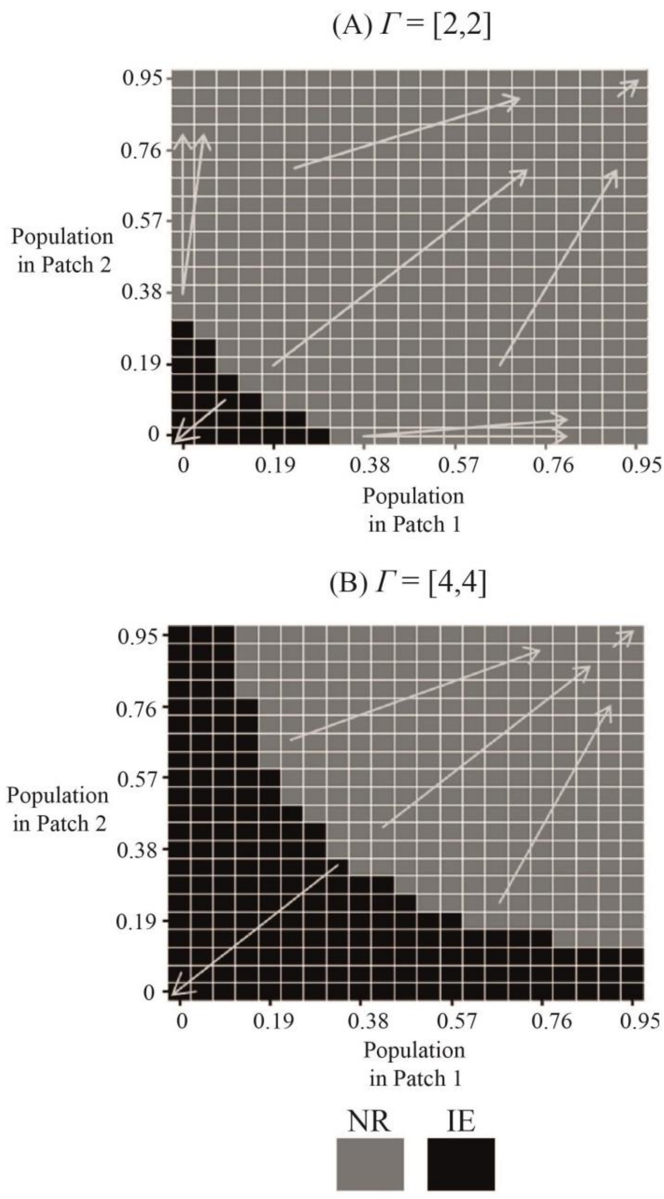

Figure 3.2: The state-dependent optimality of IE and NR for two homogenous damage cost parameterisations (given by Γ) where $\alpha_t = 0$ and $l = 1$. The arrows depict the state trajectory which results from applying the optimal intervention. In the NR area, single arrows are used to show how the invasion grows in deterministic steps to the steady state $(s_1, s_2) = (0.95, 0.95)$. Two arrows emanating from a single state show how deterministic growth increases the invasive population within an invaded patch as well as the possibility for spread to the non-invaded patch. In the IE area, removal occurs in one immediate step, with a single arrow indicating the immediate transition to $(s_1, s_2) = (0, 0)$.

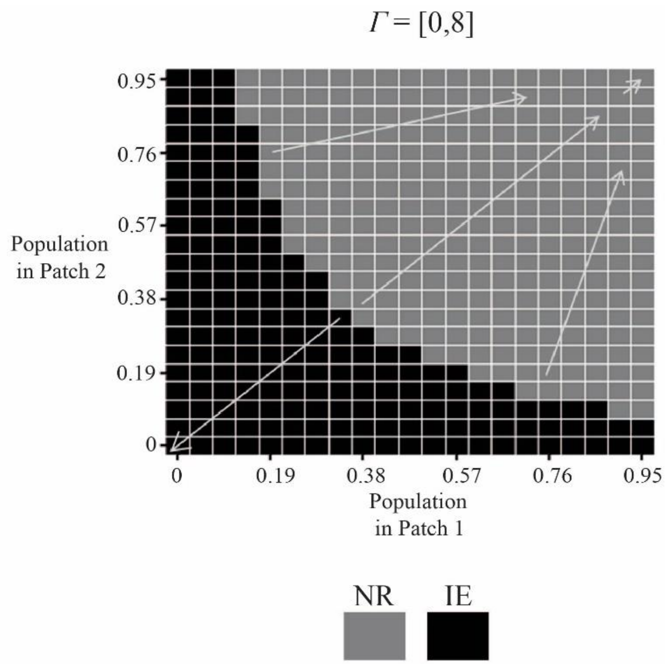

Figure 3.3: The state dependent optimality of IE or NR for a heterogenous damage cost parameterisation, $\Gamma = [0,8]$, where $\sigma_t = 0$, $l = 1$. In the NR area, multiple arrows are used to show how the invasion grows in deterministic steps to the steady state $(s_1, s_2) = (0.95, 0.95)$. In the IE area, removal occurs in one immediate step, so only 1 arrow is used to indicate the immediate transition to $(s_1, s_2) = (0, 0)$.

Redistributing the damage costs between the patches changes the switching point such that more states with greater stock sizes in Patch 2 and fewer states with lower stock sizes in Patch 1 are subject to IE. IE or NR are still the only optimal interventions. Comparing Figure 3.2(B) and Figure 3.3, shows that heterogeneity has a minor effect on the optimal policy, even in the extreme case when damage cost in one of the patches is zero. When this is the case, as in Figure 3.3, there are no direct benefits from reduced damages resulting from eradication in Patch 1. The benefits from reducing the future expected costs of damages in the other patch are sufficient to ensure that eradication of small populations in the patch with zero damage costs is still optimal.

Steeper cost specification $\mathbf{E}_2$

Interventions other than NR and IE can be optimal if the incremental cost function is steeper as in specification $\mathbf{E}_2$. We identify Partial Removal (PR) as an optimal intervention. PR is defined as the removal of some but not all of the invasion in a given patch. It does not remove all invasion in all patches. For a given state, when PR is only applied to one patch, NR is applied to the other.

Figure 3.4(A) shows a second switching threshold in addition to the switching threshold between IE and NR. The second switching threshold determines the states to which PR is applied. The application of PR maintains stock size at less than its maximum possible size and is therefore an interior solution. Hence interior solutions may be optimal if the incremental cost function is sufficiently non-linear.

Figure 3.4(B) tests the effect of redistributing damage costs between the two patches. Heterogeneity affects both the switching point between NR and IE and the switching threshold between PR and the other two interventions (NR and IE). The amount of removal under PR is also different from Panel (A). The interior solution in Panel (B) involves less removal in Patch 1 but more removal in Patch 2. Panel (B) therefore shows that heterogeneity in damage costs affects both switching thresholds. However, this effect is more marked for the threshold between NR and PR than for the switching threshold between NR and IE.

3.3.2 Introducing the possibility for containment

We now continue by adding a second type of intervention: containment. Containment limits the ability of the invasion to invade adjacent patches without affecting the size of the invasion in patches which are already invaded. Unlike in the previous section, where small changes were made to the model to identify and establish causality for the basic features of the model, this section will simply use reasonable parameters which give the most interesting results.

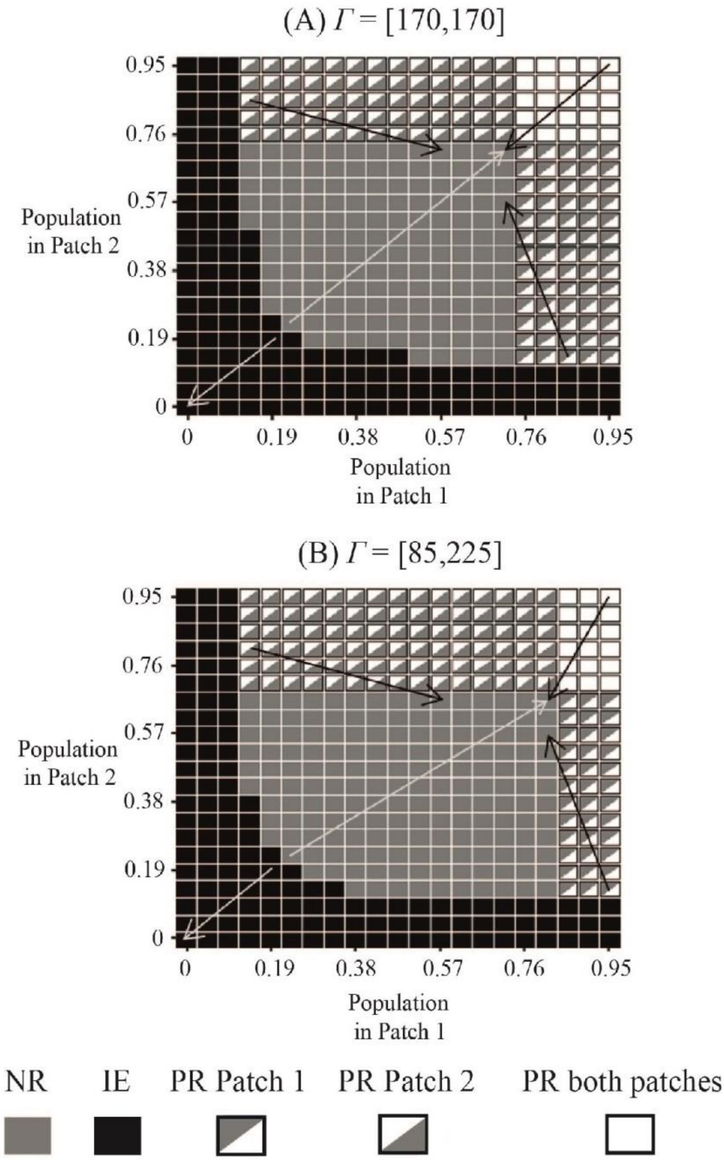

Figure 3.4: The state dependent optimality of IE, NR and PR for a heterogonous and homogeneous damage cost parameterisations where $\alpha_t = 0$ and $l = 2$. In Panel (A), PR reduces stock size to 0.713 in any patch where stock size is greater than 0.713. The steady state is therefore (0.713, 0.713) . In Panel (B), PR reduces stock size in Patch 1 to 0.808 if stock size is greater than 0.808. PR reduces stock size in patch 2 to 0.665 if the stock size is greater than 0.665. The steady state is therefore (0.808, 0.665) . The arrows depict the interventions and the location of the steady state which results from each intervention. The black arrow refers to PR.

Containment with flatter cost specification $\mathbf{E}_1$

The containment option may be used to slow down the spread of the invasion. Containment reduces the probability of spread between patches. Slowing the spread of the invasion between patches pushes damages further into the discounted future and therefore is intuitively attractive when the costs of containment are sufficiently low. This is demonstrated in Figure 3.5(A). Figure 3.5(B) demonstrates Full Removal (FR), which removes all the invasion from a single patch.

Figure 3.5(A) shows the two steady states resulting from the IE intervention (0,0) and the NR intervention (0.95, 0.95). The steady state (0.95, 0.95) also results from the patches where containment is implemented. Containment is only implemented in states which are not subject to IE and where the invasion is only found in one of the two patches. The containment option reduces the probability of spread, thus deferring damages further into the discounted future. Here, the possibility for containment has not affected the steady states.

The reduced costs of containment used in Figure 3.5(B) mean that maintaining a stochastic limit cycle using both removal and containment becomes optimal. This requires using the FR intervention. The combination of FR and containment results in two stochastic limit cycles, each of which consists of two states. The first state consists of one fully invaded patch and a non-invaded patch. The second state consists of one fully invaded patch and a patch which has been invaded in the previous time period. The first state of the stochastic limit cycle will lead to the second state according to the stochastic spread process. As soon as this occurs, FR is enacted on the newly invaded patch, thus returning the system to the first state. There are two stochastic limit cycles. The stochastic limit cycle in the bottom right of Panel (B) cycles between (0.95, 0) and (0.95, 0.0475) and the stochastic limit cycle in the top left of Panel (B) cycles between (0, 0.95) and (0.0475, 0.95).

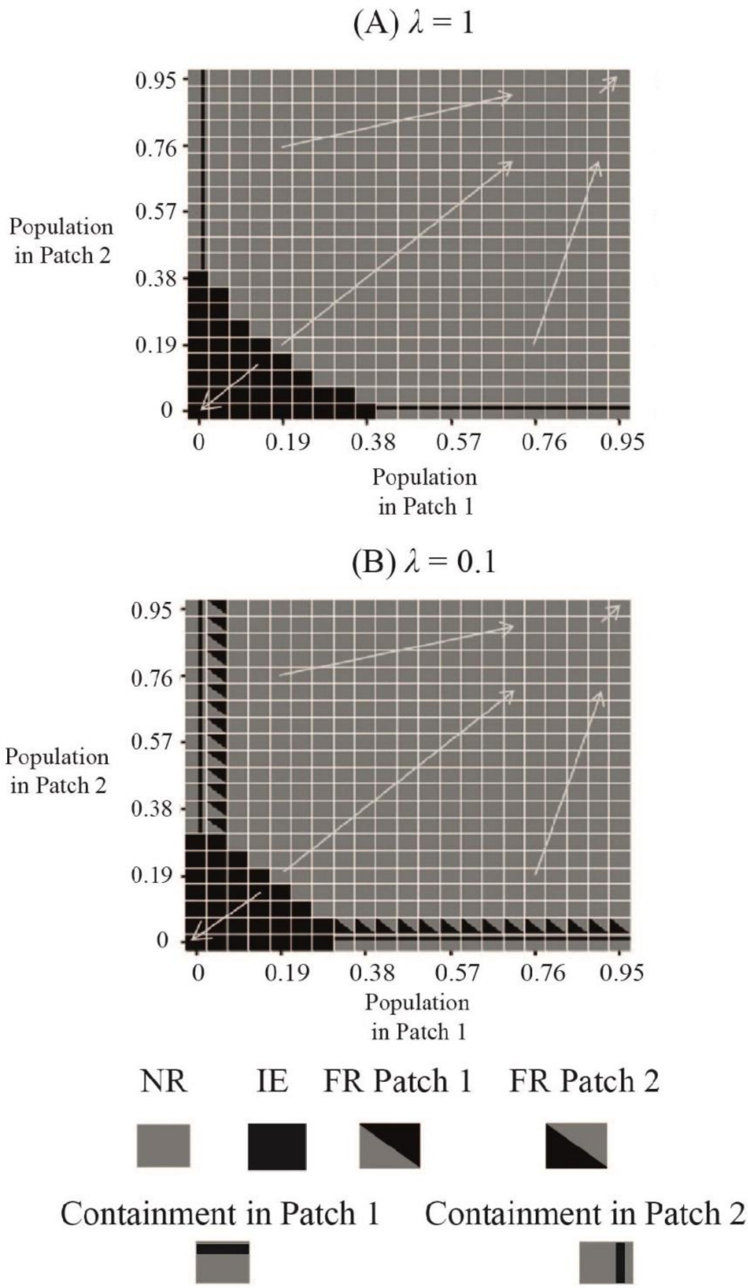

Figure 3.5: The state dependent optimality of IE, NR, FR and containment for homogeneous damage cost and varying costs of containment (given by λ) where $\alpha_t = 0$, $l = 2$ and $\Gamma = [2.4,2.4]$. FR stands for Full Removal and involves the removal of all invasion from one patch. Arrows for FR are not shown in this case to allow for ease of interpretation.

Containment with steeper cost specification $\mathbf{E}_2$

The previous section has shown how multiple types of interventions can be used to maintain a stochastic limit cycle and how containment can be used to slow the spread of the invasion. We now proceed to show that under the steeper incremental cost specification $\mathbf{E}_2$ and heterogeneous damage costs, the optimal policy for an invasion which is only found in one patch can depend heavily on the size of the population in that patch.

We will separately address three aspects of Figure 3.6. The first is seen in the bottom right of Figure 3.6. Here, containment is used for any invasion which is only in Patch 1 and is greater than 0.238. In addition, any invasions limited only to Patch 1 which are greater than 0.885 are subject to PR which reduces the stock size in Patch 1 to 0.885. This has the effect of further minimising the probability of spread. This is therefore a spread-slowing policy which is achieved by combining two types of interventions.

Eventually, the spread-slowing policy described above will lead to spread into the second patch. This leads us to the second interesting aspect. After spread into Patch 2, the invasion in Patch 2 grows until it becomes larger than 0.713 in Patch 2. At this time period, any invasion in Patch 2 above 0.713 is removed. The stock in Patch 1 will grow and remain at its maximum. In each time period then, the state of the system will change between (0.95, 0.903) and (0.95, 0.713). This is a similar effect to that seen in Figure 3.4(A) and Figure 3.4(B) whereby heterogeneity in damage costs has a large effect on the stock size in patches where the interior solution is maintained. The case of Figure 3.6 is a more extreme example because the heterogeneity results in an internal solution whereby it is only optimal to employ PR in one of the patches.

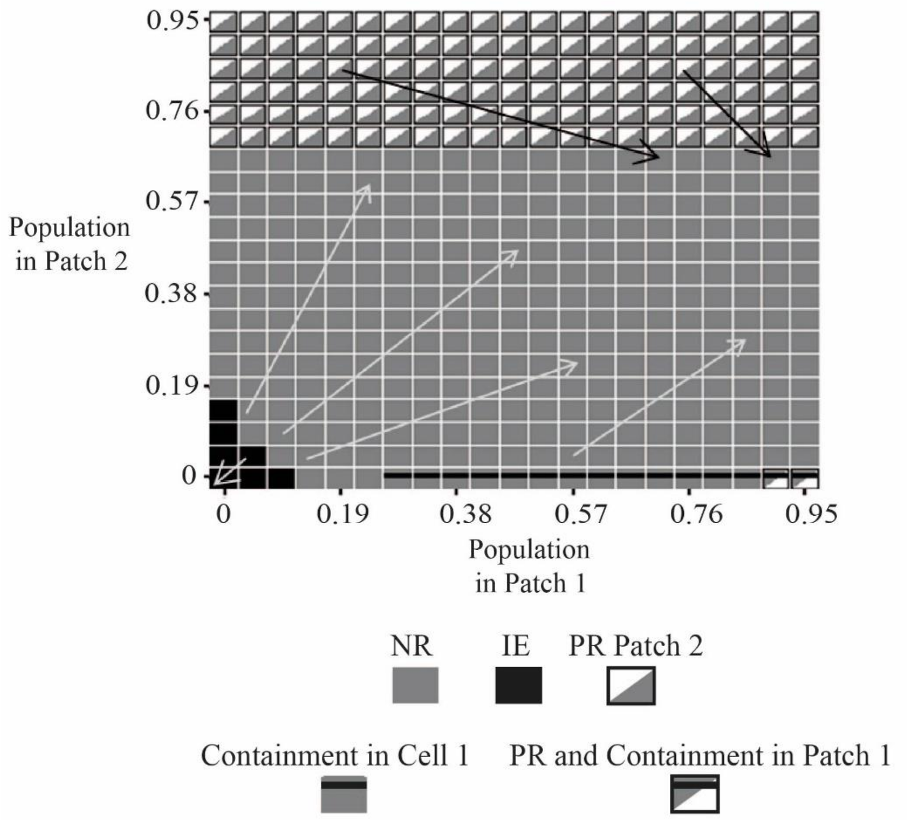

Figure 3.6: The state dependent optimality of IE, NR, PR and containment for heterogeneous damage costs $\Gamma = [5,145]$ where $l = 2$, $\alpha_t = 1$ and $\lambda = 70$. The black arrow shows the PR intervention. Arrows relating to PR in Patch 1 are omitted for ease of interpretation.

The third interesting aspect is seen in the extent in the state space over which the containment intervention is optimal. When the invasion is only found in Patch 1, invasions less than 0.1425 are eradicated, 0.1425 up to and including 0.2375 are subject to a NR intervention, and invasions greater than 0.2375 are subject to a containment intervention. The reason for this is that higher probabilities of spread without containment result in greater reductions in spread probability due to containment. Therefore, it is not necessarily optimal to implement a containment intervention in a given location if the population is not sufficiently large.

The first and third interesting aspects combined mean that there are a total of four different ways that interventions can be used to control an invasion which is found in only one patch. Small invasions are subject to IE, somewhat larger invasion are subject to NR, most invasions are subject to containment and the largest invasions are subject to both containment and PR.

3.4 Extension to three patches

This section demonstrates how the invasion grows within, and spreads between, patches in a three-patch system under the optimal policy, allowing for both containment and removal. Under a three patch system, the state is given by $\mathbf{S}_t = [s_{1,t}, s_{2,t}, s_{3,t}]$. We choose a parameterisation such that IE, NR, containment and PR are all employed depending on the state. All parameters values are provided in Appendix 3.2. Damage costs are heterogeneous between patches such that $\gamma_1 < \gamma_2 < \gamma_3$ and we employ the steeper cost specification. Due to the increased dimensionality of the three-patch case, we demonstrate the application of these interventions in a forward simulation, which shows how the state of the system changes over time, given optimal interventions and an initial state. Because invasion of a non-invaded patch is probabilistic, we present the expected manner in which the state of system will change over time i.e. the most likely transition path. Figure 3.7 shows the transition path from the smallest state which is not subject to IE. Specifically, the starting state in Figure 3.7 is $\mathbf{S}_0 = [0.19, 0, 0]$. The incrementally smaller state, $\mathbf{S}_0 = [0.095, 0, 0]$, is subject to IE such that the invasion is present only at time $t = 0$. Therefore, this scenario is not shown in Figure 3.7.

Figure 3.7 shows results which follow directly from the results for a two-patch system. As in the two-patch case, states with a smaller population are subject to IE, but this intervention is not optimal for large states, such as the initial state in Figure 3.7. Containment is employed only when the invasion is confined to Patch 1 because there are two uninvaded patches and as such, the reduced future damages resulting from reducing the probability of spread to Patch 2 are greatest. This is similar to Figure 3.5(A). If there are more patches in the system, there can be more non-invaded patches. In general, it holds that the greater the number of non-invaded patches, the greater the future damages which can be avoided by reducing the probability of spread between patches and hence, the optimality of containment increase. In Figure 3.7, once Patch 2 becomes invaded, containment is no longer optimal. It is however optimal to employ PR to limit the size of the invasion in each patch in the steady state, which is reached in the eighth period. This is a similar result to Figure 3.4(B), because the steady state in each patch varies depending on the damage cost. Just like in Figure 3.4(A), if the damage cost is homogenous between patches, the steady state population size within patches would be equal.

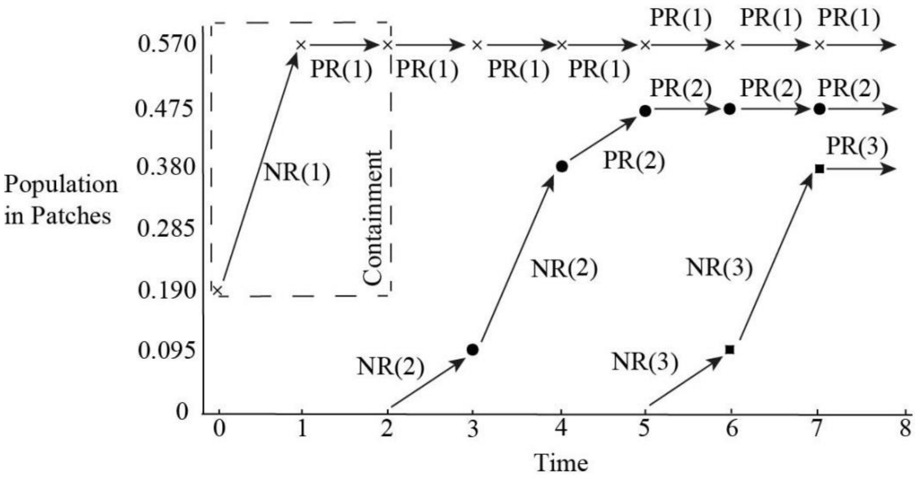

Figure 3.7: The expected growth and spread of an invasion in a three-patch system, with parameters given in Appendix 3.2. The initial state is $\mathbf{S}_0 = [0.190, 0, 0]$. A smaller initial state size $\mathbf{S}_0 = [0.095, 0, 0]$ is subject to IE (not shown). The state of the system is given by the points indicated by a cross (Patch 1), a circle (Patch 2) and a square (Patch 3). For clarity, arrows link the population size in each patch over time. Optimal interventions by patch are indicated by, for example, NR(1), meaning No Removal in Patch 1, or PR(3) meaning Partial Removal in Patch 3. The population within each patch converges to a steady state in the eigth time period. The population sizes in the steady state are maintained by PR at a level which varies according to the varying damage cost between patches.

3.5. Discussion and conclusion

This study explored a spatially explicit dynamic progamming model for the optimal control of stochastically spreading invasive species. The model is relevant for the optimal control of invasions spreading in coastal and riparian habitats. Optimal management interventions are derived in a two-patch model and the generalisability of the model is demonstrated in a three-patch setup, which shows how the invasion grows and spreads under the optimal policy in a larger system. The insights provided by the model result principally from allowing for varying stock sizes within patches. This study focuses on optimal management once the invasion has arrived and therefore does not consider policies aimed at preventing the arrival of the invasive species in the system in the first place.

This study provides new insights into the optimal control of invasive species. This is best demonstrated with respect to Epanchin-Niell and Wilen (2012), who consider two-dimensional space with a deterministic spread process, patches which are either invaded or not invaded and allow for both removal and containment. Firstly, by allowing for varying stock sizes within patches, our model facilitates the optimization of the timing of containment. In Epanchin-Niell and Wilen (2012), if it is optimal to apply containment, then containment must be applied as soon as a patch becomes invaded. We have shown that this is not neccessarily optimal. Postponing containment for a given patch until the population within that patch has reached a certain size may be optimal if the reduction in the probability of spread resulting from the containment strategy increases as the invasive population size within that patch increases. In this case then, and borrowing the terminology of Buhle et al. (2005), obtaining the most "bang for the buck" from containment can only be achieved via optimally timing its implementation.

Secondly, we identify optimal policies in the (often more realistic) case where containment can only reduce the probability of spread, not prevent it entirely. This is in contrast to Epanchin-Niell and Wilen (2012), who assume that containment is always perfectly effective. When containment reduces the probability of spread in a model with varying stock size in patches, several novel and interesting policies may result. Preventing spread entirely may still be optimal, even if containment alone cannot achieve this. Thus, policies to stop the spread can still be optimal, but require a combination of containment and removal of any new spread.

When containment alone cannot fully prevent the spread of the invasion, it is obvious that slowing the spread is a potentially optimal strategy. However, and thirdly, we identify the possibility for complex optimal policies for slowing the spread which depend on the stock size in the invaded patch. Allowing for varying stock sizes has shown that up to four disctinct interventions may be applicable depending on the population size. These vary from Immediate Eradication for small populations, to No Removal or containment for intermediate populations and finally a combination of containment and Partial Removal for larger stock sizes.

These results can be evaluated with respect to the literature on the optimal control of epidemics in metapopulations, for which two-patch models have been employed (e.g. Mbah and Gilligan, 2011; Alpízar and Gordillo, 2013 and most notably by Rowthorn et al., 2009). Rowthorn et al. (2009) consider two populations of individuals who are susceptible to infection by a disease in an optimal control model which aims to minimise the number of infected individuals. The authors find that treatment should be directed at the population with the lower level of infection. In order to compare results between this study and Rowthorn et al. (2009), let us identify which of our results are derived in a situation most comparable to that considered in Rowthorn et al. We find that comparisons to our results can best be made where damage cost is homogenous, incremental costs are flatter and containment is not possible, i.e. Figure 3.2. Figure 3.2 shows that it is optimal to immediately eradicate the invasion in both patches or to do nothing in both patches for all possible states. Hence, unlike Rowthorn et al., we find that it is not optimal to focus on the patch with the lowest stock size. The principal reason for this difference is that Rowthorn et al.'s result is derived in the case where the budget constraint is binding. There is no budget constraint in our study. When the budget constraint is non-binding, Rowthorn et al.'s results are more similar to those in our study. However, if the budget constraint is binding, then some method of prioritisation is required. In Rowthorn et al. (2009), the population with the greatest number of suseptible individuals is prioritised.

Could we then expect a similar result if a budget constraint was applied in our model? The degree to which the results will be similar depends on how steep incremental costs are. If incremental costs are linear, as in Rowthorn et al., then a binding budget constraint would prioritise patches with smaller stock sizes. As incremental costs become steeper, however, this effect would reduce, because incremental reductions in stock of the invasive species are more costly in a patch with a smaller stock. This will, to some extent, offset the benefits of prioritising a patch with the largest potential for growth in the stock size, i.e. a patch with the smallest stock size. Overall then, comparison of our study to the literature on the optimal control of epidemics in metapopulations reveals the importance of assumptions regarding the existence of a budget constraint, and that the implications of such assumptions depend heavily of assumptions regarding the nature of incremental (marginal) cost of removal (treatment).

Applications of this modelling framework need to account for two principle constraints. The first is the availability of parameters. This can be achieved using statistical habitat suitability models (Sadeghi et al., 2014) as correlates for invasibility among patches. Economic data is also required, which can be partially provided by valuation studies, as in, for example Nunes and Markandya (2008). The second is computational burden. This can be addressed by trading off the number of patches against the number of different values which the invasive population in a given patch can take (i.e. the number of elements in the vector $\mathbf{Q}$). If the system in question cannot be accurately modelled within reasonable computation time, then more advanced and efficient computational approaches, such as constraint integer programming (Achterberg, 2009, as in Epanchin-Niell and Wilen, 2012) could be employed.

In conclusion, the results show that allowing for varying population sizes within patches facilitates more accurate optimal policy prescriptions. Future research could focus on parameterizing the existing case studies for stochastically spreading species while extending the range of possible management options, for example to include the use of biological control agents (Impson et al., 2004).

Appendix 3.1 Parameterisations for Section 3.3

$\psi = 0.1$

$$\mathbf{Q} = [0, 0.0475, 0.095, 0.1425, 0.19, 0.2375, 0.285, 0.3325, 0.38, 0.4275, 0.475, 0.5225, 0.57, 0.6175, 0.665, 0.7125, 0.76, 0.8075, 0.855, 0.9025, 0.95]$$

$\mathbf{G} = [0,\ 0.048,\ 0.238\ ,\ 0.475,\ 0.618,\ 0.713,\ 0.76,\ 0.808,\ 0.808,\ 0.808,\ 0.855,\ 0.855,\ 0.855,\ 0.855,\ 0.903,\ 0.903,\ 0.903,\ 0.903,\ 0.903,\ 0.903,\ 0.95,\ 0.95]$

Appendix 3.2 Parameterisations for Section 3.4

$\mathbf{Q} = [0, 0.095, 0.19, 0.285, 0.38, 0.475, 0.57, 0.665, 0.76, 0.855, 0.95]$

$\mathbf{G} = [0, 0.38, 0.57, 0.665, 0.76, 0.85, 0.85, 0.85, 0.95, 0.95, 0.95]$

$\psi = 0.8$

$\lambda = 600$

Part B: International fisheries agreements

4 Farsightedness, changing stock location and the stability of international fisheries agreements

ABSTRACT

Changes in stock location may affect the stability of international fisheries agreements. This paper offers a theoretical analysis of the stability of Regional Fisheries Management Organisations (RFMOs) in a non-cooperative, coalition formation game based on the classic Gordon-Schaefer model. We employ a new stability concept which modifies Farsighted Stability (Chwe, 1994). We call this concept Farsighted Downwards Stability (FDS). We also employ the internal stability concept for comparison. Analytical results regarding FDS for symmetric players without changing stock location show stable Grand Coalitions for $n \leq 4$ player games and the possibility for partial cooperation. Sensitivity analysis deals with changing stock location and cost asymmetry. Stability decreases in n, increases when costs are asymmetric and increases when FDS is employed. Farsighted conjectures on behalf of RFMO members can thus help to maintain cooperation as stock location changes. However, FDS is more sensitive to changes in stock location than internal stability.

This chapter is based on the paper: Walker, A.N., Weikard, H-P., 2014. Farsightedness, changing stock location and the stability of international fisheries agreements Environmental and Resource Economics (in press). DOI 10.1007/s10640-014-9853-1

4.1 Introduction

There is a general recognition that cooperation is needed for the management of international fisheries to ensure the sustainability of stocks. With this in mind, Regional Fisheries Management Organisations (RFMOs) were set up to facilitate cooperation . The need for cooperation has sparked a recent literature concerned with the potential for, and the stability of, such agreements (e.g., Kaitala and Lindroos, 1998; Bjørndal et al., 2000; Lindroos, 2008, Pintassilgo and Lindroos, 2008; Pintassilgo et al., 2010; Breton and Keoula, 2012; Rettieva, 2012; Punt et al., 2012; Bjørndal and Lindroos, 2012). Such research is especially important because RFMO agreements are not binding or enforceable (Bjørndal et al., 2000). Much of the literature rightfully focuses on the potential for cooperation and the new-member problem (for a summary see Bailey et al., 2010).

In addition to these established research lines, recent research has begun to focus on the issue of changes in stock location which is likely due to climate change (Cheung et al., 2009). For example, mackerel stocks in the North East Atlantic have recently shifted northwards (Jansen and Gislason, 2011). This has led to unilateral setting of national fishing quotas which constitute a violation of the existing RFMO agreement (Arnason, 2012; Haraldsson and Carey, 2011). Ellefsen (2012) studies this specific problem with a calibrated model to assess the effects on the stability of the RFMO after the entrance of Iceland into the game. In general, uncertainty regarding the effects of climate change on stocks and the inflexibility of agreements to changes in stock locations have also been shown to be a significant barrier to maintaining cooperative agreements (Miller and Munro, 2004). Further motivation is provided by Munro (2008) who calls for more applied game theoretic research on this issue.

In addition to Ellefsen (2012), three other studies address changes in stock location. Ekerhovd (2010) is concerned with both the area which is under RFMO management and the Exclusive Economic Zones (EEZs) of given countries, wherein those countries have exclusive fishing rights. Ekerhovd (2010) considers changes in the shares of stock of blue whiting distributed between the high seas and EEZs and shows whether or not coalitions are stable. The scenarios of changes in stock location have a strong impact on the stability of coalitions. Further, Brandt and Kronbak (2010) consider the case of cod in the Baltic under IPCC climate change scenarios. They analyse the size of the possible set of cooperative agreements under changes to recruitment and size of the stock. They conclude that cooperative solutions are less likely under changes in stock location. The most recent work, by Ishimura et al. (2012) has been concerned with the Pacific sardine under climate variability and its exploitation by Mexico, Canada and the USA. Different cooperative and non-cooperative regimes are analysed. They conclude that unilateral efforts to maximise conservation and management benefits would not be successful under climate change. The stability of the different cooperative and non-cooperative regimes is, however, not analysed.

This paper conceptualises the stability of cooperation under changes in stock location and hence adopts a different approach from Brandt and Kronbak (2010), Ekerhovd (2010), Ishimura et al. (2012) and Ellefsen (2012). Cooperation is most beneficial when the RFMO is a "Grand Coalition" consisting of all nations with a genuine interest in a given stock. Accordingly, we examine the stability of Grand Coalitions for a fixed number of players under changes in stock location.

Changes in stock location can be included in a model by allowing for changes in the "catchability" (usually denoted by q in the standard Gordon-Schaefer model). Catchability is normally considered to represent the fishing technology and thus the productivity of fishing effort. As a stock of a constant size changes its position relative to the fishing harbours of different countries, we can consider their productivity of effort as changing. This would be due to changing sailing time before reaching fishing grounds or an increased concentration of fish in proximity to the harbour. We assume that the productivity of fishing effort is determined only by the stock location and therefore that fishing technology is identical across states. This approach is most suitable for high seas fisheries where biological change does not affect the spatial distribution of stock across EEZs. We also assume that climate change, while it affects location, does not affect other aspects of the biology of the stock.

In addition to addressing the theory of changes in stock location, and in order to address the question of how fully cooperative agreements can be stabilised, we use two solution concepts. First, we employ a variant of the farsightedness concept which is based on farsighted conjectures (Chwe, 1994). Farsighted conjectures are used in the context of a Great Fish War by Breton and Keoula (2012). In comparison to Nash conjectures, farsighted conjectures do not restrict players to remain in the coalition structure resulting from the deviation of one player. Farsighted conjectures therefore allow players to respond to deviations by making further deviations. Second, we employ the internal stability solution concept. This is based on Nash conjectures and is used most frequently in the literature. Nash conjectures do restrict players to remain in the coalition structure resulting from the deviation of another player.

This paper analyses implications of changes in stock location for the stability of Grand Coalitions under these different solution concepts. Comparing results under different solution concepts allows us to analyse the degree to which the (credible) responses to deviations, as conjectured by farsighted players, can affect the stability of Grand Coalitions under changing stock location.

Our study uses analytics to explore the characteristics of a farsighted solution concept in the symmetric setting and to derive some basic results in the asymmetric setting. A more detailed analysis of the asymmetric case is achieved via sensitivity analyses, which allow us to draw conclusions about the effects of asymmetry and changes in stock location on internal and farsighted stability in 3 and 4-player games.

This chapter makes three contributions to the literature. Firstly, we broaden the literature on asymmetric fishing games by comprehensively analysing the effects of asymmetric catchability. This builds on work by Pintassilgo et al. (2010). In turn, and secondly, this allows us to produce a theoretical framework to analyse the effects of changes in stock location on the stability of cooperation. Thirdly, and as will become clear later in this chapter, we develop a modified solution concept based on farsighted stability which addresses the problem of myopia while also being applicable in asymmetric coalition formation games which use sharing rules. We now continue into our model and analysis.

4.2 The bioeconomic model

The set of N players represents n different fishing nations i who choose effort e_i; $E = (e_1, \dots, e_n)$. We restrict effort such that $e_i \in \mathbb{R}_0^+$. Efforts affect harvests h_i and, in turn, profits Π_i. We employ the Gordon-Schaefer model of fisheries which has a long tradition in the literature. A single commercial fish stock is given as x. Stock grows according to

$$g(x) = rx\left(1 - \frac{x}{k}\right). \tag{4.1}$$

Here, $r > 0$ refers to the intrinsic growth rate of the stock and k is the carrying capacity of the ecosystem. The production function (harvest) is given by

$$h_i = q_i e_i x. \tag{4.2}$$

Here, $0 < q_i \leq 1$ is the catchability coefficient which we use to represent changes in stock location. Unlike most studies, we allow catchability to vary between players and therefore become a source of potential asymmetry in the model.

This paper analyses the steady state where growth (Equation (4.1)) is equal to total harvest, $\sum_{i=1}^{N} h_i$. This allows us to determine the steady state stock as a function of efforts and obtain

$$x = k - \frac{k}{r}\sum_{i=1}^{N} q_i e_i. \tag{4.3}$$

Fish is sold on a common market and profit is given by

$$\Pi_i = pq_i e_i x - c_i e_i, \tag{4.4}$$

where p is price and c_i is i's unit cost of effort. Costs may differ between players. This bio-economic model is used to calculate profits for any vector of efforts $(e_1, \dots, e_n)$.

4.3 The fisheries game

Because there can only be one RFMO for a given fish stock, we model RFMO stability as a cartel game. We examine the incentives to participate in an RFMO in a two stage game. In the first stage, players' strategy space is $\{join, not\ join\}$ and this determines their RFMO membership. A coalition structure is denoted by the set S of players who join where $|S| = s$. The set $N \backslash S$ contains $n - s$ singletons who do not join the RFMO. We have a Grand Coalition when $S = N$, i.e. where all players are in the RFMO. Given a coalition structure, players choose their effort levels in the second stage.

4.3.1 Choosing effort levels

Effort levels are chosen to maximise profits in a strategic setting. Coalition members cooperate by choosing effort levels to maximise joint profits. Effort is a function of the efficiency of players. We define inverse efficiency as $b_i \equiv \frac{c_i}{pq_ik}$. Further, we define $\gamma_i \equiv \frac{c_i}{q_i}$, which we term the cost-catchability ratio of a given player. The term γ_i thus denotes the cost of fishing effort adjusted for the catchability and contains all the terms of b_i which we allow to be asymmetric. Furthermore, let $l \in S$ be the member with the lowest cost-catchability ratio such that $\gamma_l \equiv \min_{i \in S} \gamma_i$. Under these definitions, the following holds:

Lemma 1.1*: Under a common market, the only coalition member whose effort is non-zero is the member with the lowest cost-catchability ratio,* $\gamma_l \equiv \min_{i \in S} \gamma_i \; \forall \; S \subseteq N$.

Players can only have a relative advantage via the individual parameters, c_i and q_i. Therefore, the player with the lowest cost-catchability ratio must be the most efficient fisher. Coalition members cooperate to maximise joint profits and therefore the most efficient fisher will assume the task of fishing for the coalition. In this way, it is always efficient for player *l* in a coalition to fish since we have assumed, for simplicity, that cost is linear in effort and therefore marginal and average costs are also constant. Non-linear costs would usually merit multiple active fishers in the coalition and Lemma 1.1 would no longer apply.

We introduce transfers between coalition members to compensate members with zero fishing effort under Lemma 1.1 and thus incentivising membership. Transfers allow the profit of player l to be shared among the members. Transfers (or “side” payments) have met much resistance in the policy world and are not implemented in direct financial terms (Munro, 2008). However, transfers are implicit in various policy instruments. Transfers can be made through bargaining over catch shares for other commercial species within an RFMO or with Individual Tradable Quotas (ITQs). Selling ITQs to the most effective member constitutes a transfer.

Given Lemma 1.1 and the choices of each player to join or not join, we can, via reaction functions, provide equilibrium effort strategies for coalition members and non-members. The reaction function and equilibrium strategy for the All Singletons structure is derived in Appendix 4.1. Because Lemma 1.1 holds for all coalition structures, the effort levels determined in the second stage for a game with coalition S will be the same as the efforts levels seen in an All Singletons structure consisting of $n - s + 1$ players.

The reaction function of a singleton in the All Singletons and partial cooperative structures is given by

$$e_i = \frac{r}{2q_i}(1 - b_i) - \frac{1}{2q_i}\sum_{j\in N\setminus\{i\}} q_j e_j. \tag{4.5}$$

The reaction function for the coalition in partially cooperative and Grand Coalition structures is given by

$$e_l = \frac{r}{2q_l}(1 - b_l) - \frac{1}{2q_l}\sum_{k\in N\setminus\{S\}} q_k e_k. \tag{4.6}$$

The equilibrium strategies for the Grand Coalition and both the coalition and singletons in partial cooperation structures can be expressed in one equation, namely,

$$e_i = \frac{(n-s+1)r}{(n-s+2)q_i}(1 - b_i) - \frac{r}{(n-s+2)q_i}\sum_{j\in((N\setminus S)\cup\{l\})\{i\}}\left(1 - b_j\right) \quad \text{for} \quad i \in (N\setminus S) \cup \{l\} \tag{4.7}$$

and $e_i = 0$ for $i \in S\setminus\{l\}$.

Equation (4.7) also represents the equilibrium strategy for the All Singletons structure in the special case that the coalition consists of only one player, i.e. $S = \{l\}$ such that $\left((N\setminus S) \cup \{l\}\right) = N$.

Even before searching for solutions to the game, the equilibrium strategies permit insights into the presence and nature of the externalities in the model. Equations (4.2) and (4.3) show how harvest is a function of stock and therefore the harvest of one player will negatively affect other players because less fish can be caught with the same effort. This negative externality offers scope for beneficial cooperation. There is no competition when only one player in a Grand Coalition fishes. This allows the player with the lowest cost-catchability ratio to maximise the profit for the whole coalition by fishing from a large stock.

Equilibrium efforts, calculated in Equation (4.7), can then be substituted into the profit function (Equation (4.4)) to obtain the partition function $V(S)$ which gives payoffs as a function of the coalition structure. The partition function is the basis for the following section where we introduce two stability concepts.

4.3.2 Stability, solution concepts and sharing rules

Coalition stability depends on how much profit a coalition generates and how that profit is shared. A cartel partition function gives the profit of the coalition and every singleton. The coalition profit is then shared between coalition members via a sharing rule, which determines the payoffs. We will first provide a general definition of a sharing rule and then the particulars of the sharing rule for the two solution concepts.

We use a sharing rule which maximises potential for cooperation, namely the "almost ideal sharing scheme" proposed by Eyckmans and Finus (2004), McGinty (2007) and Weikard (2009). This sharing scheme uses "outside options" to determine how surplus is shared. Outside options are defined as the payoff that a player will receive when he leaves a coalition. The sharing scheme demands that every player receives the value of his outside option ω_i plus a share $\lambda_i(S)$ of the surplus that the coalition generates in excess the sum of the values of the outside options $V_S(S) - \sum_{i\in S}\omega_i$ such that the payoff $V_i(S)$ of a coalition member in S is given by

$$V_i(S) = \omega_i + \lambda_i(S)\left[V_S(S) - \sum_{j\in S}\omega_j\right] \tag{4.8}$$

where $\sum_{i\in N}\lambda_i(S) = 1$ and $\lambda_i(S) \geq 0.$

In our numerical analysis, we will use coalition surplus as a measure of coalition stability. A positive (negative) surplus implies a stable (unstable) coalition. Our measure of stability is therefore defined as

$$Y_S \equiv V_S(S) - \sum_{i\in S}\omega_i. \tag{4.9}$$

How the outside option is defined depends on the stability concept used. We consider two stability concepts: Nash stability and a modified farsighted stability concept. A stability concept stipulates whether or not a player will deviate from a given coalition. A player's decision regarding deviation depends on the type of "conjecture" which is employed.

The first type of conjectures are Nash conjectures. This assumes that all players will remain in the coalition structure which directly results from a deviation. Other players may adjust their efforts but no player will enact further deviations. Therefore, for Nash stability, $\omega_i = V_i(S\backslash\{i\})$. The following inequality is a necessary condition for the Nash stability of a coalition:

$$V_i(S) \geq V_i\ (S\backslash\{i\}) \quad \forall\, i \in S. \tag{4.10}$$

The general formulation for Nash stability is that both internal stability (where no member wants to leave) and external stability (where no singleton wants to join) must hold. Here, we are concerned with the Nash stability of the Grand Coalition which cannot be enlarged and hence cannot be externally unstable. Internal stability is therefore a sufficient condition for a stable Grand Coalition.

The second type of conjecture is based on farsightedness. Farsighted conjectures do not assume that players will remain in structures imposed upon them by a deviation. Should further deviation from such structures be beneficial, then players will deviate. Whether further deviations are beneficial is based on farsighted conjectures developed by Chwe (1994).

Farsighted conjectures require a different definition of the outside option. Specifically, we introduce the Farsighted Downwards Stability (FDS) concept. This concept is a restricted version of Farsighted Equilibrium (Chwe, 1994). FDS is a pragmatic solution concept which restricts Farsighted Equilibrium such that it can be operationalized in a game with transfers and asymmetric players. We now define FDS via the concepts of ordered sequences and credible induction.

Definition 4.1.1: *A strictly ordered sequence is defined as a vector of coalition structures* $(S_1, S_2, \dots S_k)$ *which are ordered such that* $S_1 \supset S_2 \supset \cdots \supset S_k$ *where* $|S_j| = |S_{j+1}| + 1\ \ \forall\, j < k$.

Definition 4.1.2: *A coalition* S_k *can be credibly induced via an ordered sequence iff* $\forall\, S_j \in (S_1, S_2, \ldots S_{k-1})$, *there exists a player* $i \in S_j$ *such that* $V_i(S_k) \geq V_i(S_j)$ *and* $\forall\, S_m \subset S_k$, *there does not exist an ordered sequence* $(S_k, \ldots, S_m)$ *such that* $\forall\, S_j \in (S_k, \ldots, S_m)$ *there is an* $i \in S_j$ *such that* $V_i(S_m) \geq V_i(S_j)$.

Definition 4.1.3: *A coalition S satisfies FDS iff there does not exist a coalition* $S_k \subset S$ *which can be credibly induced from S.*

Intuitively then, we consider sequences of deviations from a coalition where one player after another deviates. Deviations by particular players are credible only if the payoff in the structure at the end of the sequence of deviations provides a greater payoff for every deviator. Therefore, by construction, the All Singletons structure always satisfies FDS. Structures satisfying FDS are those from which outside options are derived. These are used to calculate coalition stability as in Equation (4.9).

Definition 4.1.1 refers to sequences of deviations. Sequences allow players to "induce" certain structures (Definition 4.1.2). If a player deviates from a Grand Coalition, the payoff which results directly from that deviation alone may be very large. However, if the new coalition is not FDS, the initial deviation would induce further deviations by the remaining coalition members until an FDS structure is reached. In the FDS structure, the player who deviated first may receive a payoff lower than what he received in the Grand Coalition. As such, FDS addresses the problem of myopia in Nash conjectures (Harsanyi, 1974) and, can potentially result in a larger set of stable coalitions than under internal stability.

Although the FDS concept may not be behaviourally convincing in all settings, we would argue that it applies in our case where the main concern is the stability of the Grand Coalition. Implicitly, the FDS concept implies a punishment strategy whereby players who have deviated from the Grand Coalition are not allowed to benefit from re-joining. The FDS concept thus reflects a plausible restriction on the action space of players.

The FDS solution concept also has several pragmatic advantages in games with transfers and asymmetric players. These advantages result directly from the exclusion of external stability considerations. If coalitions in asymmetric games can be externally unstable, then a stable structure from which to draw the outside option may not exist. Furthermore, one of the structures that could be reached is the Grand Coalition itself. This is due to the potential for cycles in coalition structures.

To illustrate how cycles can occur, consider $n = 3$ and the Grand Coalition {1, 2, 3}. Player 1 considers what structure his deviation from the Grand Coalition would eventually induce in order to decide if his initial deviation is worthwhile. His initial deviation would lead to coalition {2,3}. Suppose he knows that partial coalition {2,3} is internally unstable and that player 2 will deviate, resulting in All Singletons. He also knows that player 3 prefers {1,3} to All Singletons (it is externally unstable) so he knows that coalition structure {1,3} will form. Finally, he knows too that {1,3} is externally unstable so player 2 will join the coalition. This brings us back to the Grand Coalition, and thus results in a cycle. This process is summarised in Figure 4.1. Note also that cycles cannot occur if players are symmetric. The presence of cycles mean that there is no stable structure from which player 1 can draw his outside option to decide whether his deviation from the Grand Coalition is beneficial. A further practical implication is that if outside options are not calculable, then optimal sharing cannot be implemented.[4]

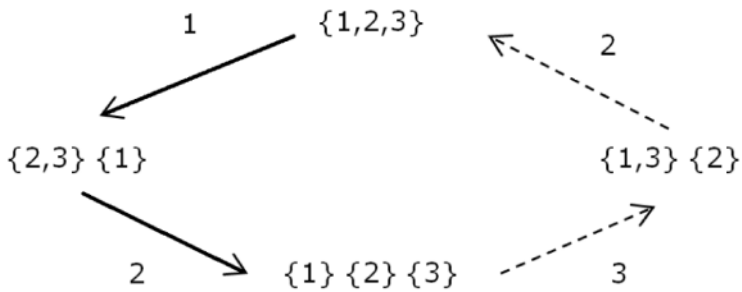

Figure 4.1: The potential for cycles. The dashed arrows indicate moves that are ruled out under FDS and thus how cycles are prevented.

The issue of external stability, asymmetric players and farsightedness has been addressed by Caparrós and Giraud-Heraud (2011). They suggest an alternative definition of external stability such that a coalition is externally stable if the addition of a player to that coalition would lead to an internally unstable coalition. Such an approach combined with optimal sharing rules would, however, not preclude the possibility of cycles.

[4] Consider also that in this example, the set of imputations whose values could be considered to inform the outside option includes the Grand Coalition itself. Needing to know V_i when we need to know V_i in order to know V_i is a paradox best avoided.

The problem of cycles and the need to define outside options is avoided if there are no transfers in the game. Transfers are not used in IEA analyses such as de Zeeuw (2008), Osmani and Tol (2009) and Biancardi and Di Liddo (2010). In these examples, players receive the benefits of cooperation directly. In our game, a requirement for stable coalitions is that benefits are realised by the most efficient player and then distributed to coalition members.

In our specific circumstances, one option for dealing with this problem is to limit the information required for the sharing rule for a given structure to that which can be provided purely by the imputation for that particular structure. Various methods to achieve this such as the nucleolus and the Shapley value are found in cooperative game theory. However, optimal sharing rules perform best to stabilize coalitions for the provision of public goods (McGinty et al., 2012). Therefore, employing methods such as the Shapley value under any solution concept would lead to a reduction in the stability of Grand Coalitions (McGinty et al., 2012). Hannesson (2011) argues that non-cooperative approaches are too pessimistic regarding the potential for collaboration. Sharing rules which are not "optimal" are therefore undesirable.

In games with asymmetric players and transfer, the FDS concept therefore represents a plausible restriction in the action space of players, prevents possible cycles and permits a consistent application of optimal sharing rules. Additionally, if Grand Coalition stability under changing stock location is improved when the FDS solution concept is employed, we can suggest that restricting the action space of players as implied by the FDS concept could be beneficial for ensuring stability. Later, we will return to the FDS concept in order to show how it can be applied to asymmetric players using computational methods.

4.3.3 Some established results

Before we continue with our analysis, we briefly review some established results using internal stability. Increasing the number of players leads to reduced internal stability of the Grand Coalition (Pintassilgo and Lindroos, 2008). Internal Stability of the Grand Coalition can however be achieved by introducing cost asymmetry into the model (Lindroos, 2008). Cost asymmetry increases the relative efficiency at which a coalition can fish. In coalitions, the most efficient player fishes. Should this most efficient player have a sufficiently large advantage, the removal of externalities resulting from coalition formation allows the most efficient player to fully exploit his advantage to the extent that he can compensate those in the coalition enough to prevent them from deviating.

4.4 Analysing FDS coalitions with *n* symmetric players

For a given *n* and symmetric players, coalition structures are sufficiently described by the number of coalition members *s*. For different numbers of symmetric players, structures can be described by a pair (n, s). In this section, we characterise the set of structures satisfying FDS in a symmetric setting.

First note that, by construction, All Singleton structures $(n, 1)$ satisfy FDS. Next, a larger coalition $s > 1$ cannot satisfy FDS if members' payoffs are less than what they get in All Singletons. Hence, it is a necessary condition for a structure to satisfy FDS that

$$\frac{1}{s} V_S(n, s) \geq V_i(n, 1). \tag{4.11}$$

The right-hand side of Inequality (4.11) represents a player's payoff in the All Singletons structure. The left-hand side represents a member's payoff in a coalition of size *s*. If this inequality holds, then members (weakly) prefer to remain in structure (n, s) rather than induce $(n, 1)$.

Inequality (4.11) can be simplified by cancelling out the economic and biological parameters. This is shown in Appendix 4.2a. We obtain the following inequality which only contains n and s.

$$\frac{1}{s(n-s+2)^2} \geq \frac{1}{(n+1)^2} \tag{4.12}$$

Inequality (4.12) trivially holds for $s = 1$. Solving for s shows that for $s > 1$, the inequality holds when.

$$s \geq \frac{3}{2} + n - \frac{1}{2}\sqrt{4n + 5}. \tag{4.13}$$

Using Inequality (4.13) we can, for a given n, identify the smallest integer s for which (4.13) holds. This gives us a set of coalition structures from which no player would deviate because any deviation would induce $(n, 1)$, which gives a smaller payoff. This set of coalition structures thus satisfies FDS. In Figure 4.2, all structures that satisfy FDS are marked by ■. Structures satisfying FDS where $s = 1$, display trivial FDS and those where $s > 1$ display non-trivial FDS.

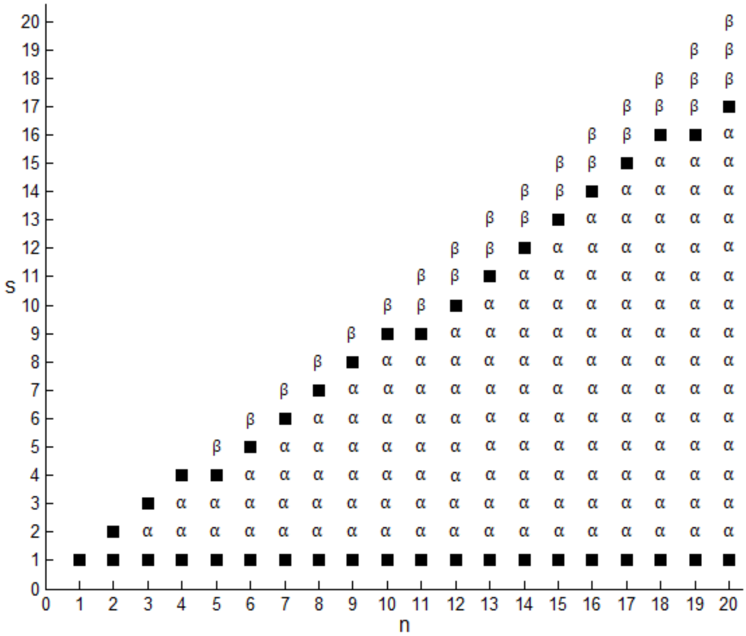

Figure 4.2: Coalition structures satisfying FDS. Pairs (n, s) marked by ■ satisfy FDS. Coalition structures marked by α and β do not satisfy FDS. From structures marked by α, the structure $(n, 1)$ will be induced. From structures marked by β, structures displaying non-trivial FDS will be induced.

For a given n, coalitions for which s is too small to satisfy Inequality (4.14) therefore do not satisfy FDS. Such structures are marked by α. To complete the characterisation of structures satisfying FDS, note that Inequality (4.13) is only a necessary condition for FDS. While all structures satisfying FDS must satisfy (4.13), satisfying (4.13) is not sufficient for FDS. For example, consider a Grand Coalition with 15 players. While (15,15) satisfies Inequality (4.14), we have not yet shown that this structure does not satisfy FDS because, as we will show, (15,13) is credibly inducible from (15,15). Hence, our final step is to prove that structures marked by β in Figure 4.2 do not satisfy FDS.

Consider the incentives to leave the Grand Coalition. A player would deviate if his payoff in the structure which is induced by his deviation (and hence satisfies FDS) provides a larger payoff than his Grand Coalition payoff:

$$V_{j\notin S}(s,n) \geq \frac{1}{n} V_N(n,n). \tag{4.14}$$

Similar to the derivations in Appendix 4.2a, when substituting the payoffs into (4.14), again, economic and biological parameters cancel and we obtain

$$s \geq 2 + n - 2\sqrt{n}\,. \tag{4.15}$$

We compare the conditions for s in (4.13) and (4.15) in Appendix 4.2b. The comparison shows that the minimum coalition size required for positive incentives to deviate from the Grand Coalition (4.15) is always smaller than the minimum size of a coalition satisfying FDS, determined by (4.13). Hence, there are incentives to deviate from Grand Coalitions marked by β because structures displaying non-trivial FDS satisfy (4.15). Furthermore, member payoffs in any coalition $s < n$ in structures marked by β are lower than the Grand Coalition member payoff. Therefore members of these partial coalitions also have an incentive to deviate. We have thus fully characterised the set of structures satisfying FDS in our symmetric fisheries game.

4.5 Analytics of asymmetry and changing stock location

We will now begin to analyse the role of asymmetry in the model. This section presents analytical results on the effects of changes in stock location. These results will allow us to understand the results to be presented in the sensitivity analysis to follow. Asymmetry in the Gordon-Schaefer model has been studied by Pintassilgo et al. (2010) who consider cost-asymmetry. It turns out that there are important differences between the effects of cost-asymmetry and catchability-asymmetry in the model. Understanding these differences allow us to understand and compare results.

Firstly, let us consider the equilibrium effort strategy given in Equation (4.7). It is immediately obvious that effort is decreasing c_i. However, taking the first derivative of Equation (4.7) with respect to q_i shows us that a threshold value of b_i exists which determines whether changes in q_i have a positive or negative effect on the equilibrium effort strategy. The threshold, where changes in q_i have no effect on the equilibrium effort strategy, is given by

$$\hat{b}_i = \frac{1}{2} - \frac{\sum_{j\in((N\setminus S)\cup\{l\})\setminus\{i\}}(1-b_j)}{2(n-s+1)}, \tag{4.16}$$

where l refers to the coalition member with the lowest cost-catchability ratio. If b_i is greater than the threshold value, equilibrium effort is increasing in q_i. Given that b_i is the inverse efficiency parameter, players who are inefficient in terms of b_i respond to increases in q_i by increasing their effort, whereas those who are efficient in terms of b_i respond to increases in q_i by decreasing their effort. Because effort is always decreasing in c_i, we can see that different asymmetries may have very different effects.

Intuitively, cost reductions or catchability increases for a given player relative to other players will increase the harvest of that player. The key difference between the marginal effects of c_i and q_i is that increases in harvests due to decreases in c_i will be achieved by increasing effort, whereas, from the derivation of (4.16), we know that increases in harvest due to increases in q_i can be achieved with a reduction in effort. Therefore, increases in catchability allow efficient players to reduce effort and costs while simultaneously increasing harvests. On the other hand, reductions in cost can only be exploited by increasing effort, which is costly, if albeit at a lower unit rate. This suggests that favourable marginal changes in q_i may be more profitable to a player than favourable marginal changes in c_i. We examine this proposition for the *n* player case, and find that

$$-\frac{\partial \Pi_i}{\partial c_i} < \frac{\partial \Pi_i}{\partial q_i} \iff q_i < c_i. \tag{4.17}$$

Therefore, marginal increases in q_i are more beneficial than marginal reductions in c_i as long as q_i is less than c_i. We provide more details on the derivation of condition (4.17) in Appendix 4.4.

4.6 Numerical and sensitivity analyses

The purpose of this section is to expand our analysis to include asymmetric players in order to analyse the effects of changing stock location. Accordingly, we must operationalize the FDS solution concept to deal with asymmetric players in a numerical setting. The main challenge in doing so is in defining players' outside options.

By construction, FDS allows us to exclude the possibility of cycles. However, introducing asymmetry requires us to deal with two additional problems in order to identify outside options and calculate stability. The root of these problems is in Definition 4.1.2, which requires that for a coalition to be unstable, *at least one player* must have an incentive to deviate. If a coalition does not generate enough profit to satisfy outside options (i.e. there is a negative surplus), all players have an incentive to deviate because the negative surplus is shared among players. Therefore, the payoff of coalition membership will be less than the outside option. Accordingly, there may be many ordered sequences which end with credibly inducible structures. In order to calculate the outside option for player *i*, we need to know which structures satisfying FDS could result from a deviation by player *i* from the Grand Coalition. We label this set of structures as the feasible set.

Definition 4.2*: The feasible set of player i, f_i, is defined as the set of all S_k which can be reached from the Grand Coalition via ordered sequences resulting in credibly inducible structures.*

An example of the feasible set for player 1 is shown in Figure 4.3 which illustrates a 4-player example where all coalition members have an incentive to deviate from the {2,3,4} structure. The structures {3,4} and {2,4} are stable such that the payoffs to player 1 in these structures are included in player 1's feasible set. {2,3} is however, not stable but the All Singletons structure is. The payoff to player 1 in the All Singletons is therefore included and those in the {2,3} structure are excluded. In this case then, f_1 has three elements because there are three structures which can be credibly induced via ordered sequences. We now need to use the elements of a player's feasible set to define the outside option of that player.

Definition 4.3 assumes that players will always deviate if there is at least one structure in the feasible set which provides a higher payoff. This reflects a cautious approach because it is less likely to lead to stable Grand Coalitions than, for example, taking the mean. In this way, stability can only be underestimated.

Finally, the FDS concept is based on static conjectures. Players foresee the eventual result of other players' actions and do not hesitate by waiting for the next period to deviate.

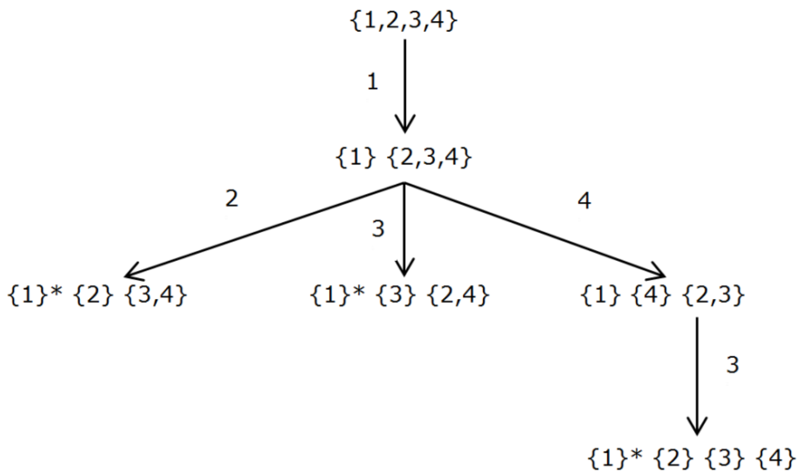

Figure 4.3: The elements of the feasible set for player 1 are marked by *.

Definition 4.3*: The outside option is defined as* $\omega_i = \max_{S \in f_i} V_i(S)$.

Having fully defined our FDS concept and illustrated how it is applied in practise, we can now continue the analysis. In reality, fisheries games are characterised by asymmetries in catchability and costs. We work on the assumption that fishing nations are asymmetric in their catchability and we wish to see how stability of Grand Coalitions is affected by changes in catchability.

4.6.1 Changes in stock location with three players

To begin, we consider the 3-player case with a specific scenario for changing stock location whereby catchability shifts entirely from Player 1 to Player 2. For simplicity, let us assume that Player 3 is not affected by the changing stock location. In this way, the sum of the catchability of the three players remains constant. Following the illustrative example, we carry out a sensitivity analysis to obtain more general results.

We define the vector $Q = (q_1, q_2, q_3)$ to denote position of the fish stock relative to the three players. In our illustrative example, we consider 80 different values of Q, which are ordered to represent gradual stock change as shown in Figure 4.4.

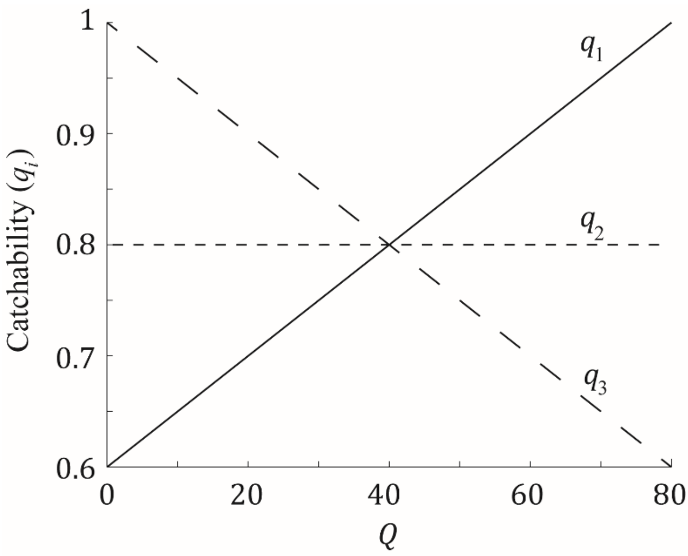

Figure 4.4: Values of catchability for 3 players across 80 ordered sets of Q.

We limit our analysis to values of q_i which ensure that players are always efficient enough to choose positive fishing effort. This allows us to isolate the effect of catchability change from changes in the number of players who are actively fishing. Therefore, we never consider values of q_i lower than 0.6. Figure 4.5 shows stability results for our illustrative example using the following parameterisations; $p = 1$; $k = 10$; $c_1 = c_2 = c_3 = 1.5$. We do not specify a value for the parameter r because r has no effect on the stability of coalitions. This is proven in Appendix 4.3.We choose this parameterisation because it allows us to fully illustrate the potential differences between the Internal Stability and FDS solution concepts.

Here, the Grand Coalition is internally unstable under all parameterisations of catchability. Precisely the opposite is true when the FDS is considered. The Grand Coalition satisfies FDS under all parameterisations of catchability.

To explain these results, note that under Internal Stability, the outside options are always drawn from the remaining 2-player coalition (regardless of its stability). Note also, that under FDS, the set of stable partial coalitions changes and hence the set of coalition structures from which outside options are drawn changes also. Figure 4.5(a) shows how, as symmetry increases (perfect symmetry exists where $Q = 40$), internal instability becomes more severe. This conforms with previously established results. Similarly, for a given set of stable partial coalitions, FDS of the Grand Coalition also decreases as symmetry increases. However, under FDS, the set of stable partial coalitions changes. These changes result in discontinuous jumps in stability. Increasing asymmetry increases the size of the set of stable partial coalitions and this decreases FDS while, for a given set of stable partial coalitions, increasing asymmetry increases FDS.

In order to make more general comparisons between the effects on changing stock location under the two solution concepts, we carry out a sensitivity analysis. We use bold type face to denote sets of values of a given parameter used in the analysis. The analysis tests over a discrete parameter space given by $\theta \times Q$ where $\theta = (\mathbf{p}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{k})$. Our approach is to determine appropriate values for the set θ and analyse the properties of each element of θ as stock location changes.

Appropriate values of θ need to allow for comparison of results with Pintassilgo et al. (2010). As such, we require a uniform distribution of b_i over a suitable range. However, the asymmetric q_i in our case precludes collecting terms p, c_i, q_i and k into the single parameter b_i in the profit function *a la* Pintassilgo et al. (2010). We therefore require a procedure which tests a uniform distribution of b_i but also specifies specific parameters for p, c_i and k, thus determining the set θ.

The results of the sensitivity analysis in Table 4.1 demonstrate that outcomes similar to those in Figure 4.5(a) occur for 34.1% (Never IS) of parameterisations for θ. Outcomes similar to those in Figure 4.5(b) occur for 18.8% (Always FDS) of parameterisations of θ. Further analysing the cost-symmetric case, the results also show that Sometimes FDS and Sometimes IS are the most common outcomes. However, Always IS never occurs but Always FDS occurs for 18.8% of parameterisations of θ in the cost-symmetric case. Additionally, Sometimes IS is less common than Sometimes FDS. Within our range of catchability changes, the use of the FDS solution concept offers significant stability improvements compared to Internal Stability.

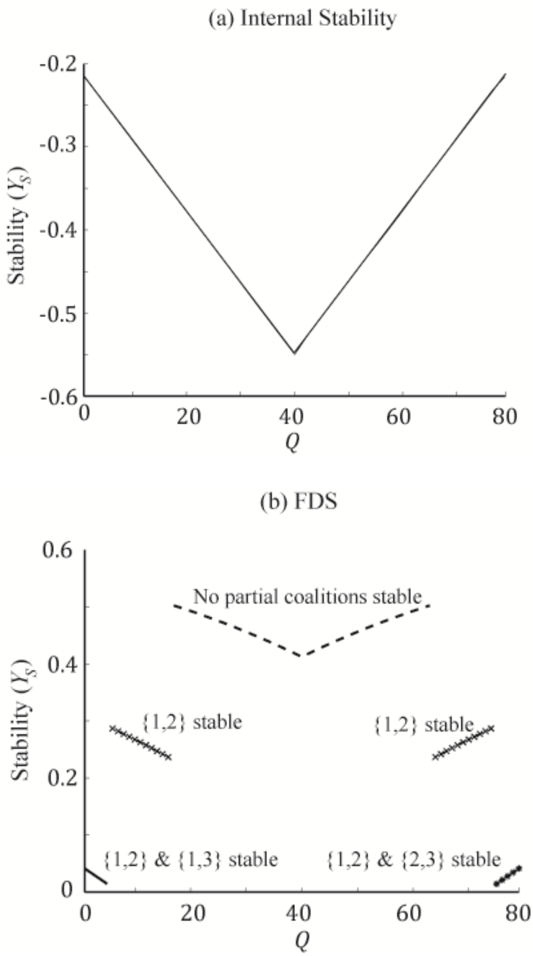

Figure 4.5: Stability values for a 3-player Grand Coalition with cost symmetry for the different parameterisations indicated in Figure 4.4. Panel (a) shows internal stability which is always negative, indicating an unstable Grand Coalition across all parameterisations. Panel (b) shows FDS. Different parameterisations result in varying FDS of partial coalitions, which in turn affects the outside options for players in the Grand Coalition and thus the stability of the Grand Coalition. For example, at parameterisation 10, the partial coalition {1,2} satisfies FDS. This means that the outside option of Player 3 leaving the Grand Coalition is determined in the structure given by coalition {1,2}.

In determining θ, we note that the variables p, c_i and k have no upper bound. However, a necessary condition for positive effort of player *i* is that b_i must be in the interval $0 \leq b_i < 1$. Therefore, we choose the values p, c_i and k such that b_i does not exceed its bounds for any tested value of q_i. Additionally, we select values for p, c_i and k such that the values of b_i are uniformly distributed in the interval $[0, 1)$. Values for p, c_i and k are chosen from sets of random draws from uniform distributions in the following intervals; $0 < c_i < 2$, $0 < p < 2$ and $1 < k < 100$. We select these intervals because they are reasonable and allow for a full range of b_i. In addition, due to the result given in (4.17), it is appropriate to allow c_i to be less and greater than q_i.

We retain our assumption that changes in stock location occur in the range $0.6 \leq q_i \leq 1$. Elements of Q are drawn from a uniform distribution in the interval of $[0.6, 1]$ and obey the criterion that $q_1 + q_2 + q_3 = 2.4$. The summation criterion ensures that total catchability is always constant and thus represents the case where catchability is redistributed between players[5].

The parameter space $\theta \times Q$ has approximately 100,000 elements which provides sufficient confidence in the results. For the cost-symmetric and cost-asymmetric case and for each element of $\theta \times Q$, we test for FDS and internal stability. The sensitivity analysis is programmed in Matlab to classify each element of θ into certain categories (see Table 4.1) depending on the stability of the Grand Coalition as stock location changes.

Table 4.1 also shows that cost-asymmetry increases stability for both solution concepts. When cost asymmetry is introduced, the potential range of the cost-catchability ratio for the three players is greater when costs are asymmetric and thus, it is more likely that the most efficient member of the coalition can satisfy outside option requirements. We therefore see an increase in stability for both the FDS and IS solution concepts. Cost-asymmetry has a greater effect under the Internal Stability solution concept, as evidenced by the larger increase in Always IS.

---

[5] In this way, we lose the ordering of Q as seen in Figure 4.4. Given our method of statistical analysis, losing ordering does not affect the interpretation of the results. If the random draws were ordered to represent a changing stock location scenario over time, as in Figure 4.4, the results would be the same as without ordering. Further, this method benefits from not presenting scenarios as in Figure 4.4 because imposing such a scenario is restrictive, particularly in the case of the four player game.

This indicates that the FDS solution concept is less reliant on cost-asymmetry to improve stability than Internal Stability.

Table 4.1: Results of sensitivity analyses for the cost-symmetric and cost-asymmetric cases in a 3-player game. There are three categories. FDS and Internal Stability (IS) are reported for each of the three categories as percentages of the elements of θ which fall into each category. Firstly, "Always FDS / Always IS"; where the stability concept is satisfied across all parameterisations of Q. Secondly, "Sometimes FDS / Sometimes IS"; where the stability concept is satisfied for at least one, but not all parameterisations of Q. Thirdly, "Never FDS / Never IS"; where the stability concept is not satisfied for any parameterisation of Q.

	Cost-symmetric	Cost-asymmetric
Always FDS (%)	18.8	38.98
Sometimes FDS (%)	81.2	61.02
Never FDS (%)	0	0
Always IS (%)	0	25.17
Sometimes IS (%)	65.9	48.16
Never IS (%)	34.1	26.67

In general then, stability increases under the FDS solution concept and in asymmetry. An interesting aspect of the results is that the FDS solution concept results in more frequent occurrences of "Sometimes FDS" in both the cost-symmetric and asymmetric cases. This means that, under the FDS solution concept, changing stock location is more likely to render a stable Grand Coalition unstable (or vice versa). Therefore, while the FDS concept results in more stability in general, it also shows more sensitivity to changing stock location. Sensitivity to changing stock location increases for FDS relative to Internal Stability because of the different way that the outside option is calculated. Using the 3-player game as an example, under Internal Stability, the outside options are always drawn from the payoffs of free-riders playing against partial coalitions. Under FDS, each outside option is calculated according to the stability of the partial coalitions. Changes in stock location can change the stability of partial coalitions and thus lead to greater variation in the outside options as stock location changes. In turn, this increases the sensitivity of stability to changing stock location.

4.6.2 Changes in stock location with four players

We now examine the effect of a unit increase in the number of players. In order to do so, we employ the same sensitivity analysis procedure as in the previous sub-section. The only changes are that $\theta = (\mathbf{p}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}_4, \mathbf{k})$, the random draws for catchability must now obey the criterion that $q_1 + q_2 + q_3 + q_4 = 3.2$ and the increased number of players increases the number of elements in $\theta \times Q$.

The results in Table 4.2, in comparison to Table 4.1, show that, as expected, increasing the number of players from 3 to 4 decreases stability for both solution concepts and in both the cost-symmetric and asymmetric cases. Cost-asymmetry increases Internal Stability and FDS. Again, FDS offers improvements in stability overall, but also increases the frequency of "Sometimes FDS" in both cost-symmetric and asymmetric cases and thus increases the sensitivity of Grand Coalition stability to changing stock location. Comparison of Table 4.1 and Table 4.2 shows that the problem of sensitivity with FDS becomes more severe as the number of players increases. There are two reasons for this. Firstly, in general, stability decreases in *n.* Secondly, in the 3-player case there are only 4 structures (3 partial coalitions and All Singletons) from which the outside options can be drawn. In the 4-player case, there are 11 structures. Thus, changes in stock location can change the stability of a greater number partial coalitions and thus lead to greater variation in the outside option.

In addition to analysing the stability properties of each element of θ, we can also analyse each element of $\theta \times Q$ individually. This allows for direct comparison to Pintassilgo et al. (2010). For four player games, considering asymmetry in c_i (which is represented by asymmetry in b_i), Pintassilgo et al. (2010) find that the Grand Coalition will be Internally Stable in 5.1% of cases. In our cost-symmetric case, the percentage of elements of $\theta \times Q$ for which the Grand Coalition is internally stable is 22% (not shown in Table 4.2). This shows that asymmetry in q_i is more likely to lead to stability than asymmetry in c_i.

Table 4.2: Results of sensitivity analysis for a cost-symmetric and cost-asymmetric case in a 4-player game. For definitions of the terms in the first column, see Table 4.1.

	Cost-symmetric	Cost-asymmetric
Always FDS (%)	0	8.87
Sometimes FDS (%)	100	78.29
Never FDS (%)	0	12.84
Always IS (%)	0	12.23
Sometimes IS (%)	48.15	51.07
Never IS (%)	51.85	36.70

The reason for this difference has already been partially explained in Section 4.5, where we established that marginal increases in q_i lead to a greater increase in profit than marginal reductions in c_i if and only if $q_i < c_i$. Therefore, when $q_i < c_i$ holds, a given degree of asymmetry in q_i can lead to greater differences in payoffs between players than the same degree of asymmetry in c_i. In our sensitivity analysis, due to our selection of parameter ranges, $q_i < c_i$ may or may not hold. When it does hold, a given degree of asymmetry in q_i can allow the most efficient coalition member to be more profitable than for a given degree of asymmetry in c_i. Thus, the most efficient coalition member is more likely to be able to satisfy outside options. This explains the increased stability in our case relative to Pintassilgo et al. (2010).

Of course, the selection of intervals for q_i and c_i in the sensitivity analysis affects this result. For example, if c_i is always greater than q_i then the Grand Coalition will be stable for more elements of $\theta \times Q$. In general, it holds that increases in q_i which reinforce existing cost advantages of the most efficient coalition member will increase stability. However, the relative magnitudes of c_i and q_i are important in determining the marginal effect of changes in q_i on coalition stability.

The results show that the FDS solution concept offers consistently more Grand Coalition stability under stock location changes in both 3 and 4-player games.

Although, the improvements in stability due to the FDS solution concept are accompanied by an increase in the sensitivity of stability to changing stock location.

4.7 Discussion and conclusion

The results of this study contribute to an improved understanding of the impacts of changing stock location on the potential for full cooperation regarding fish stocks. When players are symmetric and changes in stock location affect all players equally, the Grand Coalition satisfies FDS for $n \leq 4$. For $n > 4$, large partial coalitions can also be FDS. When considering cost-symmetry and cost-asymmetry combined with changing stock location, we find that FDS leads to an increase in stability relative to Internal Stability. However, while stable Grand Coalitions are more likely under the FDS solution concept, changes in whether a Grand Coalition is stable due to changing stock location are also more likely. In this way, the use of the FDS solution concept increases stability, but also increases the sensitivity of stability to changes in stock location.

Finally, we discuss some important issues highlighted by our results. We need to consider the positivist aspect of which solution concept best reflects reality and the normative aspect of which behaviours implied by the solution concept are preferable. The normative aspect is clear. Should policy makers wish to increase the stability of Grand Coalitions under changes in stock location, then mechanisms could be put into place to encourage further deviation, thereby forcing players who are considering a deviation to make farsighted conjectures about the effects of their deviation. In doing so, policy makers should consider that such farsighted conjectures may lead to more frequent switches been stable and unstable Grand Coalitions as a consequence of the increase in sensitivity to changing stock location associated with FDS.

The positivist question is less clear cut. In the simplest, symmetric case, Grand Coalitions are unstable for more than 4 players. In reality, there have been examples of both success and failure of fisheries agreements for various numbers of fishing states (see Munro, 2008). This offers some support for the notion that fishing countries are posing credible threats to deviation and that this facilitates stable coalitions.

There are also other approaches which could be employed to analyse the issue of changes in stock location. For example, Long and Flaaten (2011) use a Stackelberg game to analyse the potential for cooperation to manage straddling fish stocks and have found more optimistic results than in the literature based on Cournot games. Breton and Keoula (2012) employ a dynamic farsightedness concept as well as a static version. A dynamic structure has the potential to increase the stability of the Grand Coalitions when stable coalition structures are reached after a large number of deviations (for example, the All Singletons structure) and the discount rate is sufficiently low. Higher payoffs from deviations will therefore be reduced through discounting and this will help to stabilise Grand Coalitions.

Appendix 4.1 Equilibrium strategy for the All Singletons structure.

Individual profits are given by

$$\Pi_i = pq_ie_ix - c_ie_i. \tag{4.i}$$

Using the steady state condition, solving for *x* and substituting the value for *x* gives us

$$\Pi_i = pq_ie_i\left(\frac{k}{r}\left(r - q_ie_i - \sum_{j\in N\setminus\{i\}} q_je_j\right)\right) - c_ie_{i,}. \tag{4.ii}$$

The first order condition is

$$\Pi'_i = pq_ik - \frac{2pq_i^2ke_i}{r} - \frac{pq_i\sum_{j\in N_{-i}} q_je_j}{r} - c_i = 0. \tag{4.iii}$$

Solving for effort gives the reaction function

$$e_i = \frac{r}{2q_i}(1 - b_i) - \frac{1}{2q_i}\sum_{j\in N\setminus\{i\}} q_je_j \tag{4.iv}$$

where $b_i = \frac{c_i}{pq_ik}$.

We manipulate the reaction function to obtain the following two identities.

$$e_i = \frac{r(1-b_i)}{q_i} - \frac{1}{q_i}\sum_{i\in N} q_i e_i \quad (4.v)$$

$$\sum_{i\in N} q_i e_i = \frac{r}{n+1}\sum_{i\in N}(1-b_i) \quad (4.vi)$$

Substituting the first identity into the second gives us the equilibrium strategy

$$e_i = \frac{r(1-b_i)}{q_i} - \frac{r}{q_i(n+1)}\sum_{i\in N}(1-b_i), \quad (4.vii)$$

the RHS of which can be manipulated so that e_i is a function of the inverse efficiency parameters of all other players such that,

$$e_i = \frac{nr}{(n+1)q_i}(1-b_i) - \frac{r}{(n+1)q_i}\sum_{j\in N\setminus\{i\}}(1-b_j). \quad (4.viii)$$

Appendix 4.2a Derivation of Inequality (4.12)

$$\frac{1}{s}V_S(n,s) \geq V_i(n,1) \quad (4.ix)$$

$$\frac{r(1-b)}{n+1}\left[pk\left(1-\frac{n(1-b)}{n+1}\right)-\frac{c}{q}\right] \leq \frac{r(1-b)}{s(n-s+2)}\left[pk\left(1-\frac{(n-s+1)(1-b)}{n-s+2}\right)-\frac{c}{q}\right]. \quad (4.x)$$

Cancelling terms outside the square brackets and dividing both sides by pk yeilds

$$\frac{1}{n+1}\left[(1-b)-\frac{n(1-b)}{n+1}\right] \leq \frac{1}{s(n-s+2)}\left[(1-b)-\frac{(n-s+1)(1-b)}{n-s+2}\right]. \quad (4.xi)$$

Removing $(1-b)$ from inside the brackets and cancelling yields

$$\frac{1}{n+1}\left(\frac{1}{n+1}\right) \leq \frac{1}{s(n-s+2)}\left(\frac{1}{n-s+2}\right). \quad (4.xii)$$

Which simplifies to

$$\frac{1}{(n+1)^2} \leq \frac{1}{s(n-s+2)^2}\,.$$

Solving for s gives the result in (4.13). (4.xiii)

Appendix 4.2b Comparison of (4.13) and (4.15)

We examine the difference between the right hand sides of (4.13) and (4.15).

$$\left(\frac{3}{2}+n-\frac{1}{2}\sqrt{4n+5}\right)-\left(2+n-2\sqrt{n}\right) \tag{4.xiv}$$

$$\Leftrightarrow -\frac{1}{2}+2\sqrt{n}-\frac{1}{2}\sqrt{5+4n} \tag{4.xv}$$

Hence, the difference is positive (and increasing) for all $n > 2$.

Appendix 4.3 The independence of coalition stability from *r*

We will demonstrate that the parameter *r* has no effect on stability. In order to do so, we prove that the ordering of any two profit functions, regardless of coalition size or membership, does not depend on *r*. The subscript *j* and *k* are used to denote the steady state stock size and effort under different coalition sizes or membership choices and hold their usual meaning for c_i and q_i.

Consider the ordering

$$pq_je_jx_j - c_je_j > \ pq_ke_kx_k - c_ke_k. \tag{4.xvi}$$

Equations (4.3) and (4.7) from the main text are repeated below.

$$x_i = k - \frac{k}{r}\sum_{i=1}^{N} q_ie_i \tag{4.3}$$

$$e_i = \frac{(n-s+1)r}{(n-s+2)q_i}(1-b_i) - \frac{r}{(n-s+2)q_i}\sum_{j\in((N\backslash S)\cup\{l\})\backslash\{i\}}(1-b_j) \text{ for } i \in (N\backslash S)\cup\{l\} \tag{4.7}$$

Substitution of (4.7) into (4.3) shows that x is not a function of *r* because the *r* terms cancel. Effort e_i is linear in *r* and can thus be entirely cancelled from the ordering. Given that the equality of any two profits does not depend on *r,* it holds also that the payoff of a given coalition member (Equation 4.8) also does not depend on *r*. To see this, note that given Equation (4.9), $V(S) > \sum_{j\in S}\omega_j$ is a sufficient condition for stability. The argument for the independence of the ordering of any two profit functions from *r* thus holds also for payoffs.

Appendix 4.4 The relative advantage of cost versus catchability asymmetries

Beginning by substituting Equations (4.3) and (4.7) into Equation (4.4) and simplifying, we have equilibrium profit

$$\Pi_i = prk\left[B_i\frac{a}{a+1} - \frac{\sum_j B_j}{a+1}\right]\begin{bmatrix}1 - B_i\frac{a}{a+1} + \frac{\sum_j B_j}{a+1} + \\ \sum_j\left(B_j\frac{a}{a+1} - \frac{\sum_k B_k}{a+1}\right)\end{bmatrix} - \frac{c_i r}{q_i}\left[B_i\frac{a}{a+1} - \frac{\sum_j B_j}{a+1}\right], \qquad (4.\text{xvii})$$

where $B_i \equiv 1 - b_i\ \forall i$, $a \equiv n - s + 1$, $\sum_j$ sums over all $j \in N\backslash i$ and $\sum_k$ sums over all $k \in N\backslash j$.

The first differentials of the profit function are given below,

$$\frac{\partial \Pi_i}{\partial c_i} = \frac{r}{q_i(a+1)^2}\left[(2a - 2a^2)(1 - b_i) + (2a - 1)\textstyle\sum_j\left(1 - b_j\right)\right], \qquad (4.\text{xviii})$$

$$\frac{\partial \Pi_i}{\partial q_i} = -\frac{c_i r}{q_i^2(a+1)^2}\left[(2a - 2a^2)(1 - b_i) + (2a - 1)\textstyle\sum_j\left(1 - b_j\right)\right]. \qquad (4.\text{xix})$$

5 The rise and fall of the great fish pact under endogenous risk of stock collapse

ABSTRACT

Risk of stock collapse is a genuine motivation for cooperative fisheries management. We analyse the effect of an endogenously determined risk of stock collapse on the incentives to cooperate in a Great Fish War model. We establish that equilibrium harvest strategies are non-linear in stock and find that Grand Coalitions can be stable for any number of players if free-riding results in a total depletion of the fish stock. The results thus show conditions under which a Great Fish War becomes a Great Fish Pact. Importantly, this conclusion no longer holds upon dropping the standard assumption that payoffs are evaluated in steady states. If payoffs in the transition between steady states are included, the increased incentives to deviate offset the increased benefits from cooperation due to the presence of endogenous risk and the Great Fish Pact returns to being a Great Fish War.

This chapter is based on the paper: Walker, A.N., Weikard, H-P., Richter, A., 2015. The rise and fall of the great fish pact under endogenous risk of stock collapse. FEEM working paper 2015.060. Submitted to a peer reviewed journal.

5.1 Introduction

Risks of catastrophe or regime shifts, if endogenously determined, have been shown to be an important incentive for precaution in strategic resource use (Nkuiya et al., 2014; Ren and Polasky, 2014; Sakamoto, 2014). Further, such risks are relevant for understanding the decision to join climate treaties (Kolstad, 2007; Dellink and Finus, 2012; Barrett, 2013) and are likely to affect the strategic harvest choices of fishing nations (Hannesson, 2014). From 1950 to 2000, 366 fisheries collapsed and the collapses are generally attributed to over-fishing (Mullon et al., 2005). Indeed, avoiding stock collapses was one of the principle motivations for the formation of Regional Fisheries Management Organisations (RFMOs), the institutions intended to facilitate cooperation in the management of high seas fish stocks. Surprisingly, the effect of endogenously determined catastrophes on the potential for cooperation in fisheries agreements has received little attention in the literature.

In this paper, we fill this gap using the Great Fish War model of Levhari and Mirman (1980) and consider the effects of a risk of stock collapse which increases in harvest. We ask, whether an endogenous risk of stock collapse can transform the Great Fish War into a Great Fish Pact. We modify the Great Fish War model of Levhari and Mirman (1980) (henceforth LM) to estimate Markov Perfect Nash Equilibrium (MPNE) harvest functions under an endogenously determined risk of irreversible collapse such that the stock after the collapse is zero, and remains zero, for all future time periods. It should be noted that "collapse", as defined in the fisheries literature does not require the stock to be completely extinct (Cooke, 1984). Instead, we define collapse as an economic collapse, meaning that the fishery is no longer viable and no profits can be made. Our study therefore relates generally to the literature on uncertainty in resource management such as Clarke and Reed (1994) and Tsur and Zemel (1998). More specifically, our study relates to literature which considers endogenous risk of regime shift in resource games, namely, Sakomoto's (2014) analysis of the subclass of dynamic renewable resource games of Sorger (2005) and Ren and Polasky (2014), who conduct a more general analysis. These two studies show that endogenous risk can lead to either more or less aggressive resource use. Additionally, Sakamoto (2014) demonstrates the importance of considering the transition between regimes, i.e. taking off-steady-state payoffs into account. Finally, our study fits directly into the literature using the LM model. Exogenous uncertainty in the LM model has been considered in three studies. Antoniadou et al. (2013) and

Agbo (2014) consider exogenous uncertainty in stock dynamics. Fesselmeyer and Santugini (2013) consider exogenous uncertainty in the quality of the resource as well as the probability of regime shifts in the growth rate of the stock.

In our study, we compare analytically how incorporating endogenous risk affects the structure of the LM model. The inclusions of endogenous risk results in equilibrium harvest strategies which are non-linear in stock. Due to the technical difficulties in analytically solving for non-linear harvest functions (Antoniadou et al., 2013), we employ a numerical approach. Our model calculates the Internal Stability of Grand Coalitions across a range of growth and discount rates and for any number of players. In turn, this allows us to determine if an endogenous risk of stock collapse affects the potential for successful cooperation.

Our study is the first detailed exploration of non-linear harvest functions in the LM model and the first study to explicitly consider endogenous risk of stock collapse from a coalition theory perspective. We find that endogenous risk of stock collapse may provide an incentive to entirely deplete the fish stock. Because entirely depleting the stock is a response to the presence of endogenous risk, we term this "pre-emptive depletion". The effect of pre-emptive depletion on coalition stability depends on the assumptions adopted regarding how the payoff from deviation is calculated. The application of a two-stage game (d'Aspremont et al., 1983) to fisheries results in the commonly used assumption (Lindroos, 2008; Pintassilgo, 2008; Pintassilgo et al. 2010; Long and Flaaten, 2011; Ellefsen, 2012) that players receive payoffs calculated in steady state according to the coalition formed. Under this standard assumption, we find that, in general, an endogenous risk of stock collapse increases Grand Coalition stability. This is particularly so if non-cooperation would result in pre-emptive depletion. When this is the case, the incentive to cooperate is so strong that the Grand Coalition is stable for any number of players and can therefore be described as a Great Fish Pact. This study therefore suggests a solution to the "puzzle of small coalitions" (Breton and Keoula, 2014), whereby the size of theoretically stable coalitions is smaller than what is observed in reality. Notably, the puzzle of small coalitions can be solved without the use of transfer payments. Transfer or "side" payments combined with asymmetric players is a frequently invoked and powerful method which increases the number of players for which cooperation can be sustained (Kaitala and Lindroos, 1998; Kennedy, 2003; Lindroos 2008; Pintassilgo et al. 2010; Long and Flaaten, 2011; Ellefsen, 2012; Breton and Keoula, 2014). However, transfers payments have met much resistance in the policy world in general (Folmer et al., 1993) and are not implemented in direct financial terms in fisheries agreements (Munro, 2008). Further, the puzzle of small coalitions is solved without the use of sequential move games (e.g. Long and Flaaten, 2011) or alternative solution concepts, such as farsightedness (e.g. Breton and Keoula, 2012).

Thus, endogenous risk solves the puzzle of small coalitions (the Great Fish War becomes a pact). However, we find that this result is very sensitive to the assumptions implicit in the standard two-stage game. We relax the assumption that payoffs are determined in steady states by considering a transition period whereby the stock size gradually adjusts after a deviation has occurred (cf. Sakamoto, 2014). Deviators receive payoffs during this transition period ("transition payoffs"). Transition payoffs turn out to be a decisive incentive for non-cooperation. Without transition payoffs, if deviation leads to pre-emptive depletion, then the payoff of deviation is zero. With transition payoffs, the process of pre-emptively depleting provides a payoff. We find that transition payoffs motivate non-cooperation to the extent that the Grand Coalition is only stable in a two-player game, and then, only if the discount rate is sufficiently low and the stock grows sufficiently slowly. The Great Fish Pact thus returns to a Great Fish War. Overall then, the paper shows how endogenous risk of stock collapse leads to dramatic increases in the potential for cooperation but qualifies this with the important proviso that this result holds only if transition payoffs are not considered.

The following Section 5.2 describes the bio-economic model and derives and analyses the envelope condition. Section 5.3 explains how Grand Coalition stability is calculated. Section 5.4 numerically analyses the model in terms of the stability of Grand Coalitions. Section 5.5 proceeds to consider the effects of including transition payoffs. Section 5.6 concludes.

5.2 Bio-economic model

We will first describe the biology of the system and introduce the objective functions. The objective functions determine the payoffs for a given coalition membership choice, which are then used to determine coalition stability. The set N of identical players represents n nations, indexed by i. Let us first define escapement e (the stock remaining after harvest) in a given period, t as

$$e_t \equiv x_t - \sum_i h_{i,t} \ , \tag{5.1}$$

where $h_{i,t}$ is the harvest of player i in period t. The stock in the next period depends on escapement in the current period and is determined by the function $f(e_t)$ as follows:

$$x_{t+1} = f(e_t) = \beta e_t^{\alpha}, \tag{5.2}$$

where $\beta > 0$ and $0 < \alpha < 1$. If there is no harvest, x_t increases over time to its carrying capacity, which is given by

$$\bar{x} = \beta^{\frac{1}{1-\alpha}}. \tag{5.3}$$

We normalise the model such that the carrying capacity is fixed and not affected by the growth parameters α and β. Specifically, we set $\beta = \bar{x}^{(1-\alpha)}$ (from Equation (5.3)) and thus the carrying capacity $\bar{x}$ can be treated as a parameter in the model and we only need to specify α, which we term the "growth parameter". Note that lower α entails a higher growth rate.

The probability of the fish stock surviving into the next period, $0 \leq r < 1$, is endogenously determined by the escapement and is given by

$$r(e_t) = \max\left(0, 1 - \frac{\gamma\bar{x}}{e_t}\right), \tag{5.4}$$

where $0 < \gamma < 1$ and therefore $0 \leq r(e) < 1$. The parameter γ determines the critical escapement level $\gamma\bar{x}$, below which collapse is certain. For any escapement level $e_t > \gamma\bar{x}$ such that $\max\left(0, 1 - \frac{\gamma\bar{x}}{e_t}\right) = 1 - \frac{\gamma\bar{x}}{e_t}$, it is easy to see that there is a strictly positive survival probability which is increasing in escapement at a decreasing rate. This means that there is a strictly positive risk of stock collapse at all stock sizes. This is reasonable because, for certain species, pressures from habitat loss or invasive species may mean that a risk of stock collapse is present even in the absence of any fishing (Field et al., 2009; Gjøsæter et al., 2009).

The instantaneous utility function for player i is given by

$$u(h_i) = \max(0, \ln(h_i)). \tag{5.5}$$

This utility function avoids the problem of being undefined when harvest is zero, which is useful in our numerical approach. Appendix 5.1 explains and validates the choice of utility function in more detail.

The value function of player i is given by

$$V_i(x) = u(h_i) + r(e)\delta V_i\big(f(e)\big), \tag{5.6}$$

where $0 < \delta < 1$ is the discount factor. The value function depends on instantaneous utility and the value of the stock in the future, subject to discounting and risk of collapse. Both the risk of collapse and the future value of the stock depend on escapement. Escapement depends on h_i and the sum of the harvests of other players h_{-i}. Therefore $e = x - h_i - h_{-i}$. Optimal harvest varies with stock size. Therefore, harvest level is represented as a function of stock size such that $h_i = \hbar_i(x)$, which we term the harvest function. Similarly, escapement is also a function of stock size, i.e. $e = x - \hbar_i(x) - \hbar_{-i}(x)$.

We can now begin to investigate how the harvest functions of our model differ from the LM harvest functions and what drives these differences. Optimal harvest maximises the value function for a given stock size. The envelope of these maxima across all stock sizes is the envelope curve. The envelope condition is a necessary condition for the maximisation of the envelope curve and thus gives insight into the conditions under which optimal stock size and harvest are achieved.

Lemma 5.1: *The envelope condition is given by* $\frac{\partial V_i}{\partial x} = \left(1 - \frac{\partial \hbar_{-i}}{\partial x}\right)\frac{\partial u}{\partial h_i}$.

Proof: *See Appendix 5.2.*

The envelope condition shows that player i's harvest is optimal when the marginal value of the fish stock $\frac{\partial V_i}{\partial x}$ is equal to the marginal value of harvest $\frac{du}{dh_i}$, which is adjusted by the proportion of the marginal harvest of all other players. The general format of the envelope condition, as in Lemma 5.1, is identical to that of the LM model (Mirman, 1979). However, the values of the derivatives are different because the endogenous risk function $r(e)$ affects the value function V_i. Therefore, harvest levels which satisfy the envelope condition will not be identical to the LM case. We can conjecture that the value function will be particularly steep at low stock sizes because small increases in stock size lead to large reductions in risk of collapse. Furthermore, the slope of the harvest functions will depend on whether x is less or greater than $\gamma\bar{x}$. By contrast, in the LM model the slope of the harvest function is constant. For a full analysis of the effects of endogenous risk in terms of Grand Coalition stability, we employ a numerical method, which is explained in the following section.

5.3 Grand Coalition stability

We now proceed to explain how the value function is optimised and how we use these results to analyse stability. We test for Grand Coalition stability across a parameter space Ω. Elements of Ω are triples (n, γ, θ) where $\theta = (\alpha, \rho)$. The parameter ρ is the discount rate where $0 < \rho \leq 1$ and $\delta = \frac{1}{1+\rho}$. The set of players in the coalition is given by M, where $m \equiv |M|$. We consider and compare two coalition structures. The first is the Grand Coalition given by $M = N$. A coalition member may deviate and will do so immediately should this be beneficial. This results in the second coalition structure; the partial coalition $M = N/\{k\}$, where $\{k\}$ is the free-rider. Coalition members choose harvest levels to maximise their joint utility and the free-rider chooses harvest to maximise individual utility. For each element of Ω, and for an infinite time horizon, we optimise the value function $V_i(x)$ for a given Ω to derive the optimised value functions $U_i(x_t; \Omega)$ for Grand Coalition members and for free-riders. This is achieved via the Bellman equation. For a coalition member j, the optimised value function $U_j(x_t; \Omega)$ is given by

$$U_j(x_t; \Omega) = \frac{1}{m} \max_H \left\{ \sum_{j \in M} \left(\max\left(0, \ln\left(\frac{H}{m}\right)\right) + \frac{r(e_t)}{1+\rho} U_j(x_{t+1}; \Omega) \right) \right\} \quad \forall\ j \in M\ , \tag{5.7}$$

where $H = \sum_{j \in M} h_j$ is the coalition harvest. The value function for a free-rider k playing against the coalition $N \backslash \{k\}$ is given by

$$U_k(x_t; \Omega) = \max_{h_k} \left\{ u(h_k) + \frac{r(e_t)}{1+\rho} U_k(x_{t+1}; \Omega) \right\}. \tag{5.8}$$

Where $n = 1$, the optimised value functions $U_j(x_t; \Omega)$ and $U_k(x_t; \Omega)$ are both equivalent to that in the sole owner case. Optimised value functions are calculated numerically using value function iteration. The harvest functions $\hbar_i(x)\ \forall\ i \in N$ which result in optimised value functions over an infinite time horizon thus constitute Markov Perfect Nash Equilibrium (MPNE) harvest functions. MPNE harvest functions allow us to determine the steady state stock size with harvesting at the stock size x^* for which the following equality holds;

$$x^* = \beta(x^* - \textstyle\sum_{i \in N} \hbar_i(x^*))^{\alpha}. \tag{5.9}$$

Evaluating the optimised value functions at x^*_{GC} (the steady state under a Grand Coalition) and x^*_{FR} (the steady state if free-riding occurs) gives payoffs, which determine Grand Coalition stability if transition payoffs are not included. We use an Internal Stability solution concept, under which the Grand Coalition is stable if the payoff to a Grand Coalition member is greater than that of a free-rider playing against the coalition of remaining members. If transition payoffs are not included, the Grand Coalition is therefore internally stable if

$$U_j(x^*_{GC}; \Omega) \geq U_k(x^*_{FR}; \Omega). \tag{5.10}$$

The internal stability condition is also applied under the inclusion of transition payoffs, but in that case, payoffs in the transition between steady states after a deviation from the Grand Coalition are accounted for. For more details and discussion of the numerical techniques used, see Appendix 5.1.

5.4 Results of the standard two-stage game

This section presents stability results for our game under the assumption that transition payoffs are excluded. We begin by validating the numerical accuracy of our model. The validation demonstrates high statistical similarity of harvest functions from a numerical LM model with analytically derived LM harvest functions; see Appendix 5.1 for details. We thus proceed to analyse the numerical model of endogenous risk of stock collapse. We consider a range of parameters[6] for α and ρ

[6] In the original LM model the discount factor is between 0 and 1 whereas we test the discount rate between 0 and 1. This means that we test discount factors between 0.5 and 1.

such that $\alpha \in A = [0.01, 0.02, \dots, 0.99]$ and $\rho \in P = [0.01, 0.02, \dots, 1]$. We denote the set of all possible $\theta = (\alpha, \rho)$ as Θ such that $\Theta = A \times P$. The disaggregation of A and P allows us to determine the stability of coalitions across a full range of parameters and therefore to acquire insights of a similar depth to those provided by analytical results. We do not analyse $\alpha = 1$ in order to retain strict concavity in the growth function. Further, we begin by analysing a low value for the parameter, γ. Lower critical escapement levels $\bar{x}\gamma$ mean that certain and immediate stock collapse occurs for a smaller range of stock sizes. We therefore set $\gamma = 0.01$ and consider the effect of changing γ later.

Figure 5.1 presents the resource stock in steady state for the parameter space Θ using $n = 2$ as a representative example. The analysis will distinguish results for the Grand Coalition (Panel A) and the case where free-riding occurs (Panel B). Note that the free-rider case for $n = 2$ coincides with the Cournot-Nash equilibrium. For each element of Θ, multiple steady states can exist. We first present and analyse the largest stable steady state for each element of Θ and later, we will describe the different steady states which can exist for each element of Θ in more detail.

Figure 5.1 shows that the largest stable steady state is either zero, as in Region I, or positive as in all other regions. In Region I, it is optimal to fish the stock to extinction rather than waiting in the hope that stocks will increase and thus the risk of collapse will drop. We refer to this effect as "pre-emptive depletion". Pre-emptive depletion occurs when α is high (the growth rate is low) and occurs for a greater range of α as the discount rate, ρ increases. Larger α and ρ mean that the stock has less value in the future: stock regeneration is limited, and any gains occurring in the future will be discounted. Further, the presence of endogenous risk makes those future gains uncertain. Hence, the choice is made to pre-emptively deplete the stock, thus gaining an immediate and certain payoff. In all other regions, "conservative management" occurs, whereby the largest stable steady state is positive. Conservative management occurs when the value of the future (in terms of α and ρ) is greater and thus maintaining a positive steady state becomes optimal, despite the risk of stock collapse.

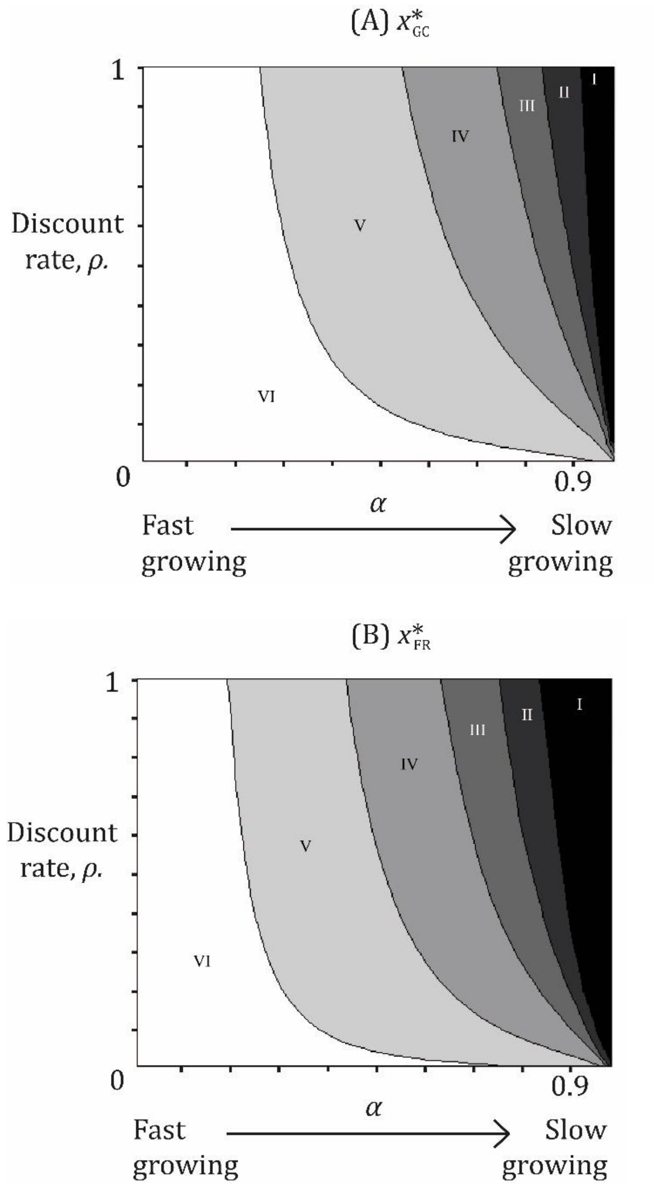

Figure 5.1: Largest stable steady states in the Grand Coalition case (Panel A) and the free-rider case (Panel B) in Θ space where $n = 2$, as an example. Using x_I^* to denote the steady state stock in Region I, x_{II}^* to denote the steady state stock in Region II and so on, the regions are defined as $x_I^* = 0$, $0 < x_{II}^* \leq 1000$, $1000 < x_{III}^* \leq 2000$, $2000 < x_{IV}^* \leq 4000$, $4000 < x_V^* \leq 7000$ and $7000 \leq x_{VI}^* < 10000$.

Pre-emptive depletion occurs in a smaller area of the parameter space in the Grand Coalition than under free-riding. In general, free-riding reduces the steady state stock and therefore increases the risk of collapse. The increased risk of collapse stimulates pre-emptive depletion for lower values of α.

To build intuition for the above result, we proceed to analyse the differences in harvest functions between cases where pre-emptive depletion occurs and where conservative management occurs. In principle, each stock size can support a certain harvest level in equilibrium, as is usually visualized in the Sustained Yield (SY) curve (Clark, 2010). In this case, the SY curve requires that the following equality holds

$$x = \beta(x - h)^{\alpha}. \tag{5.11}$$

Solving Equation (5.11) for harvest gives the SY curve as follows

$$\mathcal{y}(x) = x - \left(\frac{x}{\beta}\right)^{\frac{1}{\alpha}}. \tag{5.12}$$

The intersection of the SY curve with a given harvest function is thus a steady state, though not necessarily a stable one. The relationship between the SY curve and the harvest function determines whether pre-emptive depletion or conservative management occurs. In Figure 5.2, we provide generic figurative representations of SY curves and harvest functions under conservative management and pre-emptive depletion. We also show stock dynamics in order to aid in interpreting the steady states.

A more detailed analysis of the properties of harvest rules is given in Appendix 5.3. Both harvest functions in Figure 5.2 are linear and have a slope of 1 when $x \leq \gamma\bar{x}$. For these stock sizes, collapse is certain and therefore the entire stock is harvested immediately. A stock size of zero thus satisfies Lemma 5.1 and is a stable steady state under both conservative management and pre-emptive depletion. In the case of pre-emptive depletion, harvest is greater than growth for all stock sizes, and thus $x = 0$ is the only steady state. In the case of conservative management, harvest will be less than growth in some range of the harvest function. Therefore, both an unstable and stable steady state exist in addition to the zero steady state. Lemma 5.1 is satisfied for both stable steady states. Thus, pre-emptive depletion is formally defined as the existence of only one stable steady state, which is zero, and conservative management is defined as the existence of a positive stable steady state in addition to the zero stable steady state.

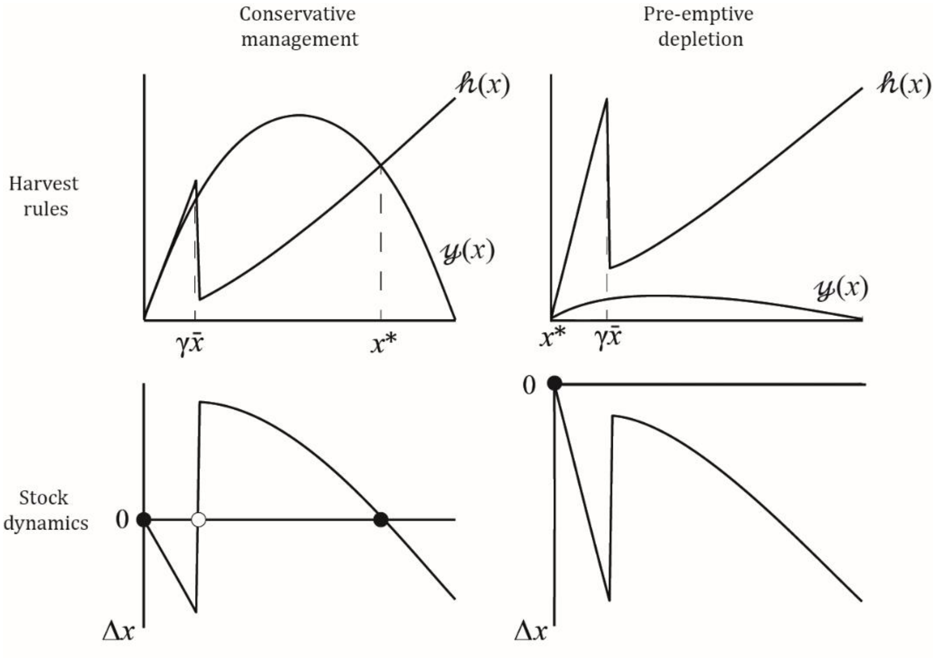

Figure 5.2: Generic representation of harvest functions, stock dynamics and steady states under conservative management and pre-emptive depletion where $\Delta x \equiv x_{t+1} - x_t$. Open circles indicate unstable steady states and closed circles represent stable steady states.

Payoffs are determined in the stable non-zero steady state if it exists (i.e. conservative management is adopted). If it does not exist (i.e. pre-emptive depletion occurs) then payoffs are zero because the steady state is zero. Payoffs are shown in Figure 5.3.

Both panels in Figure 5.3 show the same general pattern. Payoff increases as the discount rate decreases. The marginal effect of the discount rate is very pronounced at low discount rates. Also, payoff decreases as α increases. Recall, high α means that the stock grows more slowly. We also see an area for very high α where payoff is zero due to pre-emptive depletion.

To further the analysis, it is useful to formally define the threshold in the parameter space Θ which determines where payoffs change from non-zero to zero due to pre-emptive depletion – referred to as the depletion threshold. The depletion threshold is given by the borders between Region I and Region II in Figure 5.3. By comparing the relative locations of the depletion thresholds, we can see that pre-emptive depletion occurs for a larger area of Θ in the free-rider case, and the intuition is as follows. Free-riding reduces x^*, which increases the risk of stock collapse, and therefore provides greater incentives to harvest the entire stock in response to the higher risk.

Endogenous risk of stock collapse thus has profound effects on the incentives whether or not to free-ride. Free-riding reduces x^* which increases risk at an increasing rate due to the functional form of Equation (5.4). This means that free-riding leads to increases in risk which are disproportionally larger than the reduction in x^*. In turn, this risk amplification reduces the payoff of free-riding relative to Grand Coalition membership. We term this effect the "risk amplification effect"[7] of free-riding.

In order to analyse the stability of Grand Coalitions for different numbers of players *n*, we calculate payoffs in the free-rider and Grand Coalition cases in Θ space for each *n.* We can then explain how the risk amplification effect and changing numbers of players affect the stability of the Grand Coalition.

The results are shown in Figure 5.4(A) and will be discussed according to the effects of changing α, ρ and n and finally, we discuss the area marked ψ. Figure 5.4(A) shows "stability thresholds" which divide the parameter space into areas where the Grand Coalition is stable and unstable for a given number of players. Concerning α, in general, we see that for a given ρ, as α increases, the Grand Coalition can shift from being unstable to being stable. Higher α means a lower growth rate which in turns results in a lower x^*. Grand Coalitions maintain a higher x^* than coalitions when free-riding occurs. In this way, the risk amplification effect discourages free-riding disproportionally more at lower x^*. Accordingly, increasing α can result in a shift from unstable to stable.

7 The risk amplification effect is similar to the "risk reduction effect" of Ren and Polasky (2014). The risk reduction effect refers to the reduction in endogenous risk when stock increases.

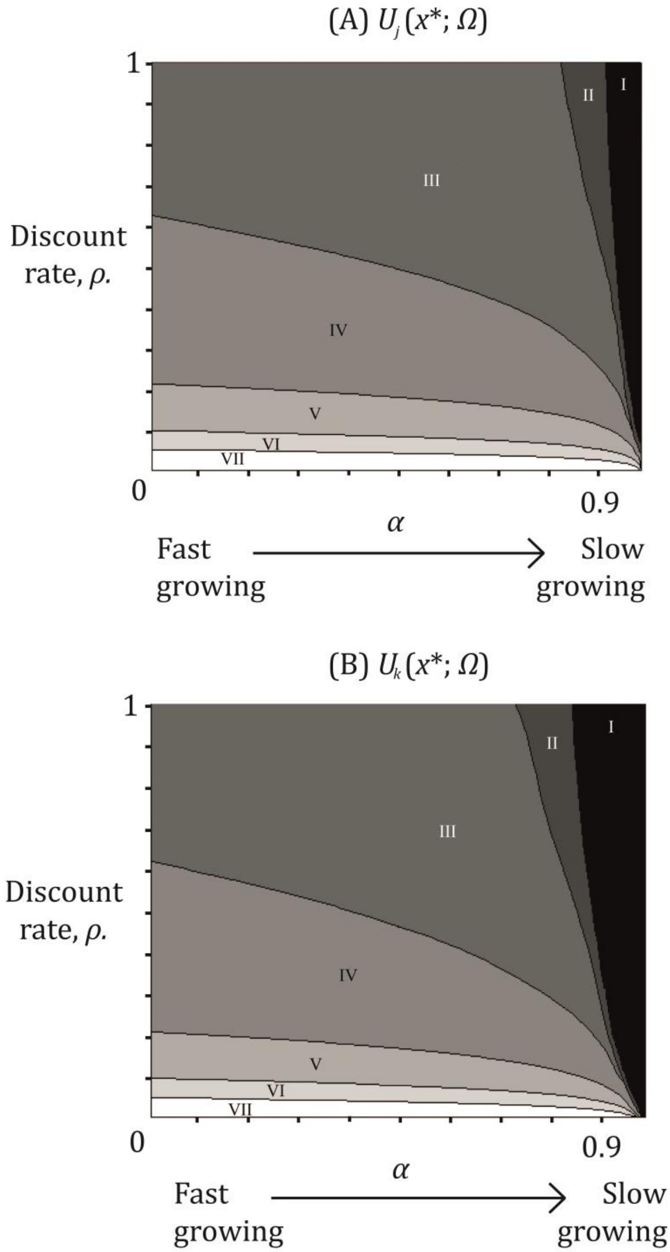

Figure 5.3: Payoffs for a Grand Coalition member (Panel A) and a free-rider (Panel B) in Θ space where $n = 2$, as an example. Using U^{I} to denote the payoff value in Region I, U^{II} to denote the payoff in Region II and so on, the regions are defined as $U^{I} = 0$, $0 < U^{II} \leq 10$, $10 < U^{III} \leq 20$, $20 < U^{IV} \leq 40$, $40 < U^{V} \leq 70$, $70 < U^{VI} \leq 110$ and $U^{VII} > 110$. Region I thus refers to parameterisations for which pre-emptive depletion occurs.

Concerning ρ, in general, we see that for a given α, as ρ increases, the Grand Coalition can shift from being unstable to being stable. This is caused, again, by the risk amplification effect. Higher ρ means that the future is less valuable. Therefore, players prefer current harvest relatively more than future harvest. Accordingly, x^* decreases, the risk amplification effect increases and concurrently, Grand Coalition stability increases. Note also that the risk amplification effect explains the curved shape of the thresholds in Figure 5.4. This is because x^*decreases in ρ and the risk amplification effect increases in x^* at an increasing rate.

Concerning n, we see that in general, increasing the number of players decreases the number of parameterisations for which the Grand Coalition is stable. Grand Coalition stability relies on internalising the externalities of fishing, which are two-fold. Firstly, harvest by one player reduces the amount of fish available for the other player in the future. Secondly harvest by one player increases the risk amplification effect. Grand Coalitions internalize these externalities, but the benefits to each player of doing so are reduced as n increases because the socially optimal catch must be shared by more members. Thus, as n increases, we see a decrease in the number of parameterisations for which the Grand Coalition is stable.

As n increases from 3 to 32, the stability threshold approaches the free-rider depletion threshold in progressively smaller steps. At $n = 32$, the stability threshold is identical to the free-rider depletion threshold. This implies that the decision to free-ride by a single player will result in pre-emptive depletion, which gives a payoff of zero in steady state. As n increases beyond 32, the socially optimal harvest must be shared by more players, but always remains non-zero, while the free-rider payoff remains zero. Hence, the stability threshold does not change for $n \geq 32$. In other words, the stability threshold has thus converged at $n = 32$.

The grey area ψ in Figure 5.4 refers to the subset of Θ for which $U_j(x^*; \Omega) = U_k(x^*; \Omega) = 0$, i.e. where pre-emptive depletion occurs in both the Grand Coalition and free-rider cases. When this is the case, stability is trivial because, when payoffs are evaluated in the steady state, there are no incentives for players to fish either in or out of the coalition. Therefore, Grand Coalitions are non-trivially stable for some values of Θ for all $n > 1$. Grand Coalitions are non-trivially stable for stocks which are slow growing, but not so slow growing that the stock is pre-emptively depleted. This result is due to endogenous risk.

The above analysis of stability raises the question, what determines the location of the stability thresholds? Also, why do the stability thresholds converge at $n = 32$? We will now demonstrate that this finding is sensitive to the parameter γ, which determines the critical escapement level $\gamma\bar{x}$ below which stock collapse is certain. We do so by increasing γ from 0.01 to 0.05. Increasing γ leads to an increase in the probability of collapse for all stock sizes. Therefore, as would be expected, increasing γ leads to an increase in the size of the area of Θ displaying pre-emptive depletion. This, in turn has an effect on the stability thresholds as demonstrated in Figure 5.4(B).

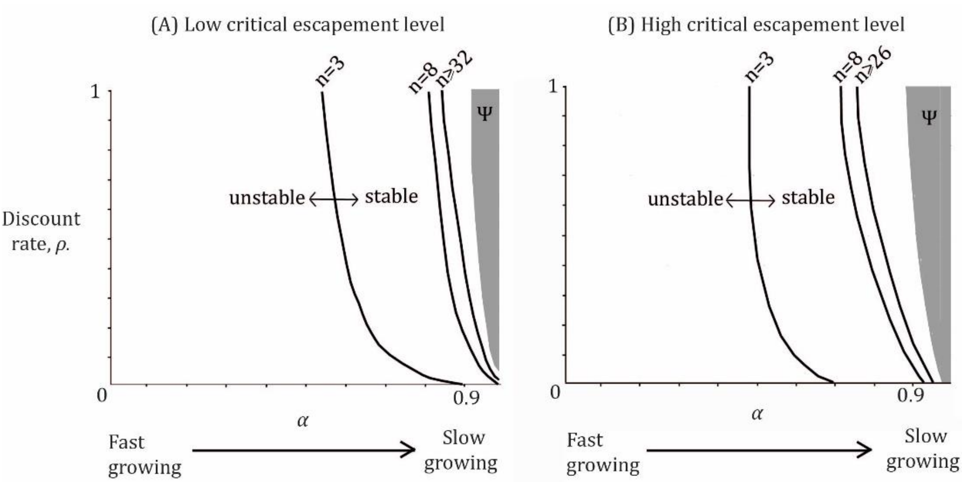

Figure 5.4: Stability thresholds between stable and unstable Grand Coalitions for selected numbers of players in Θ space with $\gamma = 0.01$ (Panel A) and $\gamma = 0.05$ (Panel B). We illustrate the interpretation of the thresholds explicitly for $n = 3$. For all stability thresholds, to the left of the stability threshold, the Grand Coalition is unstable. To the right of the stability threshold, the Grand Coalition is stable. For $n = 2$ the Grand Coalition is stable for all parameters. The grey area, marked ψ, is the subset of Θ for which $U_j(x^*;\Omega) = U_k(x^*;\Omega) = 0$, i.e. where pre-emptive depletion occurs in both the Grand Coalition and free-rider cases.

In comparison to Figure 5.4(A), Figure 5.4(B) shows that an increase in γ from 0.01 to 0.05 shifts all stability thresholds for $n > 2$ to slightly lower values of α and reduces the number of players at which the stability thresholds converge 32 to 26. To recap, the parameter γ determines the critical escapement level below which collapse is certain. Therefore, for stocks with a higher critical escapement level, we observe more pre-emptive depletion. At the same time, cooperation exists for a larger part of the parameter space because there are greater benefits to internalising the risk of stock collapse.

5.5 Including transition payoffs

The previous section has shown that a player receiving strictly positive payoffs in the Grand Coalition can receive a zero payoff upon deviation. This effect drives the possibility for a stable Grand Coalition for any number of players. However, zero payoffs from deviation are due to the assumption that the state of the stock jumps immediately from x^*_{GC} to x^*_{FR}. We therefore relax this assumption and thus account for transition payoffs.

5.5.1 Method for including transition payoffs

In order to test the effects of including transition-payoffs between steady states, we construct a forward model. Generally speaking, forward models take backwardly induced optimal control functions and applies them to a model which runs forward in time in order to fully identify the dynamics of the system. In our case then, for a given element of Θ, the forward model takes the harvest functions, $\hbar_i(x)\ \forall\ i$ corresponding to the free-rider case with a starting stock size of x^*_{GC}. In the first period, we apply the harvest functions to the stock, thus calculating utility and escapement. Escapement and the growth function determine the stock in the next period and the process is repeated until the stock size converges to x^*_{FR}. The time taken for convergence is given by T. The total payoff is given by the instantaneous utility in the Grand Coalition, plus the discounted expected sum of payoffs in the transition, plus the discounted lifetime value of the fisheries in the free-rider steady state. The payoff in the free-rider steady state is reduced as a result of these payoffs being pushed further into the future and the probability that collapse occurs during the transition period. Hence, we adjust the free-rider steady state payoffs by the function $\xi(\rho, T, R)$ where $0 < \xi(\rho, T, R) < 1$ and $R \equiv \prod_{t=1}^{T} r(e_t)$. Thus the total payoff including the transition period is given by

$$u(\hbar_k(x^*_{GC})) + \sum_{t=1}^{T} \delta^t r(e_{t-1})\, u\big(\hbar_k\big(f(e_{t-1})\big)\big) + \xi(\rho, T, R) U_k(x^*_{FR}). \quad (5.13)$$

We can thus repeat the analysis of Section 5.4, accounting for transition payoffs.

5.5.2 Results of including transition payoffs

This section presents the results of including transition payoffs. We find that the maximum size of a stable Grand Coalition is two. We find that stability only exists for a small area of the parameter space, as shown in Figure 5.5.

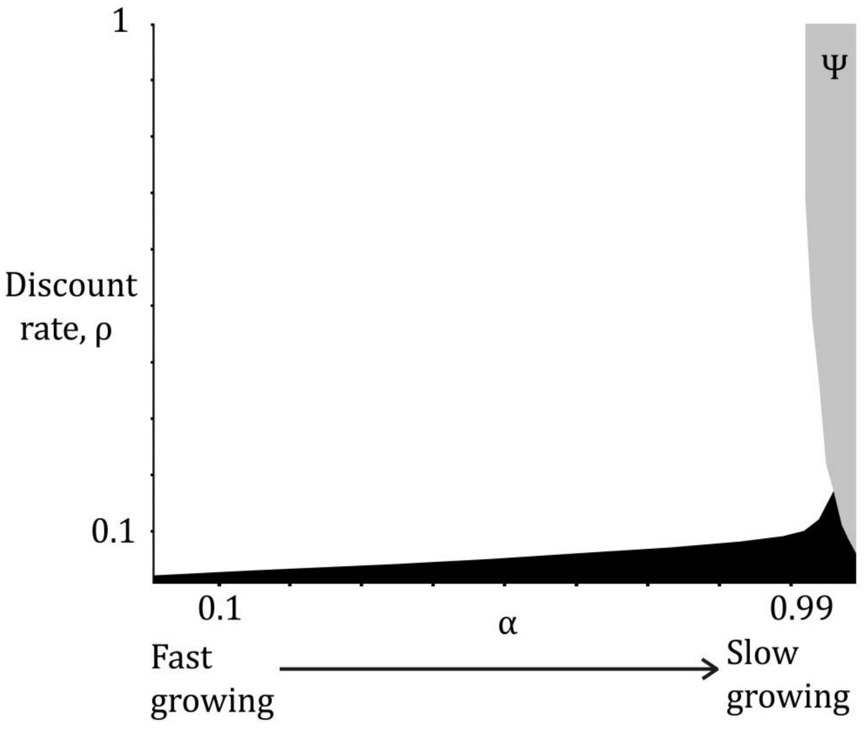

Figure 5.5: Stability of a two-player Grand Coalition in Θ space with transition payoffs where $\gamma = 0.01$. The black area shows elements of Θ for which the 2-player Grand Coalition is stable. The grey area, marked ψ, is the subset of Θ for which $U_j(x^*; \Omega) = U_k(x^*; \Omega) = 0$, i.e. where pre-emptive depletion occurs in both the Grand Coalition and free-rider cases.

The potential for stability is lower when transition payoffs are included due to the increased payoff to deviators available in the transition period. Grand Coalitions of two-players are stable when the discount rate is sufficiently low. Stability can also exist for slightly higher discount rates when α is larger. This conforms closely to the result of Kwon (2006) who studies partial coalitions and finds that partial coalitions of two players are stable only if $\alpha(1+\rho)^{-1}$ is sufficiently high. In our case however, the area of stability has a long tail which encompasses progressively lower discount rates. This is due to endogenous risk. Players who are more concerned with the future prefer Grand Coalition membership due to the reduced risk of stock collapse.

5.6 Conclusions

This study analyses the classic Levhari and Mirman model of the Great Fish War (1980) under an endogenous risk of stock collapse. The objective is to analyse the effects of endogenous risk of stock collapse on the stability of Grand Coalitions. The results of the standard two-stage game show that a risk of stock collapse increases the potential for cooperation. Further, the results show that cooperation can be sustained for any number of players if the stock is sufficiently slow growing, but not so slow growing that exploitation is not sustainable in the long run (i.e. if pre-emptive depletion occurs). Because the potential for cooperation exists for any number of players under an endogenous risk of stock collapse, the Great Fish War becomes a Great Fish Pact.

Further considering the standard two-stage game, the result relating to the growth parameter α has interesting management implications, particularly for deep-water fisheries which are often slow growing (Gordon, 2003). Slow growing stocks are more vulnerable to over-exploitation (Roberts, 2002; Neubauer et al., 2013). This paper supports this proposition for very slow growing stocks. Indeed, the results suggest that the stock would be fished to extinction. However, because Grand Coalitions are stable for slow (not very slow) growing stocks regardless of the number of players, the potential for sustainable management is somewhat less bleak.

Most importantly, the results offer counter-evidence to a long-running implicit conclusion in the literature, namely that the number of players is the most important determinant of potential for stable Grand Coalitions. This study shows that when there are more than a certain number of players, further increases in the number of players has no effect on the area of the parameter space for which the Grand Coalition is stable. The reason for this is that in previous models, increasing the number of players results in lower steady state stocks and these low steady states can be sustained *ad infinitum* with no risk that the stock might collapse. The result presented in this paper regarding the independence of stability from the number of players is entirely the result of relaxing this very common, yet inappropriate, assumption.

In general, this study contributes to the discussion regarding what makes coalitions in fisheries management stable. We observe empirically that coalitions can be stable for large numbers of players but theoretical models tend to be more pessimistic (Hannesson, 2011). Breton and Keoula (2014) refer to this as the "puzzle of small coalitions" and show that larger coalitions can be achieved by using asymmetric players in a game with first mover advantage, thus partly solving the puzzle. Asymmetric players combined with transfer payments can contribute to solving the puzzle (e.g. Pintassilgo et al., 2010), as can the type of solution concept used (e.g. Breton and Keoula, 2014). We have shown that endogenous risk of stock collapse allows the potential for cooperation for any number a players; a possibility which has not yet been identified in the literature. Additionally, this cooperation of any number of players is sustained without the use of transfer payments.

However, our results are sobering in the sense that the potential to seize transition payoffs swamps out the prospects for cooperation and hence, the Great Fish Pact returns to being a Great Fish War. Under the Great Fish Pact, farsightedness, sequential move games and transfer payments are not required to address the puzzle of small coalitions. However, because the Great Fish Pact does not hold if transition payoffs are included, farsightedness, sequential move games and transfer payments still have an important role to play in addressing the puzzle of small coalitions. Further study is required to determine the effects of these assumptions on coalition stability when transition payoffs are included.

Appendix 5.1 Utility function and numerical accuracy

This appendix discusses numerical accuracy with respect to our utility function and interpolation error.

The utility function

The utility function $u(h) = \ln(h)$ is undefined when $h = 0$. A risk of stock collapse implies that harvest level $h = 0$ may occur. We therefore require a utility function which avoids this problem but is sufficiently similar to $\ln(h)$, such that we know that differences in the stability of Grand Coalitions between our model and the LM model can be attributed solely to the presence of endogenous risk. Therefore, we use the utility function $\max(0, \ln(h))$ which is equal to $\ln(h)$ if $h \geq 1$. The range of the utility function is bounded in that it is non-negative and harvest cannot exceed the carrying capacity. We set the carrying capacity at $\bar{x} = 10{,}000$ (by setting $\beta = 10{,}000^{1-\alpha}$) such that our utility function differs from $\ln(h)$ only for a small fraction of its range. Hence, using a large values for $\bar{x}$ ensures h is extremely infrequently between zero and 1 and thus the utility function $\max(0, \ln(h))$ performs, in practise, the same as $\ln(h)$ for all $h > 0$. Where $h = 0$ however, the function $\max(0, \ln(h)) = 0$ and thus performs differently from the original LM model.

In order to evaluate whether our utility function has any effect on the outcome of the model, we numerically solve the deterministic (original) LM model with the utility function $\max(0, \ln(h))$ and evaluate the similarity of the numerically derived harvest function to the analytical solution of the original LM model. We consider the sole-owner case and consider all harvest functions in Θ space, as defined in Section 5.4. To test the similarity, we calculate a standard R^2 statistic to evaluate the extent to which the numerical harvest function can be explained by the analytical harvest function. The results are reported in Figure 5.i.

The results show that the numerical model can recreate analytical results to a high degree of accuracy. It also shows that particular areas of Θ space are more numerically challenging to estimate than others. The location of the area of largest error, consisting of the union of Regions III, IV and the larger of the two areas marked as Region II is particularly important. The accuracy of stability thresholds in this particular region is therefore somewhat reduced. Overall, the high R^2 values confer confidence in the accuracy of the numerical method, thus supporting our use of the $\max(0, \ln(h))$ utility function. While numerical accuracy is high, we cannot be sure whether the inaccuracy is due to the utility function or due to interpolation error.

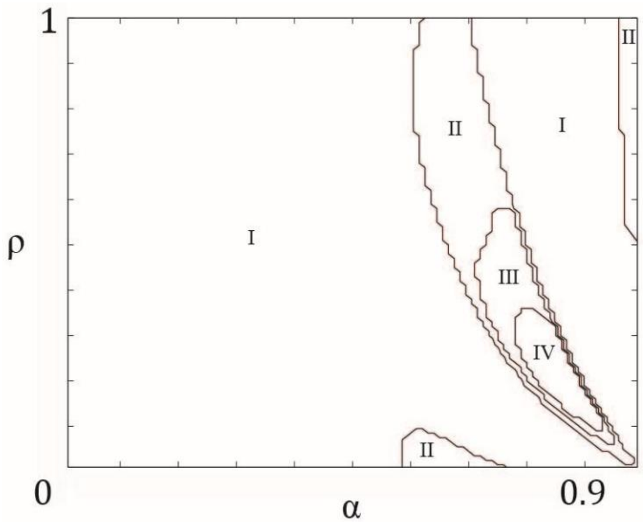

Figure 5.i: R^2 statistics in Θ space determining the accuracy of a numerical 1-player deterministic LM model. Each region (I through IV) represents a range of R^2 statistics. Using R_I^2 to denote the R^2 value in Region I, R_{II}^2 to denote the R^2 in Region II and so on, the regions are defined as $0.9998 < R_I^2 \leq 1$, $0.9994 < R_{II}^2 \leq 0.9998$, $0.9990 < R_{III}^2 \leq 0.9994$ and $0.9984 \leq R_{IV}^2 \leq 0.9990$.

Interpolation error

Error in the model may also result from interpolation error. Interpolation error results from the discretised state space. We set the state space as $x_t \in X = [0, 1000, 2000, \dots, 10000]$. In the case that the steady state x^* is less than 1000, the model increases the number of elements in the state space in order to more accurately identify the steady state. Discretisation of the state space means that $U_i(x_{t+1}; \Theta)$ is only known for each element of X. Almost always, x_{t+1} is not an element of X and therefore we use interpolation to estimate $U_i(x_{t+1}; \Omega)$. Error in interpolation means that future value in the value function deviates from its true value and this results in deviations of the harvest function from their true form. The effect of any interpolation error on the harvest function is reduced when the value function is determined by instantaneous utility relatively more than future value. Future utility has relatively less of an effect on the value function when the discount rate is high. This can be seen in Figure 5.i, where error in the numerically estimated LM model tends to be higher for lower discount rates. This suggests that some of the inaccuracy in Figure 5.i is due to interpolation error. Finally, it is useful to note that future value also has relatively less of an effect on the value function if future value is reduced due to endogenous risk. Therefore, endogenous risk has the side effect of reducing the effect of interpolation error in our model, thus increasing our confidence in the results.

Appendix 5.2 Deriving the envelope condition

Deriving the envelope condition requires determining the first order conditions of Equation (5.6) with respect to harvest and stock. The first order condition w.r.t. harvest h_i is given by

$$\frac{\partial V_i}{\partial h_i} = \frac{du}{dh_i} + \frac{\partial r}{\partial e}\frac{\partial e}{\partial h_i}\delta V_i\big(f(e)\big) + r(e)\delta\frac{\partial V_i}{\partial f}\frac{\partial f}{\partial e}\frac{\partial e}{\partial h_i} = 0. \tag{5.i}$$

From Equation (5.1) it follows that $\frac{\partial e}{\partial h_i} = -1$. Equation (5.i) therefore simplifies to

$$\frac{du}{dh_i} = \frac{\partial r}{\partial e}\delta V_i\big(f(e)\big) + r(e)\delta\frac{\partial V_i}{\partial f}\frac{\partial f}{\partial e}. \tag{5.ii}$$

Substituting $\hbar_i(x)$ for h_i and $\hbar_{-i}(x)$ for h_{-i} in Equation (5.6) such that $e \equiv x - \hbar_i(x) - \hbar_{-i}(x)$ and differentiating Equation (5.6) w.r.t. x gives

$$\frac{\partial V_i}{\partial x} = \frac{\partial u}{\partial \hbar_i}\frac{\partial \hbar_i}{\partial x} + \delta\left[\frac{\partial r}{\partial e}\left(1 - \frac{\partial \hbar_i}{\partial x} - \frac{\partial \hbar_{-i}}{\partial x}\right)V_i\big(f(e)\big) + r(e)\frac{\partial V_i}{\partial f}\frac{\partial f}{\partial e}\left(1 - \frac{\partial \hbar_i}{\partial x} - \frac{\partial \hbar_{-i}}{\partial x}\right)\right]. \tag{5.iii}$$

Substituting $\frac{\partial u}{\partial \hbar_i}$ in the above with $\frac{du}{dh_i}$, then substituting $\frac{du}{dh_i}$ with the right hand side of Equation (5.ii), then simplifying gives

$$\frac{\partial V_i}{\partial x} = \left(1 - \frac{\partial \hbar_{-i}}{\partial x}\right)\left(\frac{\partial r}{\partial e}\delta V_i\big(f(e)\big) + r(e)\delta\frac{\partial V_i}{\partial f}\frac{\partial f}{\partial e}\right). \tag{5.iv}$$

Equation (5.ii) is substituted into Equation (5.iv), thus giving the envelope condition

$$\frac{\partial V_i}{\partial x} = \left(1 - \frac{\partial \hbar_{-i}}{\partial x}\right)\frac{du}{dh_i}. \tag{5.v}$$

Appendix 5.3 The properties of harvest functions.

This appendix discusses the properties of harvest functions in terms their monotonicity and their form in relation to point $x = \gamma\bar{x}$. We present the harvest function for $\alpha = 0.99$, $\rho = 0.01$, $\gamma = 0.02$ and $x \in X = [0, 100, 200, \dots, 10000]$ in the sole-owner case in Figure 5.ii. We find that harvest functions are not necessarily monotonic, which is expected given that it may be optimal to harvest the entire stock at low resource levels. Furthermore, Figure 5.ii also indicates that the harvest function does not necessarily attain a (local) maximum where $x = \gamma\bar{x}$. It can still be optimal to harvest all of the stock immediately, even if collapse is not certain. Thus the maximum of the harvest function occurs at $x = 300$ whereas $\gamma\bar{x} = 200$.

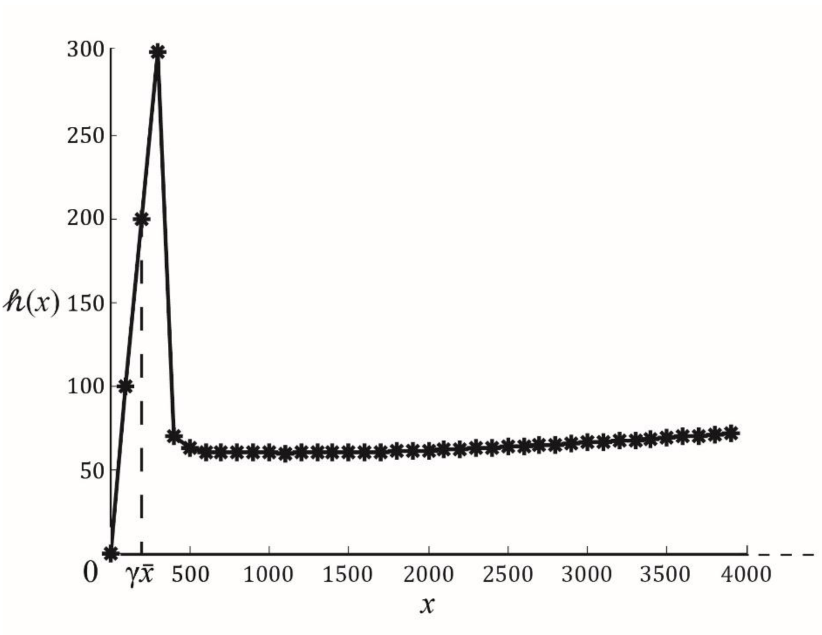

Figure 5.ii: The harvest function where $\alpha = 0.99$, $\rho = 0.01$ and $\gamma = 0.02$ demonstrates non-monotonicity. Stars are actual data points. The harvest function is defined up to $x = 10000$. We limit the x axis to focus on lower values of x.

6 Synthesis

6.1 Summary and answers to research questions

Invasive alien species spreading across the world and fish stocks responding to changes in climate and overfishing: these have been the topics of this thesis. The economic aspects of these topics are the optimality of invasive species management strategies, as considered in Part A, and the stability of International Fisheries Agreements (IFAs), as considered in Part B. The chapters in Part A separately consider the prevention and control of invasive species. Prevention is studied in the context of Ballast Water Management (BWM) and control is studied in a spatially explicit metapopulation model. The chapters in Part B separately consider the implications of changing fish stock location and an endogenously determined risk of fish stock collapse on cooperative fisheries agreements. Each chapter addresses a specific research question, as was outlined in Chapter 1.

6.1.1 Ballast Water Management

In Chapter 2, we constructed a model to analyse the optimality of BWM standards, as are set in the BWM convention (IMO, 2004). The model focused on specific features of BWM in order to understand their implications for optimal management. To do so, we constructed a model of an irreversible invasion, because irreversibility is the most appropriate general assumption for marine Invasive Alien Species (IAS) (Vitousek et al., 1997; Parker et al., 1999). The model consists principally of the expected damage function, which gives the relationship between the ballast water treatment standard and the present expected value of damage from the invasion. The damage function depends on the hazard function, which determines how the hazard rate of invasion establishment depends on the intensity of ballast water treatment. The hazard rate function includes network effects, a Minimum Viable Population (MVP) and the possibility for an Allee effect in order to answer the following research question:

Research Question 1: *What are the implications of minimum viable populations and Allee effects for optimal ballast water management standards?*

We address this research question by analysing the marginal properties of the expected damage function in terms of the conditions under which the standard adopted in the BWM convention can be optimal. We find that a discontinuous marginal expected damage function is required for the current standard to be optimal. We therefore proceed to analyse whether MVPs or Allee effects could result in discontinuity in the marginal expected damage function. We find that if the hazard function approaches zero at the limit where treatment effort approaches the level of effort required to achieve the MVP, then the marginal damage function is continuous and the current BWM standard cannot be optimal. We find that this can only occur in the presence of an Allee effect. Therefore, our model shows that an Allee effect is a critical determinant of whether the current BWM standard can be optimal.

6.1.2 Invasive species control

In Chapter 3, we considered the control problem. This concerns the situation where the invasion has already become established outside of its native range. There has been a significant amount of literature on this topic. The key issue is how to optimally control the invasion over time (e.g. Eiswerth and Johnson, 2002) and space (e.g. Sharov, 2004), in addition to other important issues such as optimization of monitoring regimes to detect invasions (e.g. Mehta et al., 2007) and how to deal with the complexities imposed by invasions spreading across jurisdictional or private property boundaries (e.g. Kovacs et al., 2014 and Zhang et al., 2010, respectively). We focus on the spatial aspects of control within a non-native range in the case that the whole of the non-native range falls under a single jurisdiction, e.g. a single nation. We consider a non-native range divided into patches (cf. Salinas et al., 2005 and Burnett et al., 2007). Each patch can contain a population of the invasive species such that the invasive population size can vary between patches (as in Burnett et al., 2007). We allow for a management intervention which reduces the population size within a selection of patches by any amount. This is an extension of modelling approaches which consider patches which either do not allow for varying invasive species populations within areas (i.e. areas are modelled in binary terms: either invaded or not invaded) or which restrict removal of invasions in a given area to complete eradication only (Carrasco et al., 2010a; Finnoff et al., 2010a; Epanchin-Niell and Wilen, 2012; Epanchin-Niell et al., 2012). In addition to removal, we allow for a second control option, termed "containment", which reduces the probability of spread between patches without affecting the population size within patches, as in Sharov (2004). Optimal application of removal and containment is derived for the case where the non-native range can be accurately described as a one-dimensional landscape, or line-network (Chadès et al., 2011). This is relevant for invasive species spreading along coastlines, and also through riparian systems. The model is used to answer the following research question.

Research Question 2: *What are the implications for optimal spatial control of invasive species when the invasive population is modelled with varying invasive population sizes within patches?*

Our results identify three implications for management. The first pertains to the timing of containment interventions. When patches can only be either invaded or non-invaded, if it is optimal to apply containment to reduce the probability spread from that particular patch, then it will always be optimal to do so immediately. When the population within a patch can vary, it becomes possible to time the application of containment depending on the size of the population within that patch. Greater population size means that the probability of spread is greater. Under our assumption that containment reduces the probability of spread by a fixed factor, containment results in a greater reduction in the probability of spread to non-invaded patches when the invasive population in adjacent invaded patches is larger. Accordingly, it may be more efficient to wait for the population to reach a certain size before implementing containment, or in other words, the timing of containment depends on population size. This is a different result from the case when invasive population size does not vary within patches over time, as in Epanchin-Niell and Wilen (2012). In that case, containment can only be applied as soon as the patch is invaded.

The second implication pertains to the policies aimed at slowing the spread of the invasion. In some cases, it is optimal to let the entire non-native range become invaded, but it is optimal to slow down the process of spread. Doing so pushes the damages of a fully invaded non-native range further into the discounted future. To optimally slow the spread, we identify the possibility for four distinct interventions that can be applied to one patch depending on the size of the population in that patch. Let us categorize the size of the invasive population in a patch as either very small, small, large or very large. Very small invasive populations are immediately entirely removed. Then, for small population sizes, the optimal policy is to do nothing and to let the invasive population grow. For large population sizes, containment is implemented to reduce the probability of spread. Finally, for very large population sizes, it becomes optimal to supplement the containment intervention with removal within the invaded patch. This provides more detailed specifications of management aimed at slowing the spread of invasions than in, for example, Sharov (2004).

The third implication pertains to policies aimed at stopping the spread of the invasion such that spread through the entire non-native range is prevented. In the setting in Chapter 3, containment is not perfectly effective; it can only reduce the probability of spread. In practice, this means that containment can only slow the invasion down. Despite this, stopping the spread can still be optimal if containment and removal policies are combined. Containment can be used to reduce the probability of spread into a non-invaded patch and then, any invasion which does occur can be immediately removed. This demonstrates that stopping the spread can still be an optimal management strategy, even if containment is not perfectly effective, as in Epanchin-Niell and Wilen (2012).

6.1.3 Changing fish stock location

In Chapter 4, we considered the effects of changing fish stock location on Grand Coalitions in fisheries agreements. We employed the Gordon-Schaefer model of fisheries to analyse the stability of Grand Coalitions and thus build on the work of Lindroos (2008), Pintassilgo and Lindroos (2008) and Pintassilgo et al. (2010) by examining changing fish stock location. We conceptualize the entire fish stock as existing at a single point in space which may change. The nations (players) in the fishery agreement relating to that fish stock are conceptualized as single points at fixed locations. When fish stock location changes, the distances between the fish stock point and the players' locations changes.

We focus on how the effects of changing stock location on the stability of IFAs are determined by the type of stability concept used. Stability concepts embody behavioural assumptions about how players respond to the cooperative choices of others, or alternatively, the "conjecture" that a given player makes about how other players will respond to changes in the cooperative choice of that given player. In the literature, it is often assumed that players adopt shortsighted conjectures, more generally referred to as Nash conjectures (e.g. Kwon, 2006; Pintassilgo and Lindroos, 2008; Pintassilgo et al., 2010; Long and Flaaten, 2011; Breton and Keoula, 2014). Under shortsighted conjectures, a given player assumes that other players will not change their cooperative choices in response to a change in the cooperative choice of that given player. An alternative stability concept is based on farsighted conjectures (Chwe, 1994), whereby a given player assumes that other players will change their cooperative choices, if such a change is beneficial, in response to a change in the cooperative choice of that given player. Farsighted conjectures have been analysed in fisheries games by Breton and Keoula (2012). Our analysis focused on the following research question.

Research Question 3: *To what extent can farsightedness stabilize IFAs in the face of changing fish stock location?*

To answer the question, we define and employ a new variant of the Farsighted stability concept (Chwe, 1994), which we term Farsighted Downwards Stability (FDS). Because the FDS concept is new, and to establish a baseline, we begin by analytically determining the properties of FDS in the simplest setting, i.e. where players are symmetric and fish stock location is constant. Where n is the number of players, we find that the Grand Coalition displays FDS for $n \leq 4$. We consider asymmetry in the costs of fishing and changing fish stock location using sensitivity analysis for the cases of three and four players. We thus test Grand Coalition stability for many different fish stock locations. We find that Grand Coalitions are more likely to be stable under the FDS stability concept than under internal stability concept. However, while stable Grand Coalitions are more likely under the FDS stability concept, changes in whether a Grand Coalition is stable due to changing fish stock location are also more likely. In this way, the use of the FDS stability concept increases stability, but also increases the sensitivity of stability to changes in fish stock location. Overall then, the results suggest that the stability of Grand Coalitions can be improved if players employ farsighted conjectures in deciding whether or not to deviate. Policy mechanisms which encourage further deviation will force players to make farsighted conjectures about the effects of their deviation, which in turn, increases the stability of the Grand Coalition.

6.1.4 Endogenous risk of fish stock collapse

In Chapter 5, we considered the effects of an endogenous risk of fish stock collapse on the stability of Grand Coalitions. Unlike in Chapter 4, we employed the Levhari and Mirman (LM) (1980) model of fisheries. In order to model a risk of fish stock collapse, we adapted the LM model such that the lower the fish stock size, the greater the risk of a fish stock collapse. Players are assumed to be symmetric. We addressed the following research question:

***Research Question 4:** What are the implications of an endogenously determined risk of fish stock collapse on cooperation in IFAs?*

We find that the implications of an endogenously determined risk of fish stock collapse depend heavily on the assumption, implicit in the standard two-stage game, that payoffs are evaluated at the steady state fish stock sizes corresponding to the size of the coalition. We find that when standard assumptions are followed such that payoffs are evaluated at steady states, then an endogenously determined risk of stock collapse increases the number of players for which a Grand Coalition is stable. This is a similar result to those obtained by Nikuiya et al. (2014), Ren and Polasky (2014) and Sakamoto (2014). In comparison to deterministic models, non-cooperation leads to not only a lower stock of fish, but also a greater risk that the fish stock will collapse in the future. Endogenous risk of fish stock collapse thus increases the stability of Grand Coalitions. In fact, the addition of endogenous risk into the model means that, for specific discount and growth rates, a Grand Coalition of any number of players is stable. This is because it is possible that the Grand Coalition is the only coalition structure for which it is optimal to attempt to maintain a non-zero steady state fish stock. If any deviation from the Grand Coalition occurs, then the fish stock will be intentionally harvested to zero. Thus, under the assumption that payoffs are calculated in the steady state, the choice for members of a Grand Coalition is between receiving some payoff (although this may be extremely small if there are many players in the coalition) or receiving no payoff at all. Accordingly, the Grand Coalition is stable for any number of players.

An important part of the incentive to stay in a Grand Coalition under the assumptions implicit in a standard two-stage game, is that no payoffs are received in the process of intentionally harvesting the fish stock to zero (transition payoffs). Transition payoffs have been shown to be important in understanding the effects of endogenous risk (Sakamoto, 2014). We include transition payoffs and find that the stability of Grand Coalitions drastically reduces. Grand Coalitions are never stable for more than two players, and are only stable for two players for specific parameterizations. This is almost exactly the same result as that in the absence of endogenous risk (Kwon, 2006). From a theoretical perspective then, the answer to the research question is that the implications depend heavily on the assumptions about transition payoffs. If transition payoffs are excluded then endogenous risk results in dramatic increases in the potential for cooperation, but this effect disappears if transition payoffs are included.

6.2. Evaluation

In this section, we discuss the appropriateness of the methodologies employed in this thesis to answer the research questions.

6.2.1 Evaluation of Part A

In Chapter 2, we employ an analytical modelling approach. The model contains no actual parameterizations for the costs of, and reduced expected damages resulting from BWM, nor does it use explicit functional forms for the hazard function. As such, we do not suggest what the actual optimal BWM standard is. This modelling approach is chosen due to the complexity and large data requirements of the problem. Information requirements are large and complex for several reasons. Firstly, IAS is a global issue. Data collection must therefore deal with the challenges of aggregating local, national or regional data which may not use consistent formats or definitions, as well as the challenges of international collaboration and data quality assurance (Ricciardi et al., 2000; Crall et al., 2006). Secondly, any data set which succeeds in addressing the first issue must contain a great amount of data. This is due to the large number of species which are currently, or may become, invasive. In Molnar et al. (2008), aggregation of data sets identified 329 invasive species, ranging from fish to plants, and algae to molluscs and crustaceans. For each species, diverse information must be collected. In general, non-exclusive terms, data should include ecological and economic impact, geographic extent, invasive potential and possible management options including their costs and efficacy.

Due to the large and complex data requirements, it would be very challenging to produce a parameterized model to estimate an optimal standard. Our method instead provides general insights into the determinants of optimal BWM policies. These insights can be used to evaluate current BWM standards and inform future research priorities. For example, our results point to the important role of Allee effects. Empirically identifying and measuring the severity of Allee effects is challenging for several reasons (Sakai et al., 2001), but given the fundamental implications of Allee effects for optimal standards, a better understanding of Allee effects in IAS is very useful. If, for example, research suggests that the most harmful IAS are unlikely to display Allee effects then this constitutes an argument that current standards are optimal. Additionally, the results provide economic support for the widely held understanding that better information on MVPs would be extremely valuable (Gollasch et al., 2007). In summary, while our method is not suitable for estimating an actual optimal standard, it can still provide useful insights for decision making in terms of prioritizing future research and understanding the implications of different features of BWM for the optimality of standards.

In Chapter 3, we employ dynamic programming, which has been frequently applied to the problem of controlling IAS (e.g. Blackwood et al., 2010; Cacho et al., 2008; Chalak-Haghighi et al., 2008 and Haight and Polasky, 2010). Dynamic programming suffers from two related drawbacks. Firstly, dynamic programming problems may not be analytically tractable. This means that closed-form solutions cannot be derived, which limits understanding of the drivers of optimal policies. Accordingly, numerical analysis is adopted. Numerical analysis of dynamic programs requires that an understanding of model outcomes be developed by experimenting with different parameterizations (e.g. Epanchin-Niell and Wilen, 2012), whereas analysis of closed-form solutions provides general insights for all parameterizations (e.g. Blackwood et al., 2010). Nonetheless, through experimentation with parameterizations, a detailed understanding of the drivers of model outcomes can be developed (as in Sanchirico et al., 2010 and Epanchin-Niell and Wilen, 2012). The second problem of dynamic programming is the curse of dimensionality (Bellman, 1957), which is essentially a problem of computational burden. Computational burden is potentially significant in our modelling approach because of the spatial set-up of the model, the stochastic spread process and the use of two management options (removal and containment). In Chapter 3, computational burden is minimised by identifying and adopting the simplest relevant spatial set-up and careful programming. Accordingly, the majority of our analysis is carried out in the simplest relevant spatial setup whereby space is divided into two patches. Two-patch models are common in the literature and have been shown to produce useful insights (Salinas et al., 2005; Rowthorn et al., 2009; Sanchirico et al., 2010). However, applied models are likely to suffer from the curse of dimensionality because two patches will be insufficiently detailed. Therefore, our method is most suitable to provide generalized management insights.

6.2.2 Evaluation of Part B

In Chapters 4 and 5, we employ game theory. Game theory, being concerned with the strategic interaction between agents, is chosen as a modelling approach because it allows for nations to formulate their decisions, not only in terms of the state of the system (the size of the fish stock), but also in terms of the decisions of other players. We employ coalition theory as a lens through which to analyse the common-pool resource problem. The advantage of coalition theory is that it gives us further insight into the common-pool resource problem. Many studies consider the common-pool resource problem in terms of the situations in which it is more or less severe. For example, Sakamoto (2014) shows that endogenous risk may lead to either more or less resource exploitation. This means that the common-pool resource problem may be worsened or ameliorated under endogenous risk. The number of players for which the Grand Coalition is stable provides a concrete measure of the severity of the common-pool resource problem. If a Grand Coalition exists for a given number of players then the common-pool resource problem does not exist for that number of players. Because an increase in the number of players generally decreases the potential for stable Grand Coalitions in fisheries management, the greater the number of players for which a Grand Coalition can be sustained, the less severe the common-pool resource problem. Accordingly, coalition theory is a popular means to analyse the common-pool resource problem (e.g. Breton and Keoula, 2012; Rettieva, 2012; Breton and Keoula, 2014). Further, coalition theory provides a useful bridge to real-world issues. Cooperative agreements play an important role in dealing with various common-pool resource problems, the most significant of which being climate change (Carraro and Siniscalco, 1998).

Specific to Chapter 4, several of the most important modelling assumptions should be mentioned for evaluative purposes. These relate to our assumptions about changing fish stock location, and to the assumptions employed in the FDS concept. In order to analyse changing fish stock location, we conceptualized the fish stock as existing at a single point in space, the location of which is defined relative to the fishing nations, which are also single points in space. We assumed that only the fish stock location point changes while other parameters are constant. However, climate change may warm more northerly waters such that they are suitable for a fish stock without warming the southerly waters so much that they become unsuitable (Cheung et al., 2009). This means that it may be more accurate to consider change in fish stock "range" rather than change in location. If fish stock range changes then it is no longer accurate to consider the fish stock as existing at a single point. Instead, describing the location of the fish stock would require consideration of the dimensions of the fish stock (i.e. the size and shape of the area which it occupies). In addition to changing fish stock location and range, the carrying capacity of the fish stock may also change, i.e. climate change alters the ecosystem such that it can support a greater fish stock. In summary, to determine the implications of changing stock location in specific cases, one must also account for possible changes in fish stock range and carrying capacity. These points are important to take into account when constructing more applied models.

The outcomes, in terms of coalition stability, resulting from the FDS concept depend heavily on Definition 4.3. Definition 4.3 is an assumption designed to deal with the following problem. The benefit of being in the Grand Coalition depends on the payoff which a player would receive outside of the coalition. However, there may be several coalition structures which could result from a deviation from the Grand

Coalition and therefore several different possible payoffs. The stability of the Grand Coalition depends on how we use the information about the set of possible coalition structures, which might result from a deviation from the Grand Coalition. Definition 4.3 assumes that players use this information in an optimistic way by assuming that they will always reach the most profitable coalition after a deviation from the Grand Coalition. In turn, this means that the stability of the Grand Coalition is lower than it would be under other assumptions such as taking the mean of the payoffs in all coalitions which could result from a deviation from the Grand Coalition. This suggests that, not only do farsighted conjectures help to stabilize Grand Coalitions in the face of changing fish stock location, but also, a healthy dose of pessimism regarding the payoffs after deviation from the Grand Coalition helps too. Pessimism can therefore also play a role in ensuring stable Grand Coalitions under fish stock location change, although this was not explicitly discussed in Chapter 4.

In Chapter 5, we employ the Levhari and Mirman (LM) model (Levhari and Mirman, 1980) to analyse the implications of an endogenously determined risk of fish stock collapse for the stability of Grand Coalitions. The most important methodological limitation of Chapter 5 is that we adopt a numerical approach. The logarithmic utility function in the LM model is unsuited to numerical approaches because it returns a computational problematic utility of minus infinity if harvest is zero. We deal with this by assuming that utility is zero when harvest is zero and by scaling the model to very large stock sizes such that harvest is never between zero and one. However, the changes to the utility function result in another problem. In order to understand the effects of endogenous risk on Grand Coalition stability in the LM model, we need to keep everything else in the model constant so that any differences in stability relative to the deterministic LM model can be attributed solely to the inclusion of the endogenous risk of fish stock collapse. We therefore test the effect of the change to the utility function on the ability of the model to reproduce the analytically derived results from the original LM model. We find that the changes to the utility function have a negligible effect. This allows us to isolate the effects of endogenous risk from changes in the utility function.

Another problem with numerical approaches is that it may be less clear what drives the results. In order to help understand what drives the results, we present results for the entire parameter space. The importance of this approach can be illustrated in the context of Chapter 3. In Chapter 3, we present results for a very limited set of parameterizations. This is principally because many variables in Chapter 3 are only bounded from below, meaning that the range of possible parameterizations is infinite. Not only is the range infinite, but there are more variables in total. In Chapter 5, all variables are bounded from both above and below and thus have finite ranges. Further, there are only three variables. This allows us to present results across much of the parameter space in a way which is reasonably easy to interpret. Most importantly, adopting a numerical approach has allowed us to let the pertinence of possible research questions dictate our line of enquiry, rather than allowing our research questions to be dictated by the confines of analytical tractability.

6.3 Reflections and policy relevance

This section provides some broader reflections about the research in this thesis. Our discussion of Part A focuses on how the results can be used in decision making and our discussion of Part B deals with broader methodological concerns relating to the economics of IFAs. Reflections for both parts end by considering the institutional perspective.

6.3.1 Discussion of Part A

As discussed in the previous section, Chapters 2 and 3 provide generalized management insights for the prevention and control, respectively, of invasions. We also justified the approach of producing generalized management insights. In order to understand the value of these generalized management insights, we must consider the settings in which decisions are made. In the context of BWM, management decisions have, up to now, been made deliberatively (Gollasch et al., 2007). An alternative to a deliberative approach is to formally calculate optimal ballast water treatment standards. As discussed, due to the complexity and magnitude of information requirements, a formal calculation is unlikely to be feasible. Proponents of evidence-based policy making may consider formally calculated ballast water treatment standards to be the best approach for deriving standards. However, literature in the domain of public administration is often critical of evidenced-based policy making (Sanderson, 2006). This literature argues that policy making is not a purely technical process (Dryzek, 1990; Fischer, 1990; Schwandt, 1997). The literature lends support to a process based on dialogue and argument (Majone, 1989). This approach is rooted in the concept of Aristotelian deliberation, whose influence can be traced through to the more modern Rawlsian concept of "reflective equilibrium" (Rawls, 1972), whereby, through debate, the judgements of stakeholders come to coincide. However, despite the rejection of the idea of policy making as a technical process, proponents of more deliberative approaches emphasize the importance of sound "scientific knowledge" (Schwandt, 2000) and "factual statements" (Majone, 1989) upon which deliberative decision making can be based. A sound theoretical understanding of the implications of the features of BWM (i.e. Allee effects) for the optimality of standards provides such "scientific knowledge" and "factual statements". These can be employed in decisions which relate not only to the strictness of the standards themselves, but also to the chosen metric of the standards and the details of the policy, such as the percentage of global shipping which is required to be covered by the agreement before it comes into force. In particular, a sound theoretical understanding of the features of BWM will be useful in formulating next steps after the current BWM convention comes into force. Such next steps should be based first on an evaluation of the current BWM convention, for which Chapter 2 provides some useful insights, principally regarding the role of Allee effects.

In Chapter 3, our contribution to decision making is achieved by establishing new generalized management insights for decision making. The value of generalized management insights is justified in the context of the challenges of formally deriving management strategies, and further, this approach is justified in the context of existing literature. In practice, when an invasion arrives in a new area and is detected, managers are faced with a series of very challenging questions. Addressing these questions can be aided by theoretically well-founded generalized management insights. The first of these questions, as addressed by Sims and Finnoff (2013), is whether or not managers should delay making decisions to gain more information about the invasive species by observing, for example, their rate of spread. Sims and Finnoff demonstrate that, while waiting and seeing is a popular option, it is rarely optimal. The generalized management insight is therefore, that if managers decided to wait and see, they should have good reason to do so (Sims and Finnoff provide such reasons). This generalized management insight is valuable to decision makers, even if they do not possess the data or the ability to formally determine whether they should act now, or wait and see. In reality, data will always be lacking to inform the formal determination of management policies and there will always be uncertainty regarding different aspects of the invasion. In this context, generalized management insights are particularly valuable. Chapter 3 contributes generalized management insights in the context of the spatially explicit metapopulation model where space is divided into patches (as in Burnett et al., 2007; Carrasco et al., 2010a; Finnoff et al., 2010a; Epanchin-Niell and Wilen, 2012; Epanchin-Niell et al., 2012 and Chadès et al., 2011). Our model shows that, under the assumption that containment reduces the probability of spread by a fixed factor, it may be optimal to delay the application of containment until the invasive population in a given patch has reached a certain size. The generalized management insight is therefore that implementing measures to reduce the spread of an invasion to another patch need to be timed according to the size of the population of the invasive in adjacent patches.

The extent to which generalized management insights are useful depends on the institutional setting in which policies are made and implemented. It is therefore pertinent to quickly discuss this setting in the European context, given the focus of research within the VECTORS project. The most recent EU legislation on the topic of invasive species is Regulation 1143/2014 of 22 October 2014 on the prevention and management of the introduction and spread of invasive alien species. Thanks to this regulation, it is now EU policy to "take all appropriate steps to encourage Member States to ratify [the BWM] Convention" (Article 5). Unfortunately, this regulation has not been greeted warmly by invasion ecologists (Genovesi et al., 2015) because this action is less concrete than that in the draft proposal (COM/2013/0620). In the draft proposal, an action plan for all member states was proposed which included actions based on an evaluation of the costs and benefits of the BWM convention (Article 11), but this was not included in the final version. There is no information available on the reason for the removal of this policy from the draft version. We can speculate that it may be because of the impracticality of estimating the costs and benefits of the BWM convention. Alternatively, it may be because of the EU's apparent unwillingness to take unilateral action on BWM relative to other countries. Australia, Canada and New Zealand all have well established domestic ballast water regulatory regimes, as do the states of California and New York, in addition to the United States' federal regulations (Albert et al., 2013). Therefore, while many EU countries have ratified the BWM convention, there is, relatively speaking, less emphasis on prevention in Europe. Fortunately, Regulation 1143/2014 represents an important step in a unified policy on management of established invasions (Genovesi et al., 2015). In general terms, the strategy in Regulation 1143/2014 is to detect new invasions early such that they can be eradicated, with "containment and control" measures employed only if eradication is unfeasible. The regulation refers to both terrestrial and marine invasive species. Given the difficulties of eradicating marine invasions and the relatively lower emphasis on BWM in Europe, containment and control measures, such as those studied in Chapter 3 of this thesis, are likely to remain important mechanisms for the management of marine IAS in Europe. This is especially the case given Descriptor 2 of Good Environmental Status under the Marine Strategy Framework Directive (2008/56/EC) which states that "non-indigenous species introduced by human activities are at levels that do not adversely alter the ecosystems". Again, this suggests that the focus of European policy is on managing existing invasions rather than preventing them.

6.3.2 Discussion of Part B

The puzzle of small coalitions (Breton and Keoula, 2014) provides a useful lens through which to evaluate our application of game theory in Chapters 4 and 5. Firstly, it is important to note that the literature supporting the puzzle of small coalitions is firmly grounded in the assumption that decision makers act as *homines economici*. Relaxation of the assumption of *homo economicus* has been shown to give rise to greater levels of cooperation in other contexts such as climate agreements (see, for example, van der Pol et al., 2012). Chapters 4 and 5 retain the assumption of *homo economicus*. This allows us, particularly in Chapter 5, to contribute to the debate on the puzzle of small coalitions according to the terms in which it was set. The puzzle of small coalitions exists because theoretical evidence suggests that only relatively small Grand Coalitions can be stable, whereas larger coalitions have been observed in reality. The General Fisheries Commission for the Mediterranean (GFCM), for example, is a Regional Fisheries Management Organization with 23 member countries (www.fao.org/gfcm/en). Theoretically, in the simplest cases where players are symmetric, the Levhari and Mirman model (1980) shows that not even a Grand Coalition of two players is stable. For the Gordon-Schaefer model with symmetric players, Grand Coalitions are not stable for more than two players (Lindroos, 2008). While there are doubts about the extent to which the nations within agreements such as the GCFM are genuinely cooperating (Hannesson, 2014), there is clearly a discrepancy between theoretical predictions and empirical observations.

Chapter 5 shows how the puzzle of small coalitions, under the standard assumption that payoffs are evaluated only in steady states, can be solved by the inclusion of endogenous risk. However, after relaxing standard assumptions by including payoffs in the transition between steady states (as in Sakamoto, 2014), only small coalitions can be maintained and thus endogenous risk no longer solves the puzzle. While the term "puzzle of small coalitions" was coined by Breton and Keoula (2014), the term is justified because many previous studies have found similar results (e.g. Lindroos, 2008; Pintassilgo and Lindroos, 2008; Pintassilgo et al., 2010; Breton and Keoula, 2012). Our result regarding transition payoffs has implications for this body of work which justifies the puzzle of small coalitions. If previous work was repeated under the inclusion of transition payoffs, then the already limited number of players for which Grand Coalitions can be stable would be even less. This suggests that relaxing the assumption of *homo economicus* is a promising avenue for solving the puzzle of small coalitions. Further work could also be carried out to explore aspects of cooperation such as issue linkage (Folmer et al., 1993), the implications for cooperation of minimum participation requirements (Long, 2009) or the opportunity costs of fishing (Jensen et al., 2015).

The puzzle of small coalitions has not yet been resolved. Therefore, there is a shortage of predictive power in our game theoretic models. However, we can still be confident in our conclusions in a relative sense. For example, while we cannot be confident in predicting the size of stable coalitions under farsighted conjectures, we can confidently say that farsighted conjectures will increase the size of stable coalitions. Therefore, we can still provide some conclusions relevant for policy or decision making. Specific to Chapter 4, we find that if fish stock location shifts towards a given player, then that player can increase their harvest while decreasing their fishing effort. This means that if fish stock location shifts towards a player, then that player can increase their revenue while also decreasing their costs. This highlights the importance of changing fish stock location, with respect to, for example, changes in the cost of fishing. If unit costs of fishing effort go down then the model shows that profit will increase overall, but the total cost of fishing increases as well. Hence, the responses of fishing nations to change in fish stock location are fundamentally different to the responses to, for example, changes in fishing costs. Specific to Chapter 5, we find that the analysis excluding transition payoffs provides important insights for the management of slow-growing stocks. Deep-water fisheries, for example, are often slow-growing (Gordon, 2003) and more vulnerable to over-exploitation (Roberts, 2002; Neubauer et al., 2013). The results show that even if a Grand Coalition is stable for slow growing fish stocks such that the fish stock is managed in a socially optimal way, it may still be socially optimal to deplete the fish stock entirely. Cooperation is then, in this case, insufficient to ensure sustainability.

It is important to consider the institutional perspective in reflecting on the work in Part B. We do so by considering the relationship between theoretical Grand Coalitions and real-world IFAs. The institutional bodies used to managed high seas fisheries (which are the type of fisheries most relevant for the analyses in Part B) are Regional Fisheries Management Organizations (RFMOs). Nominally, an RFMO is a Grand Coalition if no unregulated fishing occurs (Pintassilgo et al., 2010), which means that only members of an RFMO agreement fish the stock of concern in the agreement. The question is, if a fish stock is managed by an RFMO which does not suffer from unregulated fishing, can that RFMO be considered to be a Grand Coalition? For this to be the case, the RFMO should be maximizing the joint benefit from the fish stock of all of its members. Hannesson (2014) provides evidence that the management fish stocks under RFMOs is better described as quasi-cooperation rather than the full cooperation required for an RFMO in the absence of unregulated fishing to be truly considered a Grand Coalition. Additionally, the management of high seas fisheries by RFMOs has been criticized on the grounds of poor discard governance and surveillance (Gilman et al., 2014), their untransparent use of scientific evidence (Polacheck, 2012), and indeed the ability of RFMOs to prevent overexploitation (Cullis-Suzuki and Pauly, 2010). It is therefore clear that RFMOs are not ideal institutions for managing high seas fisheries. Therefore, the existence of an RFMO in the absence of unregulated fishing does not guarantee the socially optimal fishing implied by the term “Grand Coalition”. This suggests that Grand Coalitions are a necessary but not sufficient condition for optimal management. Indeed, in the case of common-pool resource problems in general, the central coordination implied by a Grand Coalition is not even a necessary condition (Ostrom, 1990).

An RFMO is therefore clearly not the same as a Grand Coalition. There is therefore a need to consider, not only the conditions under which all fishing nations are nominally members of an agreement, but the conditions under which these agreements can meet their stated objectives. For example, Hoffmann and Quaas (2014) consider the incentives of policy makers to set particular Total Allowable Catch (TAC) levels under heterogeneous discount rates and demonstrate how this heterogeneity results in inefficiently high TACs. Hoffmann and Quaas take the involvement of all relevant fishing nations in the TAC setting process as given. In general then, a two pronged approach is necessary. The first prong, as in this thesis, should consider the conditions for stability of Grand Coalitions, and the second prong should consider the conditions for these Grand Coalitions to achieve socially optimal outcomes, as in Hoffmann and Quaas (2014). By doing so, deeper insights can be gained into the research questions dealt with in Chapters 4 and 5.

6.4 Concluding remarks

As mentioned at the beginning of the introduction, the research in this thesis was conducted as part of the EU Seventh Framework Program project named VECTORS. The aim of VECTORS was to examine the changes taking place in European seas, their causes, and their impacts on society in order to determine the economic implications of change and to formulate management strategies. The researchers in VECTORS were predominantly biologists and ecologists, with a good handful of economists to boot. Economic research within VECTORS consisted of environmental valuation, applied models of fisheries management, general equilibrium modelling, and the work of this thesis. For further information on the economics of marine ecological change, one can consult www.marine-vectors.eu, which contains all research within the VECTORS project. Clearly then, VECTORS has involved many disciplines.

Let us define the "disciplinarity" of the VECTORS project using the definition of Choi and Pak (2006: 351), who state that "multidisciplinarity draws on knowledge from different disciplines but stays within their boundaries. Interdisciplinarity analyses, synthesizes and harmonizes links between disciplines into a coordinated and coherent whole. Transdisciplinarity integrates [different] sciences in a humanities context, and transcends their traditional boundaries". According to this categorization, the research carried out in this thesis is multidisciplinary. We have drawn on insights from biology and ecology but have stayed within the boundaries of economics. This is not to say that this is a shortcoming of the research, but rather to say that there is much potential for further integration between disciplines and that such integration would facilitate different kinds of results, and potentially, ones with more direct policy implications. The multidisciplinary research in this thesis can serve as a base from which further multidisciplinary research can be conducted, as well as serving as a basis for progressing to more interdisciplinary, and potentially transdisciplinary approaches.

Finally, it should be noted that this thesis has covered a limited range of changes in marine ecosystems and its focus, partly due to the requirements of funding sources, has been dictated by issues of particular concern in Europe. Climate change is having much more varied effects on marine ecosystems than changing stock location alone (Hoegh-Gulberg and Bruno, 2010). These include changes in productivity, changes in food-web dynamics and a greater incidence of disease. These changes are driven not only by increases in ocean temperatures, but also ocean acidification due to increased atmospheric carbon dioxide concentrations (Gattuso et al., 2015). The impacts of these changes for ecosystem services provision may be severe (Gattuso et al., 2015), and this holds particularly true for less developed countries (Allison et al., 2009). Two points should therefore be made. Firstly, changes in marine ecosystems are not confined to only changes in fish stock location, risk of stock collapse and the arrival of invasive species. Secondly, the effects of marine ecological change depend on local economic and ecological conditions. There is therefore significant scope to consider other aspects of marine ecological change, and also to consider their effects and management in other contexts, such as the context of non-European countries.

References

Achterberg, T., 2009. SCIP: Solving constraint integer programs. Mathematical Programming Computation 1, 1-14.

Adams, D., Lee, D., 2012. Technology adoption and mitigation of invasive species damage and risk: application to zebra mussels. Journal of Bioeconomics 14, 21-40.

Agbo, M., 2014. Strategic exploitation with learning and heterogeneous beliefs. Journal of Environmental Economics and Management 76, 126-140.

Albers, H.J., Fischer, C., Sanchirico, J.N., 2010. Invasive species management in a spatially heterogenous world: Effects of uniform policies. Resource and Energy Economics 32, 482-499.

Albert, R.J., Lishman, J.M., Saxena, J.R., 2013. Ballast water regulations and the move towards concentration based numeric discharge limits. Ecological Applications 23, 289-300.

Allison, E.H., Perry, A.L., Badjeck, M-C., Adger, W.N., Brown, K., Conway, D., Halls, A.S., Pilling, G.M., Andrew, N.L., Dulvy, N.K., 2009. Vulnerability of national economies to the impact of climate change on fisheries. Fish and Fisheries 10, 173-196.

Alpízar, G., Gordillo, L.F., 2013. Disease spread in coupled populations: minimising response strategies costs in discrete time models. Discrete dynamics in nature and society. Article ID 681689.

Antoniadou, E., Koulovatianos, C., Mirman, L.J., 2013. Strategic exploitation of a common-property resource under uncertainty. Journal of Environmental Economics and Management 65, 28-39.

Arnason, R., 2012. Global warming: New challenges for the common fisheries policy? Ocean and Coastal Management 70, 4-9.

d'Aspremont, C., Jacquemin, A., Gabszewicz, J.J., Weymark, J.A., 1983. On the stability of collusive price leadership. Canadian Journal of Economics 16,17–25.

Bailey, M., Sumaila U.R., Lindroos, M., 2010. Application of game theory to fisheries over three decades. Fisheries Research 102, 1-8.

Barrett, S., 2013. Climate treaties and approaching catastrophes. Journal of Environmental Economics and Management 66, 235-250.

Bellman, R., 1957. Dynamic Programming. Princeton University Press, Princeton.

Biancardi, M., Di Liddo, A., 2010. The size of farsighted stable coalitions in a game of pollution abatement. Natural Resource Modeling 23, 467-493.

Bjørndal, T., Kaitala, V., Lindroos, M., Munro, G.M., 2000. The management of high seas fisheries. Annals of Operations Research 94, 183-196.

Bjørndal. T., Lindroos, M., 2012. Cooperative and non-cooperative management of the Northeast Atlantic cod fishery. Journal of Bioeconomics 14, 41-60.

Blackburn, T.M., Lockwood, J.L., Cassey, P., 2015. The influence of numbers on invasion success. Molecular Ecology 24, 1942-1953.

Blackwood, J., Hastings, A., Costello, C., 2010. Cost-effective management of invasive species using linear-quadratic control. Ecological Economics 69, 519-527.

Brandt, U.S., Kronbak, L.G., 2010. On the stability of fishery agreements under exogenous change: An example of agreements under climate change. Fisheries Research 101, 11-19.

Breton, M., Keoula, M.Y., 2012. Farsightedness in a coalitional great fish war. Environmental and Resource Economics 51, 297-315.

Breton, M., Keoula, M.Y., 2014. A great fish war with asymmetric players. Ecological Economics 97, 209-223.

Buhle, E.R., Margolis, M., Ruesink, J.L., 2005. Bang for buck: cost-effective control of invasive species with different life histories. Ecological Economics 52, 355-366.

Burnett, K., Kaiser, B., Pitafic, B.A., Roumasset, J., 2006. Prevention, eradication and containment of invasive species: illustrations from Hawaii. Agricultural and Resource Economics Review 25, 63-67.

Burnett, K., Kaiser, B.A., Roumasset, J.A., 2007. Invasive species control over space and time: miconia calvescens on Oahu, Hawaii. Journal of Agricultural and Applied Economics 39, 125-132

Burnett, K., Pongkijvorasin, S., Roumasset, J., 2012. Species Invasion as Catastrophe: The Case of the Brown Tree Snake. Environmental and Resource Economics 51, 241-254.

Cacho, O.J., Wise, R.M., Hester, S.M., Sinden, J.A., 2008. Bioeconomic modelling for control of weeds in natural environments. Ecological Economics 65, 559-568.

Carraro, C., Siniscalco, D., 1998. International environmental agreements: incentives and political economy. European Economic Review 42, 561-572.

Carrasco, L.R., Baker, R., MacLeod, A., Knight, J.D., Mumford, J.D., 2010a. Optimal and robust control of invasive alien species spreading in homogeneous landscapes. Journal of The Royal Society Interface 7, 529-540.

Carrasco, L.R., Mumford, J.D., MacLeod, A., Knight, J.D., Baker, R.H.A., 2010b. Comprehensive bioeconomic modelling of multiple harmful non-indigenous species. Ecological Economics 69, 1303-1312.

Carrasco, L. R., Cook, D., Baker, R., MacLeod, A., Knight, J.D., Mumford, J. D., 2012. Towards the integration of spread and economic impacts of biological invasions in a landscape of learning and imitating agents. Ecological Economics 76, 95-103.

Caparrós, A., Giraud-Heraud, E., 2011. Coalition stability with heterogenous agents. Economics Bulletin 31, 286-296.

Chadès , I., Martin, T.G., Nicol, S., Burgman, M.A., Possingham, H.P., Buckley, Y.M., 2011. General rules for managing and surveying networks of pests, diseases, and endangered species. Proceedings of the National Academy of Sciences 108, 8323-8328.

Chalak-Haghighi, M., van Ierland, E.C., Bourdôt, G.W., Leathwick, D., 2008. Management strategies for an invasive weed: a dynamic programming approach for Californian thistle in New Zealand. New Zealand Journal of Agricultural Research 51, 409-424.

Cheung, W.W.L., Lam, V.W.Y., Sarmiento, J.L., Kearney, K., Watson, R., Pauly, D., 2009. Projecting global marine biodiversity impacts under climate change scenarios. Fish and Fisheries 10, 235-251.

Choi, B.C., Pak, A.W., 2006. Multidisciplinarity, interdisciplinarity and transdisciplinarity in health research, services, education and policy: 1. Definitions, objectives, and evidence of effectiveness. Clinical and Investigative Medicine 29, 351-64.

Chwe, M.S-Y., 1994. Farsighted Coalitional Stability. Journal of Economic Theory 63, 299-325.

Clark, C.W., 2010. Mathematical bioeconomics. The mathematics of conservation. Third Edition. Wiley, Hoboken, New Jersey.

Clarke, H.R., Reed, W.J.,1994. Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse. Journal of Economic Dynamics and Control 18, 991-1010.

Cooke, J.G., 1984. Glossary of technical terms. In: Exploitation of Marine Communities: report of the Dahlem Workshop on Exploitation of Marine Communities, Berlin, R.M. May, Eds.

Costello, C., McAusland, C., 2003. Protectionism, trade, and measures of damage from exotic species introductions. American Journal of Agricultural Economics 85, 964-975.

Crall, A.W., Meyerson, L.A., Stohlgren, T.J., Jarnevich, C.S., Newman, G.J., Graham, J., 2006. Show me the numbers: what data currently exist for non-native species in the USA? Frontiers in Ecology and the Environment 4, 414-418.

Crutchfield, J., 1964. The marine fisheries: a problem in international cooperation. The American Economic Review 54, 207-218.

Cullis-Suzuki, S., Pauly, D., 2010. Failing the high seas: a global evaluation of regional fisheris management organisations. Marine Policy 34, 1036-1042.

de Zeeuw, A., 2008. Dynamic effects on the stability of international environmental agreements. Journal of Environmental Economics and Management 55, 163-174.

D'Evelyn, S.T., Tarui, N., Burnett, K., Roumasset, J.A., 2008. Learning-by-catching: uncertain invasive-species populations and the value of information. Journal of Environmental Management 89, 284-292.

Dellink, R., Finus, M., 2012. Uncertainty and climate treaties: Does ignorance pay? Resource and Energy Economics 34, 565-584.

Dobbs, F.C., Rogerson, A., 2005. Ridding ships' ballast water of microorganisms. Environmental Science and Technology 39, 259-264.

Dryzek, J.S., 1990. Discursive democracy: politics, policy and political science. Cambridge University Press, Cambridge.

Duncan, R.P., Blackburn, T.M., Rossinelli, S., Bacher, S., 2014. Quantifying invasion risk: the relationship between establishment probability and founding population size. Methods in Ecology and Evolution 5, 1255-1263.

Eiswerth, M.E., Johnson, W.S., 2002. Managing Nonindigenous Invasive Species: insights from dynamic analysis. Environmental and Resource Economics 23, 319-342.

Ekerhovd, N-A., 2010. The stability and resilience of management agreements on climate-sensitive straddling fishery resources: the blue whiting (Micromesistius poutassou) coastal state agreement. Canadian Journal of Fisheries and Aquatic Sciences 67, 534-552.

Ellefsen, H., 2012. The stability of fishing agreements with entry: the North East Atlantic mackerel. Strategic Behaviour and the Environment 3, 1-24.

Epanchin-Niell, R.S., Haight, R.G., Berec, L., Kean, J.M., Liebhold, A.M., 2012. Optimal surveillance and eradication of invasive species in heterogeneous landscapes. Ecology Letters 15, 803-812.

Epanchin-Niell, R.S., Hastings, A., 2010. Controlling established invaders: integrating economics and spread dynamics to determine optimal management. Ecology Letters 13, 528-541.

Epanchin-Niell, R.S., Wilen, J.E., 2012. Optimal spatial control of biological invasions. Journal of Environmental Economics and Management 63, 260-270.

Eyckmans, J., Finus, M., 2004. An almost ideal sharing scheme for coalition games with externalities. Working Paper. Katholieke Universiteit Lueven, Center for Economic Studies.

Fenichel, E.P., Richards, T.J., Shanafelt, D.W., 2014. The control of invasive species on private property with neighbor-to-neighbor spillovers. Environmental and Resource Economics 59, 231-251.

Fernandez, L., 2006. Marine Shipping Trade and Invasive Species Management Stategies. International Game Theory Review 8, 153-168.

Fesselmeyer, E., Santugini, M., 2013. Strategic exploitation of a common resource under environmental risk. Journal of Economics Dynamics and Control 37, 125-136.

Field, I.C., Meekan, M.G., Buckworth, R.C., Bradshaw, C.J.A., 2009. Susceptibility of sharks, rays and chimaeras to global extinction. Advances in marine biology 56, 275-363.

Finnoff, D., Shogren, J.F., Leung, B., Lodge, D., 2007. Take a risk: Preferring prevention over control of biological invaders. Ecological Economics 62, 216-222.

Finnoff, D., Potapov, A., Lewis, M.A., 2010a. Control and management of a spreading invader. Resource and Energy Economics 32, 534-550.

Finnoff, D., McIntosh, C., Shogren, J.F., Sims, C., Warziniack, T., 2010b. Invasive species and endogenous risk. Annual Review of Environmental and Resource Economics 2, 77-100.

Fischer, F., 1990. Technocracy and the politics of expertise. Sage Publications, Newbury Park, California.

Folmer, H., Mouche, P.v., Ragland, S., 1993. Interconnected games and international environmental problems. Environmental and Resource Economics 3, 313-335.

Frésard, M., Boncoeur, J., 2006. Controlling the biological invasion of a commerical fishery by a space competitor: a bioeconomic model with reference to the bay of St-Brieuc scallop fishery. Agricultural and Resource Economics Review 35, 78-97.

Galil, B., Boero, F., Campbell, M., Carlton, J., Cook, E., Fraschetti, S., Gollasch, S., Hewitt, C., Jelmert, A., Macpherson, E., Marchini, A., McKenzie, C., Minchin, D., Occhipinti-Ambrogi, A., Ojaveer, H., Olenin, S., Piraino, S., Ruiz, G., 2014. 'Double trouble': the expansion of the Suez Canal and marine bioinvasions in the Mediterranean Sea. Biological Invasions 17, 973-976.

Gattuso, J-P., Magnan, A., Billé, R., Cheung, W.W.L., Howes, E.L., Joos, F., Allemand, D., Bopp, L. Cooley, S.R., Eakin, C.M., Hoegh-Guldberg, O., Kelly, R.P., Pörtner, H-O., Rogers, A.D., Baxter, J.M., Laffoley, D., Osborn, D., Rankovic, A., Rochetter, J., Sumaila, U.R., Treyer, S., Turley, C., 2015. Contrasting futures for ocean and society from different anthropogenic CO2 emissions scenarios. Science 349, 45.

Genovesi, P., 2005. Eradications of invasive alien species in Europe: a review. Issues in Bioinvasion Science 7, 127-133.

Genovesi, P., Carboneras, C., Vilà, M., Walton, P., 2015. EU adopts innovative legislation on invasive species: a step towards a global response to biological invasions? Biological Invasions 17, 1307-1311.

Gibbons, R., 1992. A primer in game theory. Harvester Wheatsheaf, Hemel Hempstead.

Gjøsæter, H., Bogstad, B., Tjelmeland, S., 2009. Ecosystem effects of the three capelin stock collapses in the Barents Sea. Marine Biology Research 5, 40-53.

Gollasch, S., David, M., Voigt, M., Dragsund, E., Hewitt, C., Fukuyo, Y., 2007. Critical review of the IMO international convention on the management of ships' ballast water and sediments. Harmful Algae 6, 585-600.

Gordon, J.D.M., 2003. Environmental and biological aspects of deepwater demersal fisheries. In: Deep Sea 2003: conference on the Governance and Management of Deepwater Fisheries. Ed. Shotton, R. Queenstown, New Zealand: FAO.

Gurevitch, J, Padilla, D.K., 2004. Are invasive species a major cause of extinctions? Trends in Ecology and Evolution 19, 470-474.

Haight, R.G., Polasky, S., 2010. Optimal control of an invasive species with imperfect information about the level of infestation. Resource and Energy Economics 32, 519-533.

Hannesson, R., 2011. Game theory and fisheries. Annual Review of Resource Economics 3, 181-202.

Hannesson, R., 2014. Does threat of mutually assured destruction produce quasi-cooperation in the mackerel fishery? Marine Policy 44, 342-350.

Haraldsson, G., Carey, D., 2011. Ensuring a Sustainable and Efficient Fishery in Iceland, OECD Economics Department Working Papers. http://www.oecd-ilibrary.org/economics/ensuring-a-sustainable-and-efficient-fishery-in-iceland_5kg566jfrpzr-en. Cited 28 Sept. 2012.

Harsanyi, J.C., 1974. An equilibrium-point interpretation of stable sets and a proposed alternative definition. Management Science 20, 1472-1495.

Haydar, D., Wolff, W.J., 2011. Predicting invasion patterns in coastal ecosystems: relationship between vector strength and vector tempo. Marine Ecological Progress Series 431, 1-10.

Herborg, L.-M., Rushton, S.P., Clare, A.S., Bentley, M.G., 2003. Spread of the Chinese mitten crab (Eriocheir sinensis H. Milne Edwards) in Continental Europe: analysis of a historical data set. Hydrobiologica 503, 21-28.

Herborg, L.-M., Weetman, D., Oosterhout , C.v., Hänfling, B., 2007. Genetic population structure and contemporary dispersal patterns of a recent European invader, the Chinese mitten crab, Eriocheir sinensis. Molecular Ecology 16, 231-242.

Hilborn, R., Orensanz, J.M., Parma, A.M., 2005. Institutions, incentives and the future of fisheries. Philosophical Transactions of the Royal Society B: Biological Sciences 360, 47-57.

Hilborn, R., 2007. Moving to sustainability by learning from successful fisheries. AMBIO: a journal of the human environment 36, 296-303.

Hoeght-Guldberg, O., Bruno, J.F., 2010. The impact of climate change on the world's marine ecosystems. Science 328, 1523-1528.

Hoffmann, J., Quaas, M.F., 2014. Common pool politics and inefficient fishery management. Environmental and Resource Economics (in press). DOI 10.1007/s10640-014-9842-4.

Horan, R.D., Perrings, C., Lupi, F., Bulte, E.H., 2002. Biological Pollution Prevention Strategies under Ignorance: The Case of Invasive Species. American Journal of Agricultural Economics 84, 1303-1310.

Huffaker, R.G., Bhat, M.G., Lenhart, S.M., 1992. Optimal trapping strategies for diffusing nuisance-beaver populations. Natural Resource Modelling 6, 71-97.

Hulme, P.E., 2006. Beyond control: wider implications for the management of biological invasions. Journal of Applied Ecology 43, 845-847.

Hulme, P.E., 2009. Trade, transport and trouble: managing invasive species pathways in an era of globalization. Journal of Applied Ecology 46, 10-18.

Hyytiäinen, K., Lehtiniemi, M., Niemi, J.K., Tikka, K., 2013. An optimization framework for addressing aquatic invasive species. Ecological Economics 91, 69-79.

IMO, 2004. International Convention for the Control and Management of Ships' Ballast Water and Sediments. International Maritine Organization. http://www.imo.org.

IUCN, 2000. Guidelines for the Prevention of Biodiversity Loss Caused by Alien Invasive Species. International Union for the Conservation of Nature. Gland, Switzerland.

Impson, F.A.C., Moran, V., Hoffmann, J., 2004. Biological control of an alien tree, Acacia cyclops, in South Africa: impact and dispersal of a seed-feeding weevil, Melanterius servulus. Biological Control 29, 375–381.

Ishimura, G., Herrick, S., Sumaila, U.R., 2012. Fishing games under climate variability: transboundary management of Pacific sardine in the California current system. Environmental Economics and Policy Studies 15, 189-209.

Jansen, T., Gislason, H., 2011. Temperature affects the timing of spawning and migration of North Sea mackerel. Continental Shelf Research 31, 64-72.

Jensen, F., Frost, H., Thøgersen, T., Anderson, P., Anderson, J.L., 2015. Game theory and fish wars: The case of the Northeast Atlantic mackerel fishery. Fish and Fisheries 172, 7-16.

Judd, K.L., 1998. Numerical methods in economics. The MIT Press, Cambridge, Massachusetts.

Kaitala, V.T., Lindroos, M., 1998. Sharing the benefits of cooperation in high seas fisheries: a characteristic function game approach. Natural Resource Modeling 11, 275-299.

Keller, R.P., Drake, J.M., Drew, M.B., Lodge, D.M., 2011. Linking environmental conditions and ship movements to estimate invasive species transport across the global shipping network. Diversity and Distributions 17, 93-102.

Kennedy, J., 2003. Scope for efficient multinational exploitation of North-East Atlantic mackerel. Marine Resource Economics 18, 55-80.

Kim, C.S., Lubowski, R.N., Lewandrowski, J., Eisweth, M.E., 2006. Prevention or control: optimal government policies for invasive species management. Agricultural and Resource Economics Review 35, 29-40.

King, D.M., Hagan, P.T., Riggio, M., Wright, D.A., 2012. Preview of global ballast water treatment markets. Journal of Marine Engineering and Technology 11, 3-15.

Knowler, D., Barbier, E., 2000. The economics of an invading species: a theoretical model and case study application. In Perrings, C., Williamson, M., Dalmazzone, S. (Eds) The Economics of Biological Invasions. Edward Elgar, Cheltnam, UK.

Knowler, D., Barbier, E., 2005. Importing exotic plants and the risk of invasion: are market-based instruments adequate? Ecological Economics 52, 341-354.

Kolar, C.S. Lodge, D.M., 2001. Progress in invasion biology: predicting invaders. Trends in Ecology and Evolution 16, 199-204.

Kolstad, C.D., 2007. Systematic Uncertainty in self-enforcing international agreements. Journal of Environmental Economics and Management 53, 68-79.

Kovacs, K.F., Haight, R.G., Mercader, R.J., McCullough, D.G., 2014. A bioeconomic analysis of an emerald ash borer invasion of an urban forest with multiple jurisdictions. Resource and Energy Economics 36, 270-289.

Kraft, D., Folmer, E.O., Meyedirks, J., Stiehl, T., 2010. Data inventory of the tidal basins in the trilateral Wadden Zee. Programma naar een rijke Waddenzee.

Kramer, A.M., Dennis, B., Liebhold, A.M., Drake, J.M., 2009. The evidence for Allee effects. Population Ecology 51, 341-354.

Kwon, O.S., 2006. Partial international coordination in the Great Fish War. Environmental and Resource Economics 33, 463-483.

Lee, H., Reusser, D.A., Frazier, M., 2013. Approaches to setting organism-based ballast water discharge standards. Ecological Applications 23, 301-310.

Leppäkoski, E., Gollasch, S., Olenin, S., 2002. Aquatic Invasive Species of Europe - Distribution, Impacts and Management. Kluwer Academic Publishers, Dordrecht.

Leung, B., Lodge, D.M., Finnoff, D., Shogren, J.F., Lewis, M.A., Lamberti, G., 2002. An ounce of prevention or a pound of cure: bioeconomic risk analysis of invasive species. Proceedings of the Royal Society of London. Series B: Biological Sciences 269, 2407-2413.

Leung, B., Drake, J.M., Lodge, D.M., 2004. Predicting invasions: propagule pressure and the gravity of allee effects. Ecology 85, 1651-1660.

Levhari, D., Mirman, L.J., 1980. The Great Fish War: An example using a dynamic Cournot-Nash solution. The Bell Journal of Economics 11, 322-334.

Lindroos, M., 2008. Coalitions in international fisheries management. Natural Resource Modeling 21, 366-384.

Liu, T.-K., Tsai, T.-K., 2011. Vessel traffic patterns in the Port of Kaohsiung and the management implications for preventing the introduction of non-indigenous aquatic species. Marine Pollution Bulletin 62, 602-608.

Lockwood, J.L., Cassey, P., Blackburn, T., 2005. The role of propagule pressure in explaining species invasions. Trends in Ecology & Evolution 20, 223-228.

Long, L.K., Flaaten, O., 2011. A Stackelberg analysis of the potential for cooperation in straddling stock fisheries. Marine Resource Economics 26, 119-139.

Mack, R.N., Simberloff, D., Lonsdale, W.M., Evans, H., Clout, M., Bazzaz, F.A., 2000. Biotic invasions: causes, epidemiology, global consequences and control. Ecological Applications 10, 689-710.

Majone, G., 1989. Evidence, argument and persausion in the policy process. Yale University Press, New Haven, Connecticut.

Mbah, M.L.N., Gilligan, C.A., 2011. Resource Allocation for Epidemic Control in Metapopulations. PLoS ONE 6, e24577.

McDermott, S.M., Irwin, R.E.,Taylor, B.W., 2013. Using economic instruments to develop effective management of invasive species: insights from a bioeconomic model. Ecological Applications 23, 1086-1100.

McGinty, M., 2007. International environmental agreements among asymmetric nations. Oxford Economic Papers 59, 45-62.

McGinty, M., Milam, G., Gelves, A., 2012. Coalition stability in public goods provision: testing an optimal allocation rule. Environmental and Resource Economics 52, 327–345.

Mehta, S.V., Haight, R.G., Homans, F.R., Polasky, S., Venette, R.C., 2007. Optimal detection and control strategies for invasive species management. Ecological Economics 61, 237-245

Mesterton-Gibbons, M., 1992. An introduction to game-theoretic modelling, Addison-Wesley, California.

Miller, K.A., Munro, G.R., 2004. Climate and cooperation: a new perspective on the management of shared fish stocks. Marine Resource Economics 19, 367-393.

Mirman, L.J., 1979. Dynamic models of fishing: a heuristic approach. In: "Control theory in mathematical economics", Liu and Sutinen, Eds. Decker, New York.

Molnar, J.L., Gamboa, R.L., Revenga, C., Spalding, M.D., 2008. Assessing the global threat of invasive species to marine biodiversity. Frontiers in Ecology and the Environment 6, 485-492.

Mullon, C., Fréon, P., Cury, P., 2005. The dynamics of collapse in world fisheries. Fish and Fisheries 6, 111-120.

Munro, G.R., 1979. The optimal management of transboundary renewable resources. The Canadian Journal of Economics 12, 355-376.

Munro, G.R., 2008. Game theory and the development of resource management policy: the case of international fisheries. Environment and Development Economics 14, 7-27.

Neubauer, P., Jensen, O.P., Hutching, J.A., Baum, J.K., 2013. Resilience and recovery of overexploited marine populations. Science 340, 347-349.

Nkuiya, B., Marrouch, W., Bahel, E., 2014. International environmental agreements under endogenous uncertainty. Journal of Public Economic Theory, (in press). DOI 10.1111/jpet.12108.

Nunes, P.A.L.D., Markandya, A., 2008. Economic value of damage caused by marine bio-invasions: lessons from two European case studies. ICES Journal of Marine Science 65, 775-780.

Olson, L.J., Roy, S., 2002. The economics of controlling a stochastic biological invasion. American Journal of Agricultural Economics 5, 1311-1316.

Olson, L.J., Roy, S., 2005. On prevention and control of uncertain biological invasions. Review of Agricultural Economics 27, 491-497.

Osmani, D., Tol, R., 2009. Toward farsightedly stable international environmental agreement. Journal of Public Economic Theory 11, 455-492.

Ostrom, E., 1990. Governing the commons: the evolution of institutions for collective action. Cambridge University Press, Cambridge.

Parker, I.M., Simberloff, D., Lonsdale, W.M., Goodell, K., Wonham, M., Kareiva, P.M., Williamson, M.H., Von Holle, B., Moyle, P.B., Byers, J.E., Goldwasser, L., 1999. Impact: Toward a Framework for Understanding the Ecological Effects of Invaders. Biological Invasions 1, 3-19.

Pimentel, D., Zuniga, R., Morrison, D., 2005. Update on the environmental and economic costs associated with alien-invasive species in the United States. Ecological Economics 52, 273– 288.

Pintassilgo, P., Finus, M., Lindroos, M., Munro, G.R., 2010. Stability and success of regional fisheries management organizations. Environmental and Resource Economics 46, 377-402.

Pintassilgo, P., Lindroos, M., 2008. Coalition formation in straddling stock fisheries: a partition function approach. International Game Theory Review 10, 303-317.

Polacheck, T., 2012. Politics and independent scientific advice in RFMO processes: a case study of crossing boundaries. Marine Policy 32, 132-141.

Punt, M.J., Weikard, H-P., van Ierland, E.C., 2012. Marine protected areas in the high seas and the impact on international fishing agreements. Natural Resource Modeling 26, 164-193.

Rawls, J., 1972. A theory of justice. Oxford University Press, Oxford.

Ren, B., Polasky, S., 2014. The optimal management of renewable resources under the risk of potential regime shift. Journal of Economic Dynamics and Control 40, 195-212.

Rettieva, A.N., 2012. Stable coalition structure in bioresource management problem. Ecological Modelling 235–236, 102-118.

Ricciardi, A., Steiner, W.W.M., Mack, R.N., Simberloff, D., 2000. Toward a Global Information System for Invasive Species. BioScience 50, 239-244.

Roberts, C.M., 2002. Deep impact: the rising toll of fishing in the deep sea. Trends in Ecology and Evolution 17, 242-245.

Rowthorn, R.E., Laxminarayan, R., Gilligan, C.A., 2009. Optimal control of epidemics in metapopulations. Journal of the Royal Society Interface 6, 1135-1144.

Ruiz, G.M., Carlton, J.T., Grosholz, E.D., Hines, A.H., 1997. Global invasions of marine and estuarine habitats by non-indigenous species: mechanisms, extent and consequences. American Zoologist 37, 621-632.

Sadeghi, R., Zarkami, R., Van Damme, P., 2014. Modelling habitat preference of an alien aquatic fern, Azolla filiculoides (Lam.), in Anzali wetland (Iran) using data-driven methods. Ecological Modelling 284, 1-9.

Sakai, A.K., Allendorf, F.W., Holt, J.S., Lodge, D.M., Molofsky, J., With, K.A., Baughman, S., Cabin, R.J., Cohen, J.E., Ellstrand, N.C., McCauley, D.E., O'Niel, P., Parker, I.M., Thompson, J.N., Weller, S.G., 2001. The population biology of invasive species. Annual Review of Ecology and Systematics 32, 305-332.

Sakamoto, H., 2014. Dynamic resource management under the risk of regime shifts. Journal of Environmental Economics and Management 68, 1-19.

Salinas, R.A, Lenhart, S., Gross, L.J., 2005. Control of a metapopulation harvesting model for black bears. Natural Resource Modeling 18, 307-321.

Sanchirico, J.N., Albers, H.J., Fischer, C., Coleman, C., 2010. Spatial management of invasive species: pathways and policy options. Environmental and Resource Economics 45, 517-535.

Sanderson, I., 2006. Complexity, 'practical rationality' and evidence-based policy making. Policy and Politics 34, 115-132.

Savage, D., Renton, M., 2014. Requirements, design and implementation of a general model of biological invasion. Ecological Modelling 272, 394– 409.

Scalera, R., 2010. How much is Europe Spending on invasive alien species? Biological Invasions 12, 172-177.

Schwandt, T.A., 1997. Evaluation as practical hermeneutics. Evaluation 3, 69-83.

Schwandt, T.A., 2000. Further diagnostic thoughts on what ails evaluation practise. American Journal of Evaluation 21, 213-219.

Scriven, R.D., DiBacco, C., Locke, A., Therriault, T.W., 2015. Ballast water management in Canada: a historical perspective and implications for the future. Marine Policy 59, 121-133.

Sharov, A.A., 2004. Bioeconomics of Managing the Spread of Exotic Pest Species with Barrier Zones. Risk Analysis 24, 879-892.

Simberloff, D., 2009. The Role of Propagule Pressure in Biological Invasions. Annual Review of Ecology, Evolution, and Systematics 40, 81-102.

Sims, C., Finnoff, D., 2013. When is a "wait and see" approach to invasive species justified? Resource and Energy Economics 35, 235-255.

Sorger, G., 2005. A dynamic common property resource problem with amenity value and extraction costs. International Journal of Economic Theory 1, 3-19.

Springborn, M.R., 2014. Risk aversion and adaptive management: Insights from a multi-armed bandit model of invasive species risk. Journal of Environmental Economics and Management 68, 226-242.

Switzer, S.E., Therriault, T.W., Dunham, A., Pearce, C.M., 2011. Assessing potential control options for the invasive tunicate Didemnum vexillum in shellfish aquaculture. Aquaculture 318, 145-153.

Troost, K., 2010. Causes and effects of a highly successful marine invasion: Case-study of the introduced Pacific oyster Crassostrea gigas in continental NW European estuaries. Journal of Sea Research 64, 145-165.

Tsolokai, E., Diamadopoulos, E., 2010. Technologies for ballast water treatment: a review. Journal of chemical technology and biotechnology 85, 19-32.

Tsur, Y., Zemel, A., 1998. Pollution control in an uncertain environment. Journal of Economic Dynamics and Control 22, 967–975.

Tsur, Y., Zemel, A., 2004. Endangered aquifers: groundwater management under threats of catastrophic events. Water Resources Research 40, 1–10.

Vilà, M., Basnou, C., Pyšek, P., Josefsson, M., Genovesi, P., Gollasch, S., Nentwig, W., Olenin, S., Roques, A., Roy, D., Hulme, P.E., 2009. How well do we understand the impacts of alien species on ecosystem services? A pan-European, cross-taxa assessment. Frontiers in Ecology and the Environment 8, 135-144.

Vitousek, P.M., Mooney, H.A., Lubchenco, J., Melillo, J.M., 1997. Human domination of earth's ecosystems. Science 277, 494-499.

Weikard, H-P., 2009. Cartel stability under an optimal sharing rule. The Manchester School 77, 575-593.

Williamson, M., 1989. Mathematical models of invasions, in: Drake, J., Mooney, H. (Eds.), Biological Invasions - a global perspective. John Wiley and Sons, New York.

Williamson, M., 1996. Biological invasions. Chapman & Hall, London.

Zhang, W., Werf, W.V.D., Swinton, S.M., 2010. Spatially optimal habitat management for enhancing natural control of an invasive agricultural pest: Soybean aphid. Resource and Energy Economics 32, 551-565.

Summary

This thesis is divided into two parts, as explained in Chapter 1, which focus on different aspects of marine ecological change. Part A considers marine Invasive Alien Species (IAS), which are taxa introduced outside of their native range. The detrimental consequences of invasions for human welfare necessitate management of IAS. There are two types of IAS management. These are (i) management of the risks that an invasion will become established, termed “prevention”, and (ii) management of already established invasions, termed “control”. Chapter 2 considers prevention of invasive species with Ballast Water Management (BWM). Vessels transport invasive species in their ballast water. BWM involves treating ballast water to reduce the risk of successful invasion establishment. Chapter 2 studies the determinants of optimal ballast water treatment standards from a theoretical perspective. Chapter 3 considers control of already established invasions from a spatial and dynamic perspective. We model a non-native habitat divided into patches, where each patch may contain a population of the invasive species, and where spread of the invasion between patches is a stochastic process. In this context, we derive optimal management policies.

The second part of this thesis: Part B, considers International Fisheries Agreements (IFAs). IFAs facilitate cooperation in the management of fish stocks. Cooperation is necessary to ensure sustainable management. Part B focuses on two issues which may affect the stability of cooperation within IFAs. These are; in Chapter 4, changes in stock location, which may occur due to climate change, and in Chapter 5, the risk of stock collapse, which may exist due to overfishing. Part B uses game theory to analyse the effects of these two issues on the stability of the Grand Coalition, which is the state of affairs where all parties cooperate to maximize their joint benefit from the fish stock.

The methods and findings of the thesis are summarized as follows: in Chapter 2 (Part A), we construct a model to study optimal BWM standards. The model is built around the assumption that invasions arriving via ballast water are irreversible, i.e. once an invasion has arrived, it is not possible to reduce the size of the invasive population to zero. The hazard rate of invasion establishment can be reduced by

setting a BWM standard. The hazard rate is also affected by the Minimum Viable Population (MVPs) of the species and the possibility of an Allee effect. An MVP exists if there is some population size below which there is an insufficient number of invasive individuals to sustain a population. An Allee effect exists if the probability that a population survives increases at an increasing rate in the size of the population. Our analysis focuses on the conditions under which a BWM standard which aims to reduce invasive populations in ballast water to below their MVPs (as is aimed for by the BWM convention) can be optimal. We find that the current aim of the BWM convention can only be optimal in the case that the hazard function (which determines the hazard rate) is not continuously differentiable around the MVP. We find that Allee effects are a requirement for a continuously differentiable hazard function. Therefore, we find that whether or not an Allee effect exists fundamentally affects whether it is optimal to aim to reduce an invasive population in ballast water to marginally below its MVP.

In Chapter 3 (Part A), we combine aspects of previous modelling approaches to provide new generalized management insights for controlling established invasions. We employ a metapopulation network consisting of patches which are arranged one-dimensionally (i.e. in a line), which is relevant, among other cases, for invasive species spreading along coastlines. We allow for the population size of the invasion within patches to be reduced, which we term "removal", and we allow for the probability of spread between patches to be reduced without affecting the population sizes directly, which we term "containment". We employ numerical stochastic dynamic programming to explore how these two interventions (removal and containment) can be optimally applied to minimize the sum of damages from the invasion and the costs of removing and containing the invasion. We find that allowing for varying stock sizes within patches facilitates optimal timing of the application of containment. We also identify two novel optimal policies: the combination of containment and removal to stop spread between patches and the application of up to four distinct policies for a single patch depending on the size of the invasion in that patch.

Chapter 4 (Part B) considers how Grand Coalitions can be stabilized in the face of changing stock location. To do so, we employ the Gordon-Schaefer fisheries model. We consider farsightedness as a mechanism by which stability of the Grand Coalition can be increased in the face of changing stock location. Farsightedness allows players to respond to deviations of other players by deviating themselves. This reduces the incentives to leave the Grand Coalition. This is in contrast to shortsightedness, whereby players cannot decide to leave the Grand Coalition in response to such a choice by another player. We begin by modifying the farsightedness concept such that it can be used in games with asymmetric players and transfer payments. We proceed to analyse the modified farsightedness concept in the case where players are symmetric (stock location does not change) in order to identify the properties of the concept in the base case. We find that farsightedness increases Grand Coalition stability with respect to shortsightedness. We proceed to analyse the extent to which farsightedness increases Grand Coalition stability, relative to shortsightedness, as fish stock location changes, using sensitivity analysis. We find that farsightedness increases the stability of the Grand Coalition, but also increases the sensitivity of stability to changes in fish stock location. Thus, for any fish stock location, a Grand Coalition is more likely to be stable if players are farsighted, but shifts between a stable and an unstable Grand Coalition will occur more frequently if players are farsighted.

In Chapter 5 (Part B), we analyse how the stability of Grand Coalitions is affected by an endogenously determined risk of stock collapse. We do so using the Levhari and Mirman (LM) fisheries model, which is adapted such that there is a risk of stock collapse which increases as the fish stock size decreases. We numerically solve the model and calculate the stability of the Grand Coalition. We find that the effect of an endogenously determined risk of stock collapse depends heavily on the assumptions made regarding how payoffs are determined. A common assumption in the literature is that payoffs are determined at the steady state fish stock. Under this assumption, endogenous risk means that for specific discount and growth rates, a Grand Coalition is stable for any number of players. This is a very different result from the original LM model whereby Grand Coalitions can never be sustained. This is because players can essentially follow two strategies in response to the risk. Firstly, they can attempt to maintain the fish stock by fishing less. In doing so they are running the risk of collapse. Secondly, they can avoid the risk by pre-emptively depleting the fish stock, i.e. harvesting the stock to zero immediately to avoid the risk. Grand Coalitions of any number of players are stable for parameterizations for which a Grand Coalition attempts to maintain a non-zero fish stock and if a deviation from the Grand Coalition would result in pre-emptive depletion. We proceed by relaxing the assumption that payoffs are determined in the steady state by allowing for deviators to obtain payoffs in the transition between steady states. In this case, only Grand Coalitions of two players are stable, and then only for certain parameterizations. The reason is that players can now gain payoffs in the process of pre-emptively depleting the stock, i.e. payoffs are received from the process of fishing the stock down to zero. This increases the benefit of deviating from the Grand Coalition. In this case, Grand Coalitions are only stable for two players for specific parameterizations.

Chapter 6 summarises the research questions formulated in Chapter 1 and evaluates the work of the thesis. Regarding Chapter 2, we justify our theoretical approach with the following two points. Firstly, BWM management is a global and complex problem, which means that the information required to formally calculate an optimal standard is prohibitively burdensome. Secondly, we argue that the complexity of BWM necessitates a sound theoretical understanding of the problem in order to evaluate the current BWM standard, and also to aid in future policy formulation. Similarly, in Chapter 3, we focus on deriving generalized management insights which are applicable to a variety of real-world cases, as opposed to deriving an optimal management strategy for a specific case. In addition to the data requirements necessary to derive such a management strategy, the complexity of such applied cases leads to potentially excessive computational burden. Chapter 3 analyses systems of two and three patches, which are likely to be too simple to analyse specific real world cases, but are sufficient to derive generalized management insights.

The game theoretic methodologies in Part B are evaluated principally in terms of the assumptions about changes in stock location in Chapter 4 and the numerical method in Chapter 5. In Chapter 4, the fish stock is conceptualised as existing at a single point in space. The location of this point is determined in relation to fishing nations, which are also conceptualised as single points in space. Changes in stock location result from rises in ocean temperatures due to climate change. Such rises in temperature are likely to lead to other changes in the fish stock such as the size of the area where the fish stock can be found and increases in the maximum fish stock size which the ecosystem can support. These other aspects of changing stock location need to be considered in evaluating Chapter 4, as well as in formulating more applied models. In Chapter 5, a numerical method is adopted to analyse the effects of an endogenous risk of stock collapse. To do so, the utility function in the

LM model is adapted such that it can be used in a numerical model. In order to isolate the effect of endogenous risk from changes in the utility function, a validation procedure is carried out by comparing analytically derived results in the deterministic case (without endogenous risk of stock collapse) to numerically derived results in the deterministic case. This reveals that changes to the utility function have a negligible effect and thus the results, in terms of the stability of Grand Coalitions can be attributed solely to endogenous risk of stock collapse.

Overall, Part A of this thesis presents new insights into the determinants of optimal BWM standards. These insights demonstrate the conditions under which the current BWM standard, which aims to eliminate the risk of invasion establishment, may or may not be optimal. Part A therefore provides a novel theoretical framework which aids in the evaluation of current, and the determination of future standards. Part A also provides new insights into the control of established invasions, by extending existing spatially explicit optimal control models. Specifically, dividing space into patches and allowing for varying invasive population sizes within patches facilitates the optimal timing of management interventions and, in general, more detailed, and thus more efficient, management strategies. Part B provides a novel analysis of the effects of changing stock location on Grand Coalitions by explicitly introducing fish stock location in the analysis, and shows how farsightedness can stabilize Grand Coalitions in the face of such changes. Part B also shows how the effects of an endogenous risk of stock collapse on the stability of Grand Coalitions depends vitally on whether transition payoffs are included. These results can form the basis for more interdisciplinary analyses, analyses of different types of marine ecological change, and analyses of these changes in different settings, such as non-European countries.

Adam Napier Walker

Wageningen School of Social Sciences (WASS)

Completed Training and Supervision Plan

Name of the learning activity	**Department/Institute**	**Year**	**ECTS***
A) Courses			
Advanced Microeconomics (ECH 32306)	WUR	2011	6
Advanced Macroeconomics (ENR 30806)	WUR	2012	6
Advanced Econometrics (AEP 60306)	WUR	2014	6
Game Theory PhD Course	WASS	2011	1.5
Techniques for writing and presenting a scientific paper	WGS	2012	1.2
Writing PhD proposal	WASS	2012	6
B) Conference Presentations			
"Optimal Control of a Stochastically Spreading Invasive Species	Natural Resource Modelling Conference, Cornell University.	2013	1
"Stability of International Fisheries Agreements under Biological Change"	International Conference 2013: Cooperation or Conflict?, Wageningen University.	2013	1
"Endogenous Risk of Stock Collapse and the Great Fish Pact"	SURED Conference, Swiss Federal Institute for Technology, Zurich.	2014	1
"Spatial Optimal Control of a Stochastically Spreading Invasive Species"	India-EU Workshop on Coastal Zone Management and Impact on Society,	2014	1
Teaching four lectures and supervising 1st year four students	WUR.	2014	3
"The Rise and Fall of the Great Fish Pact under Endogenous Risk of Stock Collapse"	European Association of Environmental and Resource Economists, Helsinki.	2015	1
Total			**34.7**

*One credit according to ECTS is on average equivalent to 28 hours of study load

The research described in this thesis was financially supported by the European Community's Seventh Framework Programme (FP7/2007–2013) under Grant Agreement No. 266445 for the project Vectors of Change in Oceans and Seas Marine Life, Impact on Economic Sectors (VECTORS).